THE UNIVERSITY OF CHICAGO

ESSAYS IN BUSINESS CYCLES, INEQUALITY AND RACIAL DISCRIMINATION

A DISSERTATION SUBMITTED TO
THE FACULTY OF THE UNIVERSITY OF CHICAGO
BOOTH SCHOOL OF BUSINESS
IN CANDIDACY FOR THE DEGREE OF
DOCTOR OF PHILOSOPHY

BY
JUNG SAKONG

CHICAGO, ILLINOIS
JUNE 2019

To Professor Gary Becker,

for encouraging me to go after big ideas and their implications.

CONTENTS

LIST OF FIGURES

LIST OF TABLES

ACKNOWLEDGMENTS

I thank my committee chair and role model in research and in life, Amir Sufi.

I thank my committee members and beacons of genius, generosity and unbreakable patience (though I may have come close), Marianne Bertrand, Raghuram Rajan and Luigi Zingales.

I thank my other advisors and mentors, Lars Hansen, Erik Hurst, Canice Prendergast and Alex Zentefis.

I thank my co-authors and work spouses, Lancelot Henry de Frahan and Paymon Khorrami.

I thank other faculty members, staff members and classmates at the University of Chicago, especially the participants of the Amir Sufi working group and the economic dynamics working group.

I thank my family.

And above all, I thank God.

ABSTRACT

My dissertation studies the interaction between business cycles in the time series and inequality in cross-section, of which racial inequality is one of the most salient dimensions.

The first chapter of my dissertation explores how heterogeneous behaviors over business cycles affect wealth inequality. 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 boomswhen expected returns are lowand sell after a bustwhen 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.

The second chapter of my dissertation finds that racial prejudice, and hence racial discrimination, is countercyclical and may partially account for the higher incidence of business cycles on racial minorities. It starts with the question: Does deteriorating economic condition cause racial prejudice to rise? Despite psychological/sociological microfoundations and multiple economic implications of the inferiority of racial prejudice, empirical evidence has been inconclusive. This paper constructs better-powered measures of local areas' racial prejudice using Google searches for racial slur and "KKK," white-on-black non-pecuniary crime, survey responses and corporal punishment at school. Across these measures, racial

prejudice correlates negatively with the local economic condition. Using predictors of local economic condition in the 2000s, I show that the relationship is causal: lower income causes higher racial prejudice in the area.

Within the context of this dissertation, the third chapter of my dissertation studies the role of racial prejudice and discrimination in determining racial economic inequality, using a natural experiment that shifts local areas' racial bias against black Americans. Following the Obama presidency, pundits and researchers have asked how having a black leader affects white Americans attitude toward black Americans. Given theoretical ambiguity, I test for causal impact of a black leader on racial attitudes using local elections of black politicians at the municipal level. Using Race Implicit Attitude Test (IAT) scores as a measure of racial prejudice and close election regression discontinuity design for causal inference, I find that electoral victory of a black leader leads to a rise in racial prejudice among white Americans against black Americans. Following a close electoral victory, the IAT score rises by about 0.03, or 7% of the average black-white difference. Simultaneously, using the same discontinuity design, black politicians electoral victory causes lower employment and higher mortgage denial for black Americans relative to white Americans. By ruling out other channels by which electoral victory could adversely affect black Americans relative economic outcome, I argue that the rise in prejudice caused black-white economic inequality to widen.

CHAPTER 1

CYCLICAL HOUSING TRANSACTIONS AND WEALTH INEQUALITY

1.1 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] [26] 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 illqiud asset like a house.[3] Partly to encourage wealth accumulation by the middle class, government policies

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. [36]

3. A long literature on present bias and its implication for savings took off starting with [93]. In an earlier

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 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

---

work with co-authors, I found in field experiments that experimentally increasing the illiquidity of a savings account attracted more savings from subjects ([18]).

4. [30]

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 ([74]).

Sorting surnames into percentiles using their historical income from the 1940 Census, I find that 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 A.2a and A.2b.[7]

The overall negative relationship between these betas and the wealth level is driven by differences between racial groups: Non-whites 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

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 ([125]).

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

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.

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 [40].

After conducting these transformations, I find that returns on housing go up 60 basis points per year between the interquartile range of the wealth distribution (Figure A.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

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 ([127]). In another earlier work, I used close electoral victory of black politicians as an instrument for local areas' racial prejudice against blacks, I found that such increase in racial prejudice caused blacks' employment to fall and mortgages to be denied more ([126]). Combining the two results, it is possible that business cycles disproportionately affect racial-minority groups through counter-cyclical racial prejudice. Also see [8]; [7].

9. By comparison, [53] look at total returns on financial wealth in Norway and find about a 1% return differential between the interquartile range per year.

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 ([40]). 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 with historical cyclicality (Figure A.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.

1.1.1 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

10. [99]; [123]

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., [27]).

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 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, [25]. 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 have 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

12. Wealth distribution is wider than income distribution ([20]). To explain why, the literature on wealth inequality finds an important role for return heterogeneity and return loading on wealth ([12]; [59]; [43]).

13. Theoretical papers have used the heterogeneous-agent model to understand wealth inequality ([20]; [43]; [71]; [73]; [119]).

14. Recent empirical work is starting to find evidence that wealth returns are increasing in income and wealth ([5] in Sweden; [53] in Norway; [60] in France; [91] and [143] in the US).

segregation by income and race, using the panel data I construct in this paper.

1.2 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 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.

1.2.1 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*

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

for measured wealth W_{it} of household i in time t. Asset classes are denoted by k with investment share θ^k_{it} and return dR^k_t, 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 [117]). 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 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 future price of housing that leads 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 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

15. [130]

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:*

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

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, [40] 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 [25].[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.

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 [61]). Other papers that document return predictability in the housing market are [116] (using Zillow Research for local house prices), [64] (using Federal Housing Finance Agency data for the price index and Department of Housing and Urban Development data for rent), [29], [42], and [40] for aggregate return.

17. Table 4 in [25] discusses heterogeneity by household financial characteristics.

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 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*

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

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, $\frac{cov\left(\log Q_{it}^{k},\log P_{t}^{k}\right)}{var\left(\log P_{t}^{k}\right)}$.

1.2.2 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 trade is an open, empirical question.

1.2.3 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 Consu-

18. [121]

19. [6]

20. This implication can be seen in [86], where a positive shock to expected house-price appreciation leads to a lower homeownership rate.

mer 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., [25]) 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:

$$\begin{aligned}
\theta_i^K &= \sum_{k \in K} \theta_i^k \\
dR_i^K &= \sum_{k \in K} \frac{\theta_i^k}{\theta_i^K} dR^k
\end{aligned}$$

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 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.

1.3 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.

1.3.1 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.

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 is 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

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.

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 results 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, 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.7c: 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.7c shows that many jurisdictions were added in 1996-1998 and then in 2004.

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

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.

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.7a (for the 1998-2013 sample) and A.7b (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.7c).[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.7d 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 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

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.

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.

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 ([74]). 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 A.1a and also in [74]). 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

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

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.

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 ([74]).

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 ([144]). 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 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).

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.

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 ([74]).

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 B.1.

The first set of tests are for averaged incomes on the CoreLogic side for owner-occupants only (the first two columns of Table B.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 B.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 B.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.

1.3.2 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 A.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]

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

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 A.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 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

percentile groups in Figure A.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.

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 A.1c shows exactly the variation that will be used in the estimation in the next sub-section.

Cyclical variation in house price: Figure A.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 ([40]). 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 [40] deals with the rent series.[34] Similarly for quantities, suppose one

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.

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 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).

1.3.3 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 A.1c on the price-residuals plotted in Figure A.1d, percentile-group by percentile-group.

For the shorter but wider sample covering 1998-2013, Figure A.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 [79].[35] The estimated betas average to zero by construction.

35. To see that I use the same variation as [79], 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, $\Delta \log q_{it} = \beta_i \Delta \log P_t + \Delta\alpha_t + \gamma_i + \Delta\varepsilon_{it}$. [79] 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 [79], are more justified in the daily-frequency setting of [79]. In the annual-frequency setting in this paper, the short sample length makes the distinction between cyclical variation and long-term

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.

The relationship is similar for the longer but narrower sample covering 1988-2013 (Figure A.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 A.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 A.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 ([74]) 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.

---

variation difficult. One solution is to obtain a longer time series, which I plan to do in future research.

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

1.3.4 Betas Between- and Within-Race

To see why this is a decomposition,

$$y_i = \sum_{r \in R} \alpha_i^r + \beta x_i + \varepsilon_i$$

$$\overline{y}_\ell = \sum_{r \in R} \gamma^r S_\ell^r + \beta \overline{x}_\ell + \overline{\varepsilon}_\ell$$

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

For a quick summary, I find that the negative beta-wealth relationship is driven by the between-race variation: Surnames that include more non-whites 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 [11]) as well as wealth and permanent income (e.g., whites may own more wealth than blacks 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 races (e.g., difference between whites and non-whites) 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 A.2a and A.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 A.3b plots the residuals $\tilde{\xi}_i$ from $\hat{\beta}_i = \tilde{\delta}_1 \, (\text{1940 income percentile})_i + \tilde{\nu}_i$ 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 A.3d and A.3f plot the equivalent values for the longer-but-narrower-1988-2013 sample; Figure A.3d is for the full period, and Figure A.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 A.3c, for the shorter but broader 1998-2013 sample. The equivalent plots for the longer-but-narrower-1988-2013 sample are presented in Figures A.3e (for the full period) and A.3g (for 1988-2002).

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

$$\hat{\beta}_i = \delta_0 \, (\text{share white})_i + \delta_1 \, (\text{1940 income percentile})_i + \nu_i$$

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 groups (see Figures A.8a and A.8b for the wealth gaps and income gaps from the Survey of Consumer Finances, for blacks and Hispanics respectively). An enormous literature studies the black-white gap in wealth (e.g., [110]); 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 ([127]). 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 blacks to lose more jobs and face more mortgage denials than whites ([126]). 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.

1.3.5 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.

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 ([74]); 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 ([52]). 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 ([26]).[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

37. While only monotonicity is required, it helps for the second transformation that housing consumption also has an income elasticity close to one ([74]).

estimate

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

where $x_i \equiv E\left[\log \text{ home value — own}\right]_i$ is the average home value among homeowners in the 2012-2013 CoreLogic assessor file (Figure A.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 A.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 ([115]). Figure A.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 A.4b is estimated from

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

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., $(\text{wealth percentile})_i = f^{-1}(x_i)$.[38] This variable will be on the x-axis in the

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 [74].

final return-differential-on-wealth-percentile relationship.

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 the 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.

In this section, I take the characteristics of the expected returns of the national housing market from [40], 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 [40] 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.

Interpreting the Return Differential by Wealth Percentiles

Figure A.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 [101]).

There is only a limited number of estimates of how portfolio returns vary by wealth level. The most comprehensive estimate is from [53], 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.

1.3.6 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 ([96]).

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.

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 ([59] 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,[40] 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 [14]). 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.

40. 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.

To be precise, what matters for the accumulation of measured wealth is the average propensity to consume (APC).[41]

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.[42] 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*

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

See the appendix for the derivation.

41. 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 ([3]). Models in which the rich prefer to save (i.e., wealth is a luxury) would exhibit decreasing APC ([31]). The interaction is insignificant in the Panel Study of Income Dynamics (PSID).

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

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 $\left(1 - c_y\right)$ (i.e., savings rate out of 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.10a), the timing of trades explains roughly a fifth of the residual wealth inequality above and beyond income inequality.

1.4 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 ([116]), and (2) Business cycles are strong predictors of expected returns ([40]).

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

1.4.1 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 ([116]). 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.9c.

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 [124]. Using the [124] housing-supply elasticity instead of the π_c leads to qualitatively the same results.

1.4.2 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 bust 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 A.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.

1.4.3 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.

Measuring CBSA-level Wealth Inequality

I first address an empirical challenge: Local-level measures of wealth inequality do no 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 [99] and [123], 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) ([123])[43]. 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 [99] and [123].[44] I 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 [123] 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.[45]

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.

43. Non-taxable fixed income claims (state/local government bonds) are tax-exempt but reported on individual tax returns since 1987 ([123]). Wealth held by individuals through trusts flows directly to dividends, realized capital gain, interest, and to Schedule E fiduciary income (rents/royalties).

44. [123] 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."

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

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.9a plots the imputed CV in net worth by CBSAs, and Figure A.9b 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 A.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 A.3b column 1) is close to the household-level relationship in Figure A.10a, computed from the Survey of Consumer Finances.

Wealth-inequality Level Results

Figure A.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 A.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 [123]) 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 A.3b reports results from:

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

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.10a and A.10b.[46] 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

46. Using the 2013 Survey of Consumers Finances (SCF), I plot log net worth against log income (Figure A.10a) and log total asset against log income (Figure A.10b). 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.

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).

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.

1.5 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.

CHAPTER 2

RACIAL PREJUDICE IS NOT NORMAL: A COLLAGE OF EMPIRICAL EVIDENCE

2.1 Introduction

Economists at least since [11] have examined how racial prejudice as a preference ("taste for discrimination") can affect economic allocation. At the same time, literatures in psychology, sociology and economics to lesser extent have argued that such social preferences change in response to the surrounding economic condition. [56], for example, argues that pro-social sentiments and virtues increase with economic growth. In this paper, I bring in several measures of racial animus to argue that worse economic conditions cause higher racial prejudice.

A causal linkage from economic condition to racial prejudice has strong implications for economics and political economy. First, it predicts social incohesion and lower tolerance of diversity as necessary consequences of recession or slowdown in economic growth. Second, economic downturn can incentivize politicians to cater to anti-minority rhetoric to capitalize on increased racial prejudice. Third, business cycle downturns have higher direct incidence on racial minorities. Fourth, because downturns are more difficult for minorities, they will sell housing and other assets at a discount, exacerbating existing wealth inequality, possibly perpetuating the racial prejudice itself.

Several theories also predict that worse economic condition should raise racial prejudice, both at the individual and the group level. At the individual level, psychology has a long tradition of theories that predict scapegoating ([46]; [104]; [15]) or in-group bias ([135]; [77]; [76]; [78]). At the group level, sociological theories of ethnic conflict argue that competition over more scarce resources would exacerbate intergroup rivalry, where race can be a salient marker of group ([21]; [111]; [112]).

Yet, despite multiple microfoundations and implications, empirical evidence has been inconclusive. Empirical tests have typically taken one of three forms. First, the oldest literature starting with [81] used time series of lynchings in the U.S. as the measure of racial animus, and tried to test whether periods of economic downturn precipitated more lynchings. Second, other studies use more recent data on hate crimes as the measure of racial animus and further employ cross-sectional variation, starting with [67]. The most recent empirical test employed self-reported survey data and show correlations with economic conditions, both in the time series and across local geographic areas (see [85] for results from the UK). The majority of empirical studies find conflicting evidence depending on the empirical specification, where the lack of robustness may be partly attributed to the fact that both lynchings and hate crimes are rare events, among other econometric issues.

In this paper, I bring in a number of more broadly applicable racial animus measures. Using them, I provide a collage of evidence to establish the causal link from economic condition to racial prejudice.

First, I find evidence that racial prejudice may be countercyclical in the present-day U.S. To measure racial prejudice in this setting, I use responses to the General Social Survey (GSS), Google searches for a racial slur and for the Ku Klux Klan, white-on-black crimes relative to white-on-white crimes in intimidation, simple assault and vandalism (types of crimes most often used to express hate), and frequency of corporal punishment in school for black students relative to white students. I validate my racial prejudice measures against the prevalence of slavery and cotton production in 1860, following [2]. In the cross-section, cities experiencing a larger drop in income saw a greater increase in these racial animus measures.

The negative correlation between income and prejudice can arise due to a natural reverse causation. As in [11], a rising racial prejudice (for whatever reason) will lead to inefficient withdrawal of resources from the discriminated minority and lead to lower output and income. To establish causation from income to prejudice, I bring in several predictors or instruments for the local economic condition. When I instrument for local business cycles using the Bartik labor demand instrument and the Saiz housing supply elasticity ([99]), relationship from income to prejudice becomes stronger. Future draft will also instrument for the economic condition using trade exposure to China ([4]) and for housing demand shock using structural breaks in the house price series ([37]).

Going further back in history, I document that states experiencing bigger income decline saw more creation of Confederate statues, since 1929. Confederate statues have recently been associated with white supremacy, spiking in early Jim Crow era and in the Civil Rights Movement era ([33]). Future draft will also contain county-level analysis during the Great Depression era using retail sales growth as a measure of local economic condition. I will also instrument for Great Depression severity using automobile purchases in the 1920s and bank runs, to establish causation. I will also bring in the historical data on lynching, to conduct finer cross-sectional analysis.

At the individual level, GSS responses to questions on racial stereotypes and animus load on income, even after controlling for fine education categories and other demographic characteristics. Future draft will further exploit questions on conditions at age 16, as well as geographic information, to more closely link the survey responses to income and not some unobservable.

Lastly, I am in the process of expanding the contemporary countercyclicality analysis to an international setting. Future draft will contain between-country analysis using the European Social Survey, and within-country analysis for selected countries using Google searches for racial slurs in the respective languages.

Throughout this paper, "race" encompasses heterogeneity not only in physical appearances, but also in social norms, preferences, income, wealth, skill and many other socioeconomic correlates that are context-dependent. All empirical patterns in this paper speak to all those dimensions of heterogeneity; some of the observed patterns may be purely functions of income differentials, but I do not control for other socioeconomic variables. Instead, I use correlation with observed, time-varying prejudice measure to attribute some of the empirical pattern to racial prejudice. I use "prejudice" to incorporate both cognitive and non-cognitive wedges.

A more precise term may have been "taste for discrimination" as in [11], which "incorporates both prejudice and ignorance." Hence, I include both persistent misinformation about certain racial groups as well as any emotional cost borne when interacting with members of those groups.

2.1.1 Literature: Theory

A long tradition of theories in psychology has argued that worse economic condition should produce more intergroup bias and animus. One of the first in this tradition, frustration-aggression theory argued that when individuals' goals are frustrated, but the source of frustration cannot be overcome, they will refract their anger against a third target, engaging in scapegoating ([46]; [104]; [15]).

Among more recent theories, the social identity theory of [135] states that intergroup bias is a response to individuals' need to establish firm social identity. It has a corollary: depressed or threatened self-esteem motivates intergroup bias ([77]). A similar theory is the subjective uncertainty reduction Theory ([76]; [78]): individuals can reduce subjective uncertainty by "[identifying] with social groups that provide clear normative prescriptions for behavior" and recessions are periods of high uncertainty. These economic condition-related implications of the psychological theories, however, lack empirical validation ([120]), even for the 60-year old frustration-aggression theory. In economics, [131] builds on the social identity theory, and argues that in the presence of behavioral types who have in-group favoritism, even rational types will also engage in discrimination against out-group optimally.

In sociology, ethnic conflicts have been understood as intergroup competition over scarce resources ([21]; [111]; [112]). [57] and [32] are two papers that build off of the sociological microfoundation. In these theories, economic downturns are a time of even more scarce resources, so competition intensifies. Individuals form coalition to secure more resources for their own group members, and race is a natural boundary along which to form groups, because race is salient and observable. In [32], race is a "boundary-enforcement device."

2.1.2 Literature: Empirical Tests

Despite several theoretical microfoundations, empirical tests have produced mixed results. There have been three main strands of empirical tests.

The oldest literature looks at lynchings in time series analysis, starting with [81]. Yet, depending on the time series techniques and data used, studies have reached opposite conclusions. The two key issues are: 1) time series tests are not ideal for uncovering causal relationships, and 2) lynchings are rare events.

A second set of empirical tests looks at hate crimes, starting with [67] in New York City. Other studies look at anti-foreigner crime in Germany ([90]), hate crimes in New York ([67]), race riots in US ([45]), and racial harassment in Britain ([49]). Hate crimes are also rare events, and often not classified as such, because of the added punishment associated with being classified as a hate crime. For example, FBI reports 6,000-10,000 hate crimes per year, but Bureau of Justice Statistics report estimate 260,000 per year.

Last set of tests use self-reported answers to surveys. The closest analogue is [85], who use survey responses from UK to show that survey respondents become more prejudiced in economic downturns. Many survey-based studies suffer from how to define questions about racial prejudice. GSS questions, for example, are often also about preferences for redistribution and ideology ([23]).

2.2 Countercyclicality in Racial Prejudice

I first begin by measuring racial prejudice in a panel. The geographical unit of analysis is a designated market area (DMA) as defined by Nielsen in defining the media market. It is the most disaggregated level at which Google search trends are available over time. Other data have been aggregated up to a DMA. The literature currently lacks panel measures of racial prejudice across the country. I propose several imperfect measures in subsection 2.2.1 and validate them in subsection 2.2.2.

I use these measures to show that racial prejudice rises in bad times, both in the time series in subsection 2.2.3 and for the cross-section of DMAs in subsection 2.2.4. One downside of my measures is that many are only available for the recent years, and as a result, the Great Recession is the only recession in the time series. The beginning of the Great Recession overlaps with the election of Barack Obama for U.S. president, which some argued has influenced the level of racial prejudice in the country, although both directions of influence have been argued. The Obama election cannot be the entirety of the story, because areas with more severe recessions saw a bigger increase in prejudice.

2.2.1 Measuring Racial Prejudice

I construct proxies for racial prejudice that are specific to the time and place using multiple data sources. In addition to widely used survey responses from the General Social Survey (GSS),[1] I use Google search trends for keywords associated with racial prejudice, crimes committed by white offenders against black victims, and school corporal punishment on black students relative to white students. These proxies are associated with other things in addition to racial prejudice; for example, the racial slur that I use for Google search trends also often features in the popular rap culture. The goal is to combine proxies that are related to racial prejudice albeit with much measurement error, so that the common component can be attributed to the racial prejudice of the time and place. I will describe each of the measures in turn.

The main prejudice measures are obtained from Google search trends for a prominent racial slur (the n-word) and for "KKK." Google shares an index of search volume at the DMA-level since 2004 at the monthly frequency. Google searches for the the n-word has been shown to predict voting against Obama better than the GSS ([132]). Google search trends have several advantages over survey responses. First, the search index is available

1. "Some of the data used in this analysis are derived from Sensitive Data Files of the GSS, obtained under special contractual arrangements designed to protect the anonymity of respondents. These data are not available from the authors. Persons interested in obtaining GSS Sensitive Data Files should contact the GSS at GSS@NORC.org"

for a geographical cross-section. While sensitive GSS data do have geographical identifiers, sample size considerations prevent local GSS measures from being too precise. Second, Googling can be done in secret and hence the anonymity suggests that Google searches can reveal racial prejudice more directly than survey responses, where the respondents may feel social pressure to give politically correct answers.

Google search trends also have some issues. First, there are other contexts in which the n-word and “KKK” can be searched for without involving prejudice. For example, the n-word sometimes appears in rap music, although more often rap usage uses the variant of the word that ends in an “a.” Fortunately, Google Trends offers most common related searches, which suggest that many uses of these search terms do involve racial prejudice. Second, for privacy reasons, only keywords with sufficient search volume can be tracked on Google Trends. Other ethnic slurs whose usage is less ambiguous are also less widely used and cannot be tracked consistently over time. Third, the time series for the search trends cannot be taken meaningfully. The type of Google users has changed dramatically over the past decade as well as the fads that dominate Googling. Since Google Trends gives me an index, how it changes over time is influenced as much by what else is searched by whom. Cross-sectional comparisons and diff-in-diff comparisons are meaningful and form the backbone of the analysis.

Figure B.1a shows the cross-sectional distribution of search frequency in 2015 for the n-word. The DMAs with the highest average searches over the 2004-2017 period are Parkersburg, WV, Twin Falls, ID and Greenwood-Greenville, MS. DMAs with the lowest average over the entire period are Macon, GA, Portland, OR and Salt Lake City, UT.

Using Google search trends for “KKK” is new. The Klan is still in operation, and this search proxies for multiple channels related to racial prejudice. Individuals interested in joining it would search for it, and for Google Chrome users, entering the first few characters of the Klan website and pressing enter pre-maturely would lead a user to search for that term on Google. Individuals who fear or are concerned about extremist groups may also search for the Klan as a notorious example, and such searches would also be related to local racial animus.

Figure B.1b shows the cross-sectional distribution of search frequency in 2015 for "KKK." The DMAs with the highest average searches over the 2004-2017 period are Presque Isle, ME, Greenwood-Greenville, MS and Parkersburg, WV. DMAs with the lowest average over the entire period are New York, NY, Washington DC and San Francisco-Oakland-San Jose, CA.

The next measure captures more extreme expressions of racial animus. I measure anti-black sentiments using crimes committed by white offenders against black victims, relative to those committed by white offenders against white victims. The data come from the National Archive of Criminal Justice Data (NACJD). Among crime categories, I take three codes only: simple assault, intimidation and destruction/damage/vandalism of property (the uniform common reporting codes are 132, 133 and 290). These crimes are non-pecuniary in nature and therefore less likely to be linked to economic downturns for directly financial reasons. They are also the most common categories of hate crime. In fact, more than 10% of crimes in these categories are hate crimes. I do not use hate crimes directly because they are too few in number and the categorization into hate crime is subjective and potentially influenced by the fact that hate crimes are sometimes punished more severely. Finally, I scale the crime ratio by black-to-white population ratio in the DMA; otherwise, areas with higher black population would mechanically see more crimes committed against blacks. The exact variable definition is:

$$\left[\frac{\text{count of crime by white offender against black victim}_{it}}{\text{count of crime by white offender against white victim}_{it}} \Big/ \frac{\text{total blacks in area}_{it}}{\text{total whites in area}_{it}}\right]$$

for DMA i in year t.

Figure B.3a shows the cross-sectional distribution of white-on-black crime relative to white-on-white crime in 2014. The DMAs with the highest average over the 1991-2014 period are Bangor, ME, Alpena, MI and Sioux City, IA. DMAs with the lowest average over the entire period are Wilkes Barre-Scranton, PA, Tucson, AZ and Jackson, MS.

The last measure is the frequency with which school teachers corporally punish (i.e. spank) students of other races relative to white students. Corporal punishment is still legal in 19 U.S. states, and relevant statistics are released by the U.S. Department of Education as a part of the Civil Rights Data Collection. In some sense, this measure best captures changes in racial prejudice. Decision to spank a student may be impulsive and thus reflective of the underlying prejudice, and it is less likely to be affected by economic conditions directly (cf. implicit-association tests). While a downturn can cause teachers to spank more overall, it is more difficult to think of why they would spank black students more than white students. When dealing with corporal punishment, I use ratio of rates of punishment, as there are large differences in the base rates. Black students are roughly twice as likely as white students to be spanked, whereas Asian students have much lower likelihood.

Crimes and spanking are less direct measures of racial prejudice, less like surveys and more like mortgage denial and non-employment. Conceptually, I imagine that there is a latent average racial prejudice associated with each DMA-year. The prejudice measure I compile are a function of the latent racial prejudice and other factors that I claim are otherwise orthogonal to the local economic conditions and household finance outcomes of interest. As I accumulate more of these measures, my proxies will converge onto the latent racial prejudice.

2.2.2 Validation of Racial Prejudice Measures

As I propose new proxies to measure racial prejudice, this section seeks to validate them against existing measures. First, I compare my measures against the relevant GSS survey responses, which are still the most widely used measures at levels of aggregation higher than the individual. Second, I compare my measures to the cross-sectional prejudice instrument proposed in [2], who use slave population share and cotton production in 1860. They argue that a legacy of slavery passes down over generations, and show that the prevalence of slavery among southern counties predicts survey responses today. Lastly, I validate my measures by showing that they capture something in common, i.e. that a few principal components can explain much of the variation.

I replicate the prejudice index constructed by [34], in addition to another index composed of questions unambiguously measuring racial prejudice.

[2] use as the instruments 1860 slave share and cotton production conditional on state fixed effects. For example, this contrasts Augusta, GA against Atlanta, GA. A strong correlate of this instrument is the contemporary black population share. This correlate is problematic for their interpretation because of an alternative hypothesis: the "racial threat" hypothesis argues that it is the current prevalence of black Americans that raises anti-black prejudice. This correlate is problematic for me, because more black Americans may proxy for prevalence of lower-skilled workers, hence raising employment cyclicality, for example. Their first solution is to simply control for current black population share. This makes instrument potentially not valid (i.e. over-control), but their coefficients are not affected much. In my validations, I also control for the contemporary black population share.

Table B.1d regresses the DMA's average prejudice measures (obtained as DMA fixed effects) against the area's 1860 slave share of the population from Census 1860, accessed via IPUMS. All regressions include state fixed effects following [2] as well as contemporary black population control. First three columns show that my three average prejudice measures are positively related to slave share in 1860. The next three columns show that areas with higher prevalence of slavery have higher black population share even today and tend to be more rural. Again following [2], Table B.1b regresses my average prejudice measures against slave share instrumented using per capita cotton production in 1860.

2.2.3 Time Series Evidence

I first present time series evidence where possible. First, using the GSS, I find that racial prejudice picked up in the U.S. during the Great Recession. Following [34], I standardize questions in the GSS relevant to racial attitudes and sum over them to create an index of prejudice against blacks. The [34] measure is plotted in Figure B.5a, for all respondents (blue line) and for white respondents only (red line). They show that the secular decline in racial prejudice rebounds starting in 2006-2008 and stays high. The original [34] measure aggregates across different sets of questions asked in different years. For robustness, I take only the set of questions asked in years 1996-2016, and repeat the exercise for all respondents (green line) and for white respondents (orange line). The rebound in the Great Recession is again visible. The consistently asked questions in this recent index are:

Do you oppose a preference in hiring and promotion? (affirmative action)

In general, how close do you feel to blacks?

Agree? The government is obligated to help blacks.

Would you vote for a law that says a homeowner can refuse to sell to blacks, or one that says homeowners cannot refuse to sell based on skin color?

Would you object to sending your kids to a school that had few/half/most black students?

Agree? Italians, Jews, and other minorities overcame prejudice and worked their way up. Blacks should do the same without special favors.

The magnitude of the rebound in the Great Recession is comparable to the cross-sectional difference in the measures in 1996 between mid-Atlantic (NJ, NY and PA) and East South Central (AL, KY, MS and TN) census divisions.

Unlike Google search trends, white-on-black crime and corporal punishment measures do have time series interpretations. Figure B.3b plots the year fixed effects from regressing the white-on-black crime to white-on-white crime ratio on year fixed effects and DMA fixed effects. I report year fixed effects instead of averages as the crime panel is highly unbalanced; NACJD's coverage of crimes increases over time. First thing to notice is an upward trend, which may reflect several forces, from increasing non-employment of white men to the secular decline in residential segregation. More relevant for the present paper, the deviations from the linear trend show some increases in both the 2001 recession and the Great Recession.

Next, Figure B.4a plots the ratio of corporal punishment hazard for students of selected minority races relative to white students. Each series is plotted on a separate y-axis, because the base hazard rate ratios range from a tenth for Asian students to twice as much for black students. While it is difficult to infer dynamics from data collected every 2-3 years, punishment hazard for all three minority groups rose in 2009.

2.2.4 Cross-sectional Evidence

As mentioned above, analysis based on the Great Recession in the time series is confounded by the Obama election. Therefore, for the main analysis, I turn to cross-sectional relationships. For Google search trends for the n-word and "KKK" and white-on-black crime, I run the following empirical specification:

$$y_{it} = \alpha_i + \alpha_t + \beta \log \left(\text{per capita income}\right)_{it} + \varepsilon_{it}$$

for DMA i and year t. I use per capita income from the Bureau of Economic Analysis (BEA) as my measure of local business cycles.

Figures B.2a, B.2b and B.3c show diff-in-diff bin scatters of the three prejudice measures against the DMA's log per capita income. Panel observations are grouped into 100 quantiles by residuals of per capita income after taking out DMA fixed effects and year fixed effects. Each dot represents the average of the group for the two variables, residualized by the two levels of fixed effects. Observations are weighted by the DMA's average number of households. Averages of prejudice measure and per capita income have been added back in generating the plots.

All three plots show that declining local income is associated with more searches of n-word and "KKK" as well as more crimes committed by whites against blacks. The coefficient estimates can be found in the OLS columns of Table B.2a. Table B.2b repeats the empirical specification, but only for the two years 2006 and 2009, using only the cross-section of changes over the Great Recession. Results show that areas hit harder by the Great Recession saw a greater increase in the prejudice measures.

The negative relationship is causal. To rule out confounding explanations such as general optimism, I instrument for local business cycles using a Bartik labor demand instrument generated from worker shares in 2-digit NAICS industries and the Saiz housing supply elasticity multiplied by the national house price index. The former is a widely used instrument for local business cycles, while the latter is particularly relevant for the housing boom and bust, following [100] and [99]. The first columns in Tables B.2a and B.2b show strong first stages. The "IV" columns show the causal effect of local income on racial prejudice. Income decrease still causes prejudice measures to rise, although magnitudes are significantly larger than the OLS counterparts.

There are a few possibilities for why the IV estimates are larger. First is simple measurement error: per capita income is not the most precise measure of local business cycle. Yet the discrepancy seems too large to be justified by measurement error of this sort. Second and more likely, the business cycle instruments are endogenous to racial prejudice proxies. In particular, Tables B.2a and B.2b report the Hansen's J statistic from over-identification tests and the associated p-values. For the Google search trend measures, Hansen's J suggests that my two business cycle instruments cannot jointly satisfy the exogeneity criterion. A possible culprit may be spatial migration patterns.

Corporal punishment rates are available only at the state-level. Figure B.4b plots the change in black students' corporal punishment rate to white states' from 2006 to 2009 against the log income change from 2006 to 2009. Each dot is a state; there are only a handful of observations because most states do not see corporal punishment in both years. Even with the few dots, it seems that states with bigger drop in income saw relatively more blacks students spanked.

Lastly, I use GSS to test for countercyclicality. Given that the public GSS only identifies the nine census divisions, here I use year-question-census division level panel. Furthermore, I restrict the analysis to a set of questions that are unambiguously suggestive of racial prejudice. For example, the time series plot included questions about affirmative action that may or may not be about racial prejudice and more about political ideology. The set of questions used in this panel test are:

How strongly would you object if a family member brought a black friend home for dinner?

Do you think blacks should have as good a chance as anyone to get any kind of job, or do you think white people should have the first chance at any kind of job?

Do you think there should be laws against marriages between blacks and whites?

How would it make you feel if a close relative of yours were planning to marry a black?

If a black with the same income and education as you have moved in to your block, would it make any difference to you?

Would you vote for a law that says a homeowner can refuse to sell to blacks, or one that says homeowners cannot refuse to sell based on skin color?

If your party nominated a black for president, would you vote for him if he were qualified for the job?

The answers are converted to a binary variable with 1 indicating racial prejudice. These questions are not asked in every year. The empirical specification is:

$$y_{iqt} = \alpha_{iq} + \alpha_t + \beta \log\left(\text{per capita income}\right)_{it} + \varepsilon_{iqt}$$

for census division i, question q and year t. The result is plotted in a bin scatter in Figure B.5b for 100 bins. Again, when income falls, GSS respondents give more prejudiced answers.

The cross-sectional evidence from Google search trends for the n-word and "KKK," white-on-black crime, corporal punishment and GSS survey responses all suggest that racial prejudice rises when income falls. The rest of the paper studies what this phenomenon implies about households' employment, credit and asset holdings.

2.3 Countercyclicality in Economic Gaps

Housing sales is a channel through which business cycle incidence translates to wealth inequality. It is well known that blacks lose jobs more in down times and that they are less likely to be able to get credit. I replicate those findings in the last two subsections. Job loss and tighter credit can lead black households to disproportionately sell assets. To study this, I match housing deed records from DataQuick to the owners' races using their last names. I find that minority households were more likely to decumulate housing in economic downturns, in the time series as well as in the cross-section. These findings are presented in the first two subsections.

2.3.1 Housing Sale: Data

I identify race of owners in deed records from DataQuick by merging the last names of individual buyers and sellers to the Census tabulation of surnames in 2000 and 2010. Recent waves of the Decennial Census include tabulations of surnames with more than 100 people, along with their composition by major race groups ([144]). 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 a 85.8% coverage.

On the housing data side, DataQuick compiles all deed transfer records, including the date, amount and nature of the transaction, and the names of the buyers and sellers. Using geographic jurisdictions that have consistent coverage during the sample period of 1998-2013, the data cover all transactions for properties in areas covering more than 60% of the addresses in U.S. The geographic areas covered by DataQuick are more likely to be urban and are not representative of the U.S. as a whole.

Using last names to identify race has the disadvantage of measurement error when it comes to surnames with multiple races associated with them, such as Williams or Jackson. Existing studies of race in the housing market use either survey data or deed records matched to HMDA loan applications, which contain a race flag. Relative to survey data, my approach here has a power advantage due to the administrative nature of the data. Identifying race from HMDA race flag is useful for studying the gross purchases of housing assets, but not the sales or stock of holdings. Using surnames allows me to construct high-frequency snapshots of real estate holdings and both purchase and sales. Wealth accumulation via real estate is a function of three factors: whether to purchase, what to buy (i.e. return), and when to sell if ever. Looking at gross purchases, gross sales and segregation (one dimension of which assets they buy) allows me to study all three factors.

Last names are more reliable identifiers for Asian and Hispanic households, and less so for black and white households. Some common, predominantly white surnames include Wood (247k), Olson (164k) and Snyder (159k), with associated Census 2000 population counts in parentheses. Common, predominantly black surnames include Smalls (11k), Batiste (6k) and Gadson (5k). Common, predominantly Hispanic surnames include Garcia (858k), Rodriguez (804k) and Martinez (775k). Common, predominantly Asian surnames include Nguyen (310k), Kim (194k) and Patel (145k). Common, predominantly American Indian surnames include Begay (16k), Yazzie (14k) and Benally (5k). Thus far I used 90% as definition of predominant. Using 80% as the definition, common black surnames are Washington (163k), Alston (28k) and Ruffin (15k).

One data issue is that while I know the total population by last name for the nation as a whole, I do not know which families live in which areas. When I construct regional measures, I have to assume that each surname is as likely to be anywhere there are racial groups that are affiliated with the surname. For example, if I only look at housing owned by black households as defined by surnames that are more than 80% black, I will omit the housing owned by black households with the surname Jefferson but will capture those with the surname Smalls. In theory I have to scale local holdings by the actual population that those surnames represent. In practice I will have to use actual black population of the area and scale by the national average of what fraction of blacks are covered by the last names I use. I am in essence assuming that whether a given black American owner has surname

Jefferson or Smalls is orthogonal to everything I do.

Using yearly holdings and transaction data for 1998-2013, I study the sales dynamics by race, both gross and net, in the next subsection. In the appendix, I study the residential segregation dynamics.

2.3.2 Housing Sale: Countercyclicality

Figures B.7a and B.7b show the per capita stock of owner-occupied housing and other real estate holdings, for four major race groups. Each group contains population with surnames for which at least 80% belong to that race. I linearly intrapolate population by surname using the 2000 and 2010 Census data on last names and use the population estimates to scale the real estate holdings to arrive at the per capita holdings. To interpret magnitude, take white primary residence holding at .1 per capita. Multiplying by the average white household size takes it to .3 primary residence per household. Since my DataQuick data cover 60% of addresses, assuming representativeness takes it up to .5 per household. The remaining gap to 70-75% homeownership rate for white Americans is because the DataQuick sample oversamples urban areas where homeownership rate is lower, and matching is imperfect.

Looking at the stock of housing, black and Hispanic households accumulate housing until 2007, then their holdings begin to drop. The cyclicality stands out more when looking at net changes, as in Figures B.7c and B.7d. Black and Hispanic households' housing accumulation dips sharply in 2007-2008 and does not recover for many years. Given that these were the onset of the Great Recession and that returns to holding housing were the highest, some of the decline is housing accumulation may be due to worsening economic condition for these racial minority households.

This black/Hispanic-white divergence is not limited to the surnames with one predominantly associated race. Figure B.8 plots the percentage change in real estate holdings for 2003-2006 (blue) and 2007-2010 (red), against what fraction of the population with the given surname is black or Hispanic. Relative to either quintile, which were plotted in Figure B.7, surnames with both races lie monotonically in the middle.

Going back to Figures B.7c and B.7d, Asian households start slowing down accumulation earlier in 2006, and their pace of acquisition does not really slow down much into the recession. Part of this pattern is due to inflow of foreign Chinese buyers ([109]).

Figures B.9c and B.9d plot gross sales as a fraction of pre-existing holdings, reflecting a propensity to sell. Relative to the average turnover, propensity to sell rises for black and Hispanic households, in particular for Hispanic households' primary residence. The rise in propensity to sell for black households is not immediately evident because they tend to have lower average propensity to sell and turnover is highly procyclical. But relative to white and Asian households, black households' propensity to sell does rise in the recession years.

On the other side of the transaction, Figures B.9a and B.9b plot gross purchases per capita, again for primary residence and other real estate. These plots show that gross purchases are procyclical, and more so for black and Hispanic households. Figure B.9b again shows a rise in secondary real estate by Asian households, some of which reflect the inflow of foreign Chinese buyers. These numbers can be compared to transactions matched to race by using HMDA data.

In the cross-section, rising income is associated with more housing accumulation by black households. Figures B.10a and B.10b plot bin scatters of the net change in housing owned by black households (scaled by total number of households in the area) against log per capita income of the area, for owner-occupied housing and for non-owner-occupied. Both owner-occupied and secondary housing owned by black households respond positively to local income.

2.3.3 Employment Gap

I document that blacks lose jobs more in recessions, in the time series and in the cross-section. The aggregate pattern is well known: blacks work in lower-income jobs and lower-income jobs have more cyclical income relative to most of the income distribution except the very top ([19], [72], [85]). In this section, I plot black-white gap without correcting for non-race factors. In this section, any black-white gap documented partially reflects income/skill differences and may have nothing directly to do with race. I just display the raw differences here.

More of the labor market discrimination literature focus on wage differential as opposed to employment differential ([94]). Given that my interests are cyclical and wages may be sticky over cycles, here I focus on employment differential. [94] document falling racial prejudice reported on surveys ("The survey results suggest that strong prejudice is an increasingly peripheral explanation for racial inequalities in the labor market"), but more modern tests such as implicit-association tests still suggest a role for prejudice.

Data on employment gap come from Quarterly Workforce Indicators (QWI) and refer to all private employment at the county-level, aggregated to the DMA. Figure B.11a plots the fraction of workers transitioning to non-employment in a given year, for black workers (blue) and white workers (red). The green dots plot the difference. The transition to non-employment for both groups as well as the difference all exhibit strong countercyclical pattern.

Figure B.11b plots the gap again (blue), in addition to the within-2-digit-NAICS transition to non-employment gap (red). That is, the red line takes the within-industry transition rate for each race and aggregates them using common, total employment shares by industry. That the two lines coincide implies that 2-digit industry cannot explain the countercyclical gap: blacks are not disproportionately hired in cyclical sectors. Figure B.12a plots the industry shares by race in 2004. The bars add up to 1 for each race group. Blacks (green bars) are hired in some cyclical industries, such as accommodation and food services (72), but they are less represented in some of the most cyclical industries such as construction (23) and finance and insurance (52).

The countercyclicality of black-white employment gap is true also in the cross-section. Figure B.12b plots a bin scatter of the black-white gap in transition to nonemployment against log income per capita of a DMA, taking out year fixed effects and DMA fixed effects. A ten percentage point decrease in per capita income leads to a 0.8 percentage point increase in the black-white gap in fraction of workers transitioning to non-employment.

2.3.4 Mortgage Origination Gap

With more cyclical employment, black households may demand more credit in downturns. However, mortgage approval and origination are also more cyclical for black households relative to white households ([142]). The contraction in mortgage credit supply is partly in response to higher unemployment among black households. The contraction has other clear non-racial components: black households have lower wealth and can be more vulnerable to credit rationing, which rises in economic downturns.

Data on mortgage denial come from the Home Mortgage Disclosure Act (HMDA) data. I calculate rejection rate as the number of mortgage applications for owner-occupied housing that were rejected, divided by all loans that were originated or denied. Figure B.13a plots the raw rejection rate for black (blue) and white (red) households, with the difference in green dots. Controlling for income year-by-year (dashed lines in Figure B.13b) shifts level of rejection rate gap, but the countercyclicality remains for all racial minorities.

Since rejection rate is affected by who applies, I also look at mortgage origination per capita. Figure B.14a plots the approval rate gap over time by number of applications (blue). It also plots the black-white gap in total dollar value of mortgages originated for owner-occupied housing per capita. Rejection rate gap rose from around 15% to over 20% in the Great Recession. Spikes in earlier business cycle downturns are also there. From the timing, it appears that a drop in approval rates precedes a drop in the amount originated for black households, possibly because black households begin to apply less. Throughout this paper, the main credit measure will be the gap in per capita mortgage origination for owner-occupied housing.

The cyclicality of mortgage credit also holds in the cross-section. Figure B.14b plots a bin scatter of the black-white gap in mortgage origination against log income per capita of a DMA, taking out year fixed effects and DMA fixed effects. A ten percentage point decrease in per capita income leads to a 5.8 percent decrease in the origination gap.

2.4 Conclusion

Using various measures of racial prejudice in an area, I find that economic downturns reverse racial tolerance. While intuitive, countercyclical prejudice has far-reaching consequences in several areas of economics.

CHAPTER 3

THE OBAMA EFFECT: EFFECT OF BLACK ELECTORAL VICTORY ON RACIAL PREJUDICE AND INEQUALITY

3.1 Introduction

The Barack Obama presidency has motivated dicussions as to how having a black leader affects white Americans' attitude toward black Americans in general. Empirical analysis of black leaders has been challenging due to two issues: 1) relative to population, black politicians are disproportionately few in numbers, and 2) black politicians are not elected randomly. This paper seeks to overcome those issues by looking at a much broader set of local elections of black politicians, and – aided by the larger sample size – by using close elections to establish causality from electing black leaders to racial attitudes and actual racial economic inequality.

Black electoral victory can affect whites' racial attitudes through multiple mechanisms, with differing direction of effect. First, a black politician could dispel negative stereotypes associated with black Americans. Second, having elected a black politician, white Americans can feel justified in holding and expressing racial bias via the self-licensing effect ([105]). In particular, [50] found that endorsing Obama led participants to express views that favored whites over blacks. The idea that Obama election has led to a resurgence in racial prejudice has been circulated in popular media as well.[1]

In this paper, I estimate a causal impact of a black politician's electoral victory on white Americans' racial prejudice and black Americans' relative economic outcome, using local elections and a close election regression discontinuity design.

Data on local elections come from Our Campaigns, a Wikipedia-like website compiling

1. For example, see Blake, John (2016, November 19) "This is what 'whitelash' looks like." CNN. https://www.cnn.com/2016/11/11/us/obama-trump-white-backlash/index.html; "A reflection on Barack Obama's presidency." (2016, December 24) The Economist. https://www.economist.com/christmas-specials/2016/12/24/a-reflection-on-barack-obamas-presidency.

electoral information. I use data on any US election with sub-federal constituency. This includes US federal House representatives, mayors, city council members, county executives and county council members. Given the disproportionately low prevalence of black politicians, extending the set of politicians beyond the commonly considered representatives and mayors drastically increases the sample size. The race of the candidates is classified using: 1) tags supplied on Our Campaigns, 2) candidates' surnames and corresponding racial distribution from the Census, and 3) facial recognition of the candidates' photos.

In order to measure racial prejudice in observation data, I use Race IAT score data from Project Implicit Database, compiled by [145]. Since its development in [68], the IAT has been a widely used test of subconscious or implicit racial bias. Project Implicit Database collects voluntarily completed online IAT surveys, along with demographic and other respondent information. Since the online IAT surveys are voluntary, the pool of respondents is both self-selected and highly unrepresentative of the US population. I project out the demographic information to get at variation in local IAT scores not driven by composition. The residual scores are then aggregated to the county level. Given about quarter million completed surveys per year, the resulting panel data have informative local variation. Finally, the results are unchanged when using the raw scores or composition-adjusted scores, suggesting that selection is not affecting the identification strategy.

The empirical strategy starts with a standard difference-in-difference estimator, comparing the jurisdiction of black winners with surrounding areas. The difference-in-difference estimate shows that there was no differential change in racial prejudice in areas affected by the election relative to surrounding areas. However, the difference-in-difference is likely biased, for example because black politicians are more likely to be elected in areas where prejudice against blacks is falling.

To overcome this identification challenge, I use a standard close election regression discontinuity, looking at the 3-year period after an election. The 3-year period before an election is used for placebo tests. I only look at elections where the top two candidates include one black and one white candidate. Results in this draft are obtained by defining "close election" as those with less than 10% vote margin between the winner and the runner-up. Optimal bandwidth (for example using [83]) is wider, and results obtained using optimal bandwidth are not qualitatively different.

I find that a black politician's victory causes racial prejudice to rise as measured by the IAT score. The average effect size is around 0.03, which is about 7% of the raw average gap in IAT scores between white and black Americans, where higher IAT score indicates more racial bias. There is no difference in IAT scores in areas with a black winner versus a white winner, before the election.

Turning to economic inequality measures using the same regression discontinuity design, black politicians' electoral victory causes black workers to transition more into non-employment (flow), fewer blacks to be in employment (stock), originate less mortgage amount and be denied more mortgages, relative to white counterparts.

In the last part of the paper, I interpret close election of a black politician as a time-varying instrumental variable for local racial prejudice. I highlight how this approach of identifying discrimination is different from what has been done in the literature so far.

3.1.1 Literature Review

The election of Barack Obama in 2008 has spurred an explosion of literature looking at the two-way interaction between black representation in politics and attitudes toward black Americans (see [114] for a comprehensive review). Academic study of the causal direction from salient black political leaders to racial attitudes is smaller but growing fast, including in economics (for example, [44]). There are multiple theoretical channels by which visible figures of a minority group can affect views toward that group in general, and theories differ on even the direction of the effect.

On one hand, seeing a black exemplar may reduce bias toward blacks by dispelling negative stereotypes and showing that black Americans are capable. In contemporary American, black males are stereotyped as criminals and black females as undeserving recipients of welfare ([82]). Seeing a black American rise to a position of prominence and competence may dispel such stereotype. In studying gender attitudes in India and female politicians in India, the authors find that randomized rise of female politicians do indeed reduce gender bias and lead to positive outcomes for girls ([9]; [10]; and [38]).

On the other hand, seeing a black leader may also increase bias toward blacks by convincing majority whites that they need no longer watch out for racism or by raising the threat they perceive as the dominant group. [105] argues that one act of racial tolerance endows the individual with a moral license to express biased views in other dimensions, exhibiting the so-called self-licensing effect. [50] finds that endorsing Barack Obama during the 2008 presidential election made his supporters more likely to display racial bias against black Americans. [89] further finds that even knowing of others' un-prejudiced behavior leads individuals to freely express biased views. Theories of intergroup bias based on social dominance argue that black politicians may raise the perceived threat to white Americans' racial dominance and motivate them to engage in backlash ([111]; [129]).

Against this backdrop of ambiguous theoretical prediction, empirical studies examining racial attitudes following the election of Barack Obama find results in both directions, depending on sample and methodology (for example, [65][2]; [88]; [137] ;[136]). A key challenge is that the Obama election is one event, limiting the statistical power of empirical tests. The innovation of this paper is to extend the set of elections to local and municipal levels, where there are more black electoral victory to potentially establish statistical power. The key tradeoff is that local elections are less salient.

At the mayoral level, there is a large political science literature examining the impact of black mayors, although the focus is predominantly on actual policy impact. The key question is whether active political participation matters for voicing the preferences of minority

2. Using a survey of 20,000 respondents, they find, "From the summer of 2008 through Obama's inauguration in 2009, there was a gradual but clear trend toward lower levels of white prejudice against blacks." "... this change in attitudes did not last."

constituents, as opposed to passive participation in the form of voting. These studies often seek to test whether black mayors enact more liberal policies ([87] ; [1]; [128]; [98]; [80]; [108]; [51]; [22]). Most recent studies find that the racial identity of mayors has no direct policy impact due to constraints that municipal executives face, except in a few policing practices. Studies looking at party or gender find similar results ([55]; [54]). These studies suggest that if election of black politicians has economic impact, it is less likely to be through actual public policy changes. And that elections with black candidates receive more attention as measured by higher white and black voter turnout suggests a possible psychological channel if black electoral victory affects economic outcomes of black and white constituents ([140]).

3.2 Data

This study draws from multiple sources of data: 1) data on political elections and candidates from Our Campaigns, 2) data on local measures of racial prejudice from Project Implicit (Implicit Association Test scores), Google searches and interracial crimes, and 3) data on interracial economic disparity from Quarterly Workforce Indicators (for employment) and Home Mortgage Disclosure Act (for mortgage origination). The following sub-sections describe each in detail. To make sense of the various data components, it is helpful to think of political elections and candidates as data on the forcing variable or instruments (i.e. Z), racial prejudice measures as data on the independent variable of interest or the outcome of the first-stage (i.e. X), and interracial economic disparity measures as data on the ultimate or second-stage outcome variable (i.e. Y).

3.2.1 Election Data

Political elections provide the ideal setting for studying the effect of a salient exemplar for both substantive and methodological reasons. Substantively, politicians are highly visible, partly because they have to campaign to attract votes. Methodologically, election outcomes are uncertain and provide an identification strategy to study the causal effect of having an exemplar.

Data on local electoral outcomes come from Our Campaigns, a Wikipedia-like website that aggregates electoral information. I use data on any US election with sub-state constituency. This includes US federal House representatives, mayors, city council members, county executives and county council members. Given the disproportionately low prevalence of black politicians, extending the set of politicians beyond the commonly considered representatives and mayors drastically increases the sample size. Table C.1a shows the in the years since 2003, there were 202 black congressional victories in US House, 12 black gubernatorial victories and 215 black mayoral victories in the Our Campaigns data using the racial identification scheme to be described below. If I wanted to use close elections between black and white candidates, defining close election as those with a vote margin of 10% or less between the top two candidates, that leaves 68 elections and 39 black victories among the three commonly studied, most visible offices. Including state legislatures, county councils and city councils, along with other elected municipal offices (e.g. county president) increases these numbers six- to sevenfold, with 501 close black-white elections and 247 black victories in such elections. Of course, such sample size gain is traded off against the lower visibility and salience of these local offices in determining statistical power. The tradeoff will be explored and exploited in the empirical specifications below.

To compile the election data used, I scrape the Our Campaigns website for the following information: most recent map of the jurisdiction (map in Figure C.1a and the associated longitude-latitude coordinates), history of electoral races with date, type, candidates and their vote counts and shares (Figure C.1b), and candidate information, in particular the user-supplied tags for race and photos for facial identification of candidate race (Figure C.1c). Since I am expanding the set of politicians studied beyond what is commonly done in the literature, identifying candidates' races is a challenge, which I address below.

Identifying the Race of a Candidate

Candidates' race information is provided for a small subset of those on the website, so I use two other methods to identify candidates' races. Whenever the website contains direct information on race, I use that information. This small subset is also useful for judging the accuracy of the other two methods. Table C.1b shows how the candidates classified according to tags are classified using the other two methods of surnames and facial recognition. Note that the user-supplied tags can be wrong too.

The second source of racial identification is the candidates' last name. Surnames are a widely used source of demographic information, even outside of the US (for example, [106] uses surnames in India to identify religious and caste identity). Using the Census 2000 and 2010 surnames files ([144]; [41]), I classify a candidate as belonging to one racial group, if 80% or more of people with that candidates' last name belong to that race. For example, candidates with the surnames Little or Smalls are classified as being black, while candidates with surname Hansen are classified as being white. Surnames are highly informative for Hispanic and Asian Americans, but black and white Americans tend to share common surnames. This can be seen in the low success rate of classifying black and white candidates to their races in Table C.1b.

Also note that for surnames that are common among black and white American, using a common threshold makes it more likely to classify a surname to be white. Since there are more white Americans in general, the fraction white is high overall. So even if relative to base population fraction, a surname might be "more black," I use the absolute probability that a person with a given surname belongs to each race.

Audit studies have exploited strong racial associations of first names ([17]). But to be used in observational data, I need a comprehensive list of first names that are predominantly of one or another race. One recent list of black and white first names uses Census sample ([138]). But since this database is based on a sample, there are only 17 first names that are 80% or more black. The most populous of these, Latoya, is based off of 93 observations.

Given this data restriction, I do not exploit the information content of candidate first names, although information is being lost.

The last source of racial identification is the candidates' photo. Whenever a candidate has an associated mug shot, I use facial recognition to infer the candidate's race. To this end, I use free software provided by Face++. Their algorithm classifies race to Asian, black and white. I can check the accuracy of facial recognition by comparing the race from face recognition with the tag information. The algorithm classifies about 92% of white candidates (according to tags) as being white, while it classifies only about 68% of black candidates (according to tags) as being black (Table C.1b). Given the noise, I use the other sources of information first if they are available, before resorting to facial recognition.

Facial recognition is a mature field of computer science, and while younger than other fields of facial recognition, extracting racial identity from faces too has a large literature behind it (see [58] for summary of techniques). The value of facial recognition can be seen in the last column of Table C.1b. Among the vast majority of candidates with no user-supplied tag, 4% of them can be classified as being black American using facial recognition. Even with the noise in classification, the sample gain is useful when creating a database of politicians by race.

When racial information from these three sources disagree, I give precedence to the tag information, surname, then facial recognition.

Creating Jurisdiction Cross-walks

The website Our Campaigns contains interactive maps, with the most recent constituency highlighted using coordinates. I additionally scrape these set of coordinates to construct cross-walks from constituencies to zip codes and counties. The website supplies only the most recent jurisdiction, but political office jurisdictions change over time. While using the most recent jurisdiction for past terms will create measurement error in which zip codes and counties were affected by those offices, in this paper I use only the most recent jurisdictional boundaries, due to the difficulty of acquiring historical jurisdictions for local offices.

3.2.2 Racial Prejudice Data

Several proxies for racial prejudice have been used in the literature, almost always as a static measure. In this paper, I am interested in examining what changes the racial attitudes of a local area. The idea that the underlying racial attitude can change over time is not new. [11] wrote, "Another proximate determinant is geographical and chronological location: discrimination may varyfrom country to country, from region to region within a country, from rural to urban areas within a region, and from one time period to another. Finally, tastes may differ simply because of differences in personality." [63] wrote, "I had a student who used to take the IAT every day. It was the first thing he did, and his idea was just to let the data gather as he went. Then this one day, he got a positive association with blacks. And he said, 'That's odd. I've never gotten that before,' because we've all tried to change our IAT score and we couldn't. But he's a track-and-field guy, and what he realized is that he'd spent the morning watching the Olympics."

Below, I introduce data sets that can be used to produced local panels of racial attitudes.

Implicit Association Test (IAT)

The Race IAT is a widely used test of racial bias ([68]). Respondents are shown either pictures of faces (black or white) and words (good or bad). They use the same set of buttons on the keyboard to classify the faces into black or white categories, and words into good or bad categories. IAT is based on the premise that if a respondent has a stronger association in his mind between being white and being good, the classification exercise will take longer when he has to use the same button to classify a face as being black and a word as being good. The main measures, the D score, is the difference in time it takes to classify, when black faces and good words are paired together (i.e. use the same button) versus when black faces and bad words are paired together.

Since its inception in [68], there has been an explosion of studies using the IAT to test various implicit attitudes. Recent meta-studies find that the Race IAT is a good predictor of racial discrimination ([113]). Criterion measures (official term for the intergroup behaviors) include brain activity, response time, microbehavior, interpersonal behavior, person perception, and policy/political preferences. Explicit measure utilize includes one separate category for "feeling thermometer." Key object in meta-studies is "IAT criterion-related correlation (ICC)." Main results for IAT in Tables 1 & 2. Weaker explicit measure results (explicit-criterion correlation) in Table 5.

Most recent meta-study of the IAT's predictive power ([92]): 1) implicit measures work well regardless of ICC moderator (p26) - "absence of theoretical predictors" (p22); 2) standard IAT is superior (p29). p50 - tighest range of ICC for Race IAT vs. other IATs. Three concepts for which use IAT: attitude stereotype, identity. Key conceptual moderators: social sensitivity, controllability (i.e. automatic activation). Found IAT correlated with behaviors thought to be controllable. "... univariate meta-regressions showed that implicit measures were equally associated with measures of intergroup behavior irrespective of social sensitivity, controllability, conscious awareness, or target concept. In fact, contrary to the widespread notion that implicit measures are not associated with highly controllable behaviors, the present meta-analysis found a sizable number of large ICCs for such behaviors, including self-reported enrollment intentions in mathematics classes (), self-reported career aspirations (), and voting behavior ()."

Also [69]: "for socially sensitive topics, the predictive power of self-report measures was remarkably low and the incremental validity of IAT measures was relatively high." "IAT measures had greater predictive validity than did self-report measures for criterion measures involving interracial behavior and other intergroup behavior."

I get data on online IAT scores from the Project Implicit Database ([145]). The data span years 2003-2017, with roughly a quarter million completed tests per year. The online survey also collects democragraphic variables (age, sex, race, education), political ideology and religious information, explicit racial bias questions comparable to those asked in the General Social Survey (GSS), and self-reported zip code. Figure C.2b shows the distribution of the raw IAT scores, where higher numbers indicate more racial bias.

The IAT has been used widely in economics as well as other social sciences as a prominent measure of racial bias ([122]; [39]). [39] in particular also uses the Project Implicit Database.

Since the online surveys are voluntary, the sample is self-selected and highly unrepresentative of the US population. To adjust for selection, I project the IAT scores on age, race, education (9 buckets), gender, and experimental variables (month, hour, weekday as well as order of experiment), before aggregating them to the race-county-month level. Figure C.2a plots the time series average of this composition-adjusted responses, for white respondents. The issues introduced due to the voluntary nature of data collection are addressed more fully below.

The rich demographic and other details self-reported in the Project Implicit Database show how racial prejudice as measured by the IAT score co-varies with demographics. By far the most salient is race: black respondents' average IAT score is -0.04 against 0.39 for non-Hispanic whites (Figure C.3a). Asians & Hispanics are in-between but closer to whites. The other important individual characteristic is political ideology: 0.39 for conservative against 0.28 for liberals (Figure C.3e). By geographic region, the average score is higher in Northeast and Midwest (0.33) than in South and West (0.30) if including everybody; among only white respondents, the average score is highest in South (0.40) and lowest in West (0.35) (Figure C.3f).

Two widely cited co-variates for racial bias are gender and education (for cognitive ability, see [75]). By gender, the average IAT score is 0.34 for males against 0.30 for females (Figure C.3d). By education, including everyone, the average IAT score is 0.32 for those with college degree or higher and also 0.32 for up to high school graduate (Figure C.3c). Among only white respondents above age 25, the average score is 0.37 for those with college degree or higher, and higher at 0.426 for up to high school graduate. The pattern by age is highly nonlinear, first falling from age 20 to 40, then rising again. By religiosity, among white respondents, those that are not at all religious have the lowest average score (0.35); among all races, those that are strongly religious have the lowest average score (0.27), largely because black respondents are more religious.

A major concern with using the online IAT database is that individuals self-select into taking the test online. I take several approaches to address potential issues that arise as a result. Such self-selection is a feature of most measures of amorphous entities such as racial prejudice. For example, Google searches for racial slurs is another measure of racial bias that is affected by who chooses to use Google and for that purpose ([132]). Yet, the self-selection issue is potentially larger in the case of the IAT because it is a test explicitly designed for racial bias.

First, potential issues can be classified into four categories. Self-selection by survey-takers introduces two sets of issues, which can be broadly summarized as the level and change in selection. The level of selection makes the sample not representative of the broader population, such that external validity claims are impeded. The change of selection is potentially more problematic in that any results I find may be driven not by actual treatment effect of the shocks I consider but by changing selection. Both issues can be further broken out into observable and unobservable kinds. Below, I describe each issue in more detail and how I address each.

Lack of representativeness along the observable dimension can be fully addressed by assigning weights to observations. To check the degree of unrepresentativeness and to assign weights, I use distributions by gender, age and education buckets from the 2008-2012 5-year American Community Survey (ACS), accessed via Integrated Public Use Microdata Series (IPUMS). Table C.2a shows the demographic distribution in the 2008-2012 ACS and the 2008-2012 IAT data in the first two columns. The IAT sample is highly over-represented in the 19-29 age range, and highly under-represented in the ages above 40, for example. To make the sample more representative of the US population along gender, age and education, I impute weights as $\frac{\text{ACS population share}}{\text{IAT sample share}}$ for gender $\times$ age $\times$ education bins, fully interacted. Single-year ages were used, and educational attainment has been grouped into categories shown in Table C.2a. Figures C.4a and C.4b show the weighted distribution of IAT scores for the entire sample 2003-2017 for all races and for white respondents only (red lines) against unweighted distributions (blue lines). The distributions are largely the same, with a slight shift towards the right. The main regression results will be replicated using this re-weighted score.

Lack of represenativeness along the unobservable dimension is more difficult to fully deal with. This issue makes external validity difficult to establish. If estimates do not differ between potentially more and less self-selected subsamples, that reduces the concern that there is much selection along the unobservable dimension.

Selection along observables can be fully addressed by including controls on the right-hand side. Along with the raw IAT score, I will use a composition-adjusted average IAT score, when I project out dummies for age (single-years), nine education buckets and gender. Regression estimates with this composition-adjusted average IAT scores will also be reported below.

Selection along unobservables can be dealt with by running placebo tests by putting observable demographic variable on the left-hand side with the treatment on the right-hand side. The idea is that if the sample changes along unobservable dimensions in response to treatment, such sample changes should partly be reflected along some observable variables. Alternatively, the same regression specifications can be run with and without these observable controls on the right-hand side. If observables change in response to treatment, the treatment effect on outcomes of interest with and without controls should be the same. The distinction between the solution to selection along observable and unobservable dimensions is: for the first, even if estimates change when controls are included, the estimate from the regression with controls is correct; for the second, difference in estimates between the regressions with and without controls already signifies selection along unobservable dimension that cannot be corrected. For all estimation results, estimates with both raw IAT scores and composition-adjusted scores (i.e. comparable to running regressions with controls) will be reported. Just to preview, the composition adjustment with sex, age and education has no effect on the estimates.

Second, I compare responses of those who were assigned to take the test against those who completely voluntarily took it. How responses differ between the more and less self-selected subgroups within the sample will be informative in evaluating how the self-selection of the sample as a whole may bias the results I find. Using the question, "What brought you to this website?" I classify a response as being mandatory if the respondent chose "Assignment for work" or "Assignment for school." I classify a response as as being voluntary if the respondent chose "Recommendation of a friend or co-worker," "Mention or link at a non-news Internet site," "Mention in a news story (any medium)" or "My Internet search for this or a related topic." About 62% of the respondents gave a response to this question since the question was asked in 2006. Relative to respondents who took the test as an assignment, those who took it entirely voluntarily are more likely to be white (71% vs. 64%), more likely to have a bachelor's degree or higher (29% vs. 16%), less likely to be female (51% vs. 61%), more likely to be liberal (58% vs. 39%), more likely to be not religious (39% vs. 27%) and older (average age of 33 vs. 26).

Based on raw level and across all races, those who took the test completely voluntarily have an average Race IAT D score of 0.30, slightly lower than 0.31 for those who took it as a part of an assignment (Table C.2b). Among white respondents, the average IAT D score is 0.36 among those who took it completely voluntarily, and 0.39 among those who took it as a part of an assignment. The IAT score gap between the assignment and completely voluntary groups remains at 0.01 for all races, even after projecting out demographic and experimental controls. The gap for white respondents also remains at 0.03 even after projecting out controls.

I basically make three adjustments to the raw scores by race: 1) projecting out dummies for sex, age and education; 2) assigning weights to be representative along sex, age and education; and 3) checking robustness with only mandatory respondents, i.e. those who replied “Assignment for work” or “Assignment for school” to the question “What brought you to this website?”

Other Prejudice Measures

To complement the IAT scores in measuring time-varying local racial prejudice, I bring in other proxies using multiple data sources. In addition to widely used survey responses from the General Social Survey (GSS), I use Google search trends for keywords associated with racial prejudice, crimes committed by white offenders against black victims, and school corporal punishment on black students relative to white students. These proxies are associated with other things in addition to racial prejudice; for example, the racial slur that I use for Google search trends also often features in the popular rap culture. The goal is to combine proxies that are related to racial prejudice albeit with much measurement error, so that the common component can be attributed to the racial prejudice of the time and place. I will describe each of the measures in turn.

Apart from the General Social Survey, the other proxies are not pure elicitation of preferences. In designating them as proxies for racial prejudice, I follow the criteria from [70] in their indirectment measurement of civic capital: “For an outcome-based measure to qualify as a good indicator of civic capital, the relationship between the input (civic capital) and the measured output should be stable and unaffected by other factors, such as legal enforcement.”

[70] also discuss common movement as a criterion for judging the value of proxies: “Consistent with the idea that these measures are capturing the same underlying norms, they tend to be highly correlated... Hence if one were to rely on measurse of this sort in applied work, one could gain some insights by obtaining several indirect indicators and looking at common components (see [134]).”

General Social Survey (GSS). [34] create an index of local racial prejudice using the GSS surveys. They take a set of questions that I replicate the prejudice index constructed by [34], by standardizing and summing over the responses to the same set of questions they use. Unlike all the other measures including the IAT scores, the GSS has the benefit of being a representative survey.

Google search trends. The two measures are Google search trends for a prominent racial slur (the n-word) and for "KKK." Google shares an index of search volume at the DMA-level since 2004 at the monthly frequency. Google searches for the the n-word has been shown to predict voting against Obama better than the GSS ([132]). Google search trends have several advantages over survey responses. For example, Googling can be done in secret and hence the anonymity suggests that Google searches can reveal racial prejudice more directly than survey responses, where the respondents may feel social pressure to give politically correct answers.

Google search trends also have some issues. First, there are other contexts in which the n-word and "KKK" can be searched for without involving prejudice. For example, the n-word sometimes appears in rap music, although more often rap usage uses the variant of the word that ends in an "a." Fortunately, Google Trends offers most common related searches, which suggest that many uses of these search terms do involve racial prejudice. Second, for privacy reasons, only keywords with sufficient search volume can be tracked on Google Trends. Other ethnic slurs whose usage is less ambiguous are also less widely used and cannot be tracked consistently over time. Third, the time series for the search trends cannot be taken meaningfully. The type of Google users has changed dramatically over the past decade as well as the fads that dominate Googling. Since Google Trends gives me an index, how it changes over time is influenced as much by what else is searched by whom. Cross-sectional comparisons and diff-in-diff comparisons are meaningful and form the backbone of the analysis.

Using Google search trends for "KKK" is new. The Klan is still in operation, and this search proxies for multiple channels related to racial prejudice. Individuals interested in joining it would search for it, and for Google Chrome users, entering the first few characters of the Klan website and pressing enter pre-maturely would lead a user to search for that term on Google. Individuals who fear or are concerned about extremist groups may also search for the Klan as a notorious example, and such searches would also be related to local racial animus. The DMAs with the highest average searches over the 2004-2017 period are Presque Isle, ME, Greenwood-Greenville, MS and Parkersburg, WV. DMAs with the lowest average over the entire period are New York, NY, Washington DC and San Francisco-Oakland-San Jose, CA. The DMAs with the highest average searches over the 2004-2017 period are Parkersburg, WV, Twin Falls, ID and Greenwood-Greenville, MS. DMAs with the lowest average over the entire period are Macon, GA, Portland, OR and Salt Lake City, UT.

White-on-black non-pecuniary crime. The next measure captures more extreme expressions of racial animus. I measure anti-black sentiments using crimes committed by white offenders against black victims, relative to those committed by white offenders against white victims. The data come from the National Archive of Criminal Justice Data (NACJD). Among crime categories, I take three codes only: simple assault, intimidation and destruction/damage/vandalism of property (the uniform common reporting codes are 132, 133 and 290). These crimes are non-pecuniary and therefore less likely to be linked to economic downturns for directly financial reasons. They are also the most common categories of hate crime. In fact, more than 10% of crimes in these categories are hate crimes. I do not use hate crimes directly because they are too few in number and the categorization into hate crime is subjective and potentially influenced by the fact that hate crimes are sometimes punished more severely. Finally, I scale the crime ratio by black-to-white population ratio in the DMA; otherwise, areas with higher black population would mechanically see more crimes committed against blacks. The exact variable definition is:

$$\left[\frac{\text{count of crime by white offender against black victim}_{it}}{\text{count of crime by white offender against white victim}_{it}} \Big/ \frac{\text{total blacks in area}_{it}}{\text{total whites in area}_{it}}\right]$$

for DMA i in year t. . The DMAs with the highest average over the 1991-2014 period are Bangor, ME, Alpena, MI and Sioux City, IA. DMAs with the lowest average over the entire period are Wilkes Barre-Scranton, PA, Tucson, AZ and Jackson, MS.

School spanking. The last measure is the frequency with which school teachers corporally punish (i.e. spank) students of other races relative to white students. Corporal punishment is still legal in 19 U.S. states, and relevant statistics are released by the U.S. Department of Education as a part of the Civil Rights Data Collection. In some sense, this measure best captures changes in racial prejudice. Decision to spank a student may be impulsive and thus reflective of the underlying prejudice, and it is less likely to be affected by economic conditions directly. While a downturn can cause teachers to spank more overall, it is more difficult to think of why they would spank black students more than white students. When dealing with corporal punishment, I use ratio of rates of punishment, as there are large differences in the base rates. Black students are roughly twice as likely as white students to be spanked, whereas Asian students have much lower likelihood.

Crimes and spanking are less direct measures of racial prejudice, less like surveys and more like mortgage denial and non-employment. Conceptually, I imagine that there is a latent average racial prejudice associated with each local area and time. The prejudice measure I compile are a function of the latent racial prejudice and other factors that I claim are otherwise orthogonal to the local economic conditions and household finance outcomes of interest. As I accumulate more of these measures, my proxies will converge onto the latent racial prejudice.

Validation of IAT

The Race IAT D score will be the main measure of racial prejudice in this paper, mainly given its straightforward interpretation and the rich auxiliary demographic information the database provides. In this section, I validate the IAT as a proxy for racial prejudice at the local level, in three ways: comparison to black-white economic disparity (which some argue is partly an outcome of racial discrimination), comparison to the other racial prejudice proxies at both the county and DMA levels, and comparison to historical slavery, which [2] argue is an instrument for today's racial prejudice level.

The comparisons are cross-sectional. Panel comparisons (i.e. difference-in-difference) are coming. Note that none of the correlations below should be interpreted causally, with the possible exception of the historical slavery instrument a la [2].

First, Table C.3c regresses the other prejudice proxies against the area's average IAT score, at either DMA or county level. With the exception of the white-on-black crime measure, average IAT score is higher in areas where the other prejudice proxies are also higher: those areas search more for the n-word and "KKK" on Google, respond to the GSS questionnaire in a potentially more racially biased way, and spank black students more than white students. The opposite correlation for white-on-black crime will be explored further.

Second, I compare my measures to the cross-sectional prejudice instrument proposed in [2], who use as the instruments 1860 slave share and cotton production conditional on state fixed effects (see also [107] for a discussion of the legacy of slavery). They argue that a legacy of slavery passes down over generations, and show that the prevalence of slavery among southern counties predicts survey responses today. The first column of Table C.3a regresses the county's average IAT D score today against the area's 1860 slave share of the population from Census 1860, accessed via IPUMS. All regressions include state fixed effects following [2].

A strong correlate of this instrument is the contemporary black population share. This correlate is problematic for their interpretation because of an alternative hypothesis: the "racial threat" hypothesis argues that it is the current prevalence of black Americans that raises anti-black prejudice. This correlate is problematic for me, because more black Americans may proxy for prevalence of lower-skilled workers, hence raising employment cyclicality, for example. Their first solution is to simply control for current black population share. This makes instrument potentially not valid (i.e. over-control), but their coefficients are not affected much. In my validations, I also control for the contemporary black population share in the second colummn.

Finally, following [2], the third column regresses my average prejudice measures against slave share instrumented using per capita cotton production in 1860.

To give more background on the 1860 data on slavery intensity, 1,117 counties in slave states with non-missing data on slave share of the population. Weighted by 1860 population, average slave population share is 0.32, with a median of 0.30. The highest is 0.925, with the following percentiles: 1% is 0.03, 25% is 0.109, 75% is 0.509, and 90% is 0.65. Counties in the same state sometimes do have different slave population share. For example, in Georgia, Cherokee county (FIPS code 13057) had 0.106, while Clarke county (FIPS 13059) had 0.505. Both counties are in Atlanta CBSA. For another example of larger areas, in Alabama, Mobile county (FIPS 1097) had 0.277, while Montgomery county (FIPS 1101) had 0.66. This difference is comparable to the interquartile difference of around 0.4. States with the most intense slavery by population share are: SC (0.57), MS (0.55), LA (0.47), AL (0.45), FL (0.44), and GA (0.44).

Lastly, I validate my measures by showing that they capture something in common, by extracting a latent state.

3.3 Empirical Methodology and Prejudice Results

This section lays out the empirical specifications designed to examine the reduced form impact of black electoral victory on both racial prejudice measures and racial economic disparity. To summarize, I use three sets of designs: 1) standard difference-in-difference around the election date, comparing counties affected to neighboring counties, 2) differential exposure to most salient black mayoral elections, defining exposure using ex ante level of racial prejudice, and 3) a regression discontinuity design (RDD) using close election of black winners over white contenders.

All the estimation techniques start from a difference-in-difference, using pre- & post-period of 3 years and compared against surrounding geographies in the same state. Exposure analysis compares these estimates against the pre-period level of racial prejudice, with the conjecture that areas that have a higher level of racial bias against blacks will respond more sensitively to the election of a black officeholder. This conjecture is later verified in heterogeneity analysis using regression discontinuity. Close election RDD compares the difference-in-difference estimates against the corresponding vote margin of the election, and then takes the discontinuity where vote margin is 0. These RDD estimates can also be plotted against the pre-period level for heterogeneity. All these regressions can be run with all elections or just the mayoral (or the set of most salient) elections, to examine the tradeoff between sample size and signal-to-noise ratio based on how salient a set of elections are.

While a difference-in-difference estimator is the most transparent, there is a clear identification issue that black electoral victory is not random. In particular, black politicians are more likely to win if white voters' racial prejudice is in decline. This creates a negative bias in the difference-in-difference estimate. Close election regression discontinuity is meant to overcome this identification challenge, relying on a vast literature arguing that winning a closely contested election is as good as random ([51]).

3.3.1 Difference-in-difference Design

I start with the standard difference-in-difference estimator. I first compile a data set at the election-county-year level following [66] for difference-in-difference estimators with multiple events. Using only elections with a black winner, and for election i, county j and year t, I estimate

$$Y_{ijt} = (\text{in jurisdiction } i)_{ij} + (\text{after election } i)_{it} + \Big((\text{in jurisdiction } i)_{ij} \cdot (\text{after election } i)_{it}\Big) + \eta_{ijt}$$

With fixed effects, following [66]:

$$Y_{ijt} = \alpha_{ij} + \alpha_{it} + \Big((\text{in jurisdiction } i)_{ij} \cdot (\text{after election } i)_{it}\Big) + \eta_{ijt}$$

For each election i where a black candidate won the election, I include the 3-year window before and after the election, and include all counties in the same state as the jurisdiction associated with the election.

Of course, black politicians do not get elected randomly. Previous literature has documented factors that predict black politicians' election ([98]). The main concern is negative selection: black politicians are more likely to be elected in areas where dominant white voters' racial bias against blacks is decreasing. This will bias my estimate to find that areas where a black politician comes into office will experience a decrease in racial prejudice.

3.3.2 Regression Discontinuity Design

The empirical strategy is a standard close election regression discontinuity, looking at the 3-year period after an election. The 3-year period before an election is used for placebo tests. I only look at elections where the top two candidates include one black and one white candidate. Results in this draft are obtained by defining "close election" as those with less than 10% vote margin between the winner and the runner-up. Optimal bandwidth (for example using [83]) is wider, and results obtained using optimal bandwidth are not qualitatively different.

For observations at election i, geography (e.g. county) j, and event time (e.g. month or year) t, I run

$$Y_{ijt} = \alpha + \gamma_1 \mathbf{1}\left\{\text{vote margin} > 0\right\}_{it} + \delta_0 \left[\text{vote margin}\right]_{it}^{-} + \delta_1 \left[\text{vote margin}\right]_{it}^{+} + \eta_{ijt}$$

When estimating heterogeneous treatment effect by election type or location characteristic, the specification is:

$$\begin{aligned} Y_{ijt} = \sum_k \Big\{ & \alpha^k + \gamma_1^k \mathbf{1}\left\{\text{vote margin} > 0\right\}_{it} + \delta_0^k \left[\text{vote margin}\right]_{it}^{-} \\ & + \delta_1^k \left[\text{vote margin}\right]_{it}^{+} \Big\} \mathbf{1}\left\{\text{in sub-group}\right\}_{ij} + \eta_{ijt} \end{aligned}$$

As with any reduced form identification scheme, the estimate γ_1 from this RDD is a local average treatment effect (LATE). Generalizing it to all elections is problematic, since close elections are likely to be different from other elections in many dimensions. The bigger external validity issue is if areas that have close elections between black and white candidates are systematically different from those that do not. I explore this issue further by looking at heterogeneity of treatment effects.

Since [95] used close elections to test the median voter theory of [48] by getting at the causal impact of incumbency on future policy, there has been a large literature utilizing close election RDD as an identification scheme to estimate the causal impact of political victory, with well-established econometrics methodology ([84]; [24]).

Table C.5 summarizes the main RDD results for white Americans' racial attitude. The first two columns of Table C.5 show that in the 3-year window before the election, there is no difference in IAT scores between areas where the black candidate will narrowly win and those where she will lose (this is shown graphically in Figure C.7a).

After the election, having a black winner causes IAT score among whites to increase by about 0.03 (last two columns of Table C.5 and Figure C.7b). The estimate is similar whether we use raw IAT score or the composition-adjusted one. This estimate of 0.03 corresponds to about 7% of the raw average gap in the IAT score between all black and white respondents in the Project Implicit Database.

Figures C.8a and C.8b plot the discontinuity estimates for each quarter relative to the election event at 0 (Figure C.8a for the raw IAT and Figure C.8b for the composition-adjusted IAT). There is stable zero difference leading up to the election, but discontinuity rises in the quarters following the election, to come back down eventually. Both the pre-period placebo and the time series of discontinuity plots suggest that the RD strategy is picking up a causal estimate.

3.3.3 Heterogeneity

This section explores the heterogeneity in treatment effects. The first dimension of heterogeneity is the level of racial prejudice. While the initial conjecture is that in areas with higher level of prejudice, seeing a black leader would lead to stronger backlash among whites (a point also made in assessing external validity in [9]), the major confound is that more prejudiced areas also tend to be more black and hence are more likely to be familiar with black politicians. Given the still low sample size of black-white races, I examine each area characteristic separately.

Another interesting dimension of heterogeneity is the economic condition of the area. In a downturn, white workers may feel less economically secure and perceive higher threat from black electoral victory ([21]).

Table C.7a presents heterogeneity results using

$$Y_{ijt} = \sum_k \Big\{ \alpha^k + \gamma_1^k \mathbf{1}\{\text{vote margin} > 0\}_{it} + \delta_0^k [\text{vote margin}]_{it}^- \\ + \delta_1^k [\text{vote margin}]_{it}^+ \Big\} \mathbf{1}\{\text{in sub-group}\}_{ij} + \eta_{ijt}$$

where the group k is defined by the below- and above-median areas sorting by average IAT score level, black population share, or average income. Columns (3)-(6) show that treatment effect is not heterogeneous by either the black population share or income.

Columns (1)-(2) of Table C.7a shows that the rise in prejudice is entirely concentrated in areas with above-median level of racial prejudice. In fact, the treatment effect in below-median prejudice level areas is insignificant and negative, consistent with [9].

Table C.7b repeats the heterogeneous treatment effect analysis using the level of IAT, with black-white economic disparity variables. While statistical power drops as expected with heterogeneous treatment effects, point estimates consistently suggest that the negative effect of black electoral victory on black relative economic outcome is concentrated in high-prejudice areas.

3.4 Effect on Racial Gaps in Employment and Credit

3.4.1 Data

Data on employment gap come from Quarterly Workforce Indicators (QWI) and refer to all private employment at the county-level. The two main measures are likelihood of transitioning to non-employment among those employed (a flow measure), and employment ratio (a stock measure).

Data on mortgage denial come from the Home Mortgage Disclosure Act (HMDA) data. I calculate rejection rate as the number of mortgage applications for owner-occupied housing that were rejected, divided by all loans that were originated or denied. Since rejection rate is affected by who applies, I also look at mortgage origination per capita. The highest frequency of the HMDA data is year. Geographically, in one version I aggregate up to the county using the state and county variable; in a more disaggregated version, I use the tract variable to aggregate to zip codes. The tract variable is missing only for about 30 million out of about 500 million loan observations (i.e. the tract variable has value of "NA"). The HMDA TS file has variable "rzip" but that is likely the lender's zip code.

For all of these measures, I take the difference between black and white individuals in a given county as my economic gap measures.

For main measure, I will use the one that can be adjusted most promptly to reflect whatever change in underlying attitudes. For employment, this is the relative transition to non-employment. For mortgage, this would be relative denial rate.

3.4.2 Difference-in-difference

Table C.6a shows the difference-in-difference estimates for employment and mortgage variables. As with the prejudice measures, economic gaps show no change with the election of a black politician, most likely given that black politicians are more likely to be elected where white voters' racial bias is declining.

Turning to identified regression discontinuity estimates, however, Table C.6b shows patterns consistent with a negative causal effect on racial prejudice.

3.4.3 Regression Discontinuity

Racial Economic Inequality

Then, looking at economic inequality measures, Table C.6b shows that black politicians' electoral victory causes black workers to transition more into non-employment (flow), fewer blacks to be in employment (stock), originate less mortgage amount and be denied more mortgages, relative to white counterparts.

3.5 Causal Identification of Racial Discrimination

3.5.1 Instrumental Variable (IV) Estimator

I start from the standard regression discontinuity exclusion restriction that whether the black candidate wins or loses in a close election is uncorrelated with local conditions, except through its effect on local racial prejudice:

$$E\left[\mathbf{1}\left\{\text{vote margin} > 0\right\}_{it}\,\varepsilon_{ijt}|\text{IAT}_{ijt}\right] = 0$$

where ε_{ijt} is the unobservable factors in

$$Y_{ijt} = \beta_0 + \beta_1 \text{IAT}_{ijt} + \varepsilon_{ijt}$$

for some black-white economic disparity measure Y_{ijt} in county j in year t, surrounding an election i.

This exclusion restriction motivates the following two-stage least squares (2SLS) estimator.

$$\begin{aligned} Y_{ijt} &= \beta_0 + \beta_1 \hat{\text{IAT}}_{ijt} + \tilde{\delta}_0\left[\text{vote margin}\right]_{it}^{-} + \tilde{\delta}_1\left[\text{vote margin}\right]_{it}^{+} + \varepsilon_{ijt} \\ \text{IAT}_{ijt} &= \alpha + \gamma_1 \mathbf{1}\left\{\text{vote margin} > 0\right\}_{it} + \delta_0\left[\text{vote margin}\right]_{it}^{-} + \delta_1\left[\text{vote margin}\right]_{it}^{+} + \eta_{ijt} \end{aligned}$$

The linear controls on either side of the vote margin enter both first stage and second stage regressions.

3.5.2 IV Estimation Results

The main IV estimation results are reported in Table C.8. Given the exclusion restriction, the first row estimates show the causal impact of 1-point increase in the composition-adjusted IAT score among whites in the county. To interpret the magnitude of the estimate for mortgage rejection rate, a 0.1 increase in racial bias among whites as measured by the IAT would lead to the black-white rejection rate gap to widen by 6.6 percentage points. Most of the estimates are borderline insignificant at the 5% level, largely due to a lack of instrument strength as can be seen in Figure C.7b.

For a naive and surely wrong back-of-the-envelope calculation, the mortgage rejection rate estimate of 0.658 along with the average IAT score of 0.39 implies that if whites' racial bias against blacks as measured by the IAT fell to 0, the rejection rate gap would fall by roughly 25 percentage points, equivalent to the actual rejection rate gap. Such extrapolation likely over-estimates the true counterfactual given endogenous responses and that my IV estimate is a LATE, but nevertheless demonstrates that the IV estimate is sizable.

3.5.3 Relation to Literature

Example papers that show racial discrimination in various real life settings: [16]; [47]; [62]; [118]; [133]

In this section, I describe conditions under which I can interpret visible black politicians' close electoral victory as an instrumental variable for racial prejudice. This methodology allows for a different way to identify racial discrimination, under potentially more plausible exclusion restrictions.

The economics literature on racial discrimination using observational data mostly runs regressions of the following form ([35]):

$$Y_i = \beta \text{black}_i + X_i \Gamma + \tilde{\varepsilon}_i$$

where Y_i is some outcome in the market in which we are studying discrimination (e.g. mortgage rejection), black_i is individual i's race, and X_i is a vector of controls such as income. In terms of difference between black and white individuals and taking within-group averages, an equivalent expression is:

$$\text{gap} \equiv E\left[Y_{i,\text{black}} - Y_{i,\text{white}}\right] = \beta + E\left[X_{i,\text{black}} - X_{i,\text{white}}\right]\Gamma$$

In either expression, the discrimination coefficient $\hat{\beta}$ is a residual from controlling for other characteristics. There is a well-known over-controlling problem here: if an area is highly prejudiced in both the mortgage market and the labor market, blacks may both earn less than whites and be rejected mortgages more frequently. In such a setting, it is possible for the specification to estimate $\hat{\beta} = 0$, even though mortgage market too practices racial discrimination. In other words, $\hat{\beta}$ from such specifications can only tell us how discriminatory the market for Y_i is relative to the markets for the controls in X_i.

Audit studies are a way to get around this issue by essentially "shocking" black_i by experimentally varying the perceived race of the applicant. By submitting identical applications but only varying the name of the applicant, these studies estimate

$$Y_i = \beta \widehat{\text{black}}_i + \varepsilon_i$$

where $\widehat{\text{black}}_i$ can be thought of as the perceived probability that an applicant is black given that the name is Lakisha and Jamal as opposed to Emily and Greg ([17]).

A handful of papers take a different approach. Starting from the insight that taste for discrimination is a combination of the target of discrimination and the intensity of the racial prejudice, they estimate regressions of the form:

$$Y_{ij} = \delta \text{black}_i + \beta \text{black}_i \times \text{prejudice}_j + \varepsilon_{ij}$$

by using local measures of prejudice_j. [34] uses General Social Survey responses to get geographical variations in prejudice_j. [2] instruments for prejudice_j using historical prevalence of slavery driven by cotton production intensity in 1860.

In this paper, I extend this second approach. I associate whole time-places with a level of racial prejudice and find proxies for it. Then, without including other controls,

$$Y_{ijt} = \delta \text{black}_i + \beta \text{black}_i \times \widehat{\text{prejudice}}_{jt} + \varepsilon_{ijt}$$

or equivalently in black-white differences,

$$\text{gap}_{jt} = \delta + \beta \hat{\text{prejudice}}_{jt} + \tilde{\varepsilon}_{jt}$$

That is, if black households' mortgage rejections rise in an area as people in that area exhibit higher racial bias as measured by the Race IAT, I associate that co-movement with racial discrimination. Conceptually this method compares different geographical areas as opposed to different markets (for example, mortgage market against labor market as above). As a result, I cannot conclude which market in a geographical area exhibits prejudice; I can only say that city X has more discrimination than city Y. By instrumenting for prejudice_{jt} using close electoral victory of black politicians, I can make a causal claim for β. To my knowledge, this is the first time-varying instrument for racial prejudice in observational data.

3.6 Conclusion

The Obama presidency has motivated questions as to how having a visible black leader affects white Americans' attitude toward black Americans. Given the theoretical ambiguity, I test for causal impact of a black leader on racial attitudes using local elections of black politicians at the municipal level. Using Race Implicit Attitude Test (IAT) scores as a measure of racial prejudice and close election regression discontinuity (RD) design for causal inference, I find that electoral victory of a black leader leads to a rise in racial prejudice among white Americans against black Americans. Following a close electoral victory, the IAT D score rises by about 0.03, or 7% of the average black-white difference. Simultaneously, using the same RD design, black politicians' electoral victory causes lower employment and higher mortgage denial for black Americans relative to white Americans. By ruling out other channels by which electoral victory could adversely affect black Americans' relative economic outcome, I argue that the rise in prejudice caused black-white economic inequality to widen.

BIBLIOGRAPHY

[1] F Glenn Abney and John D Hutcheson Jr. Race, representation, and trust: Changes in attitudes after the election of a black mayor. *Public Opinion Quarterly*, 45(1):91–101, 1981.

[2] Avidit Acharya, Matthew Blackwell, and Maya Sen. The political legacy of american slavery. *The Journal of Politics*, 78(3):621–641, 2016.

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

[4] David H Autor, David Dorn, and Gordon H Hanson. The china syndrome: Local labor market effects of import competition in the united states. *The American Economic Review*, 103(6):2121–2168, 2013.

[5] Laurent Bach, Laurent E Calvet, and Paolo Sodini. Rich pickings? risk, return, and skill in the portfolios of the wealthy. 2016.

[6] Nicholas Barberis, Robin Greenwood, Lawrence Jin, and Andrei Shleifer. X-capm: An extrapolative capital asset pricing model. *Journal of financial economics*, 115(1):1–24, 2015.

[7] Patrick Bayer, Marcus Casey, Fernando Ferreira, and Robert McMillan. Racial and ethnic price differentials in the housing market. *Journal of Urban Economics*, 102:91–105, 2017.

[8] Patrick Bayer, Fernando Ferreira, and Stephen L Ross. The vulnerability of minority homeowners in the housing boom and bust. *American Economic Journal: Economic Policy*, 8(1):1–27, 2016.

[9] Lori Beaman, Raghabendra Chattopadhyay, Esther Duflo, Rohini Pande, and Petia Topalova. Powerful women: does exposure reduce bias? *The Quarterly journal of economics*, 124(4):1497–1540, 2009.

[10] Lori Beaman, Esther Duflo, Rohini Pande, and Petia Topalova. Female leadership raises aspirations and educational attainment for girls: A policy experiment in india. *science*, 335(6068):582–586, 2012.

[11] Gary S Becker. *The Economics of Discrimination*. University of Chicago press, [1957] 1971.

[12] Jess Benhabib, Alberto Bisin, and Mi Luo. Wealth distribution and social mobility in the us: A quantitative approach. Technical report, National Bureau of Economic Research, 2015.

[13] Martin Beraja, Erik Hurst, and Juan Ospina. The aggregate implications of regional business cycles. Technical report, National Bureau of Economic Research, 2016.

[14] David Berger, Veronica Guerrieri, Guido Lorenzoni, and Joseph Vavra. House prices and consumer spending. *The Review of Economic Studies*, 85(3):1502–1542, 2017.

[15] Leonard Berkowitz. *Roots of aggression: A re-examination of the frustration-aggression hypothesis*. Atherton Press, 1969.

[16] Marianne Bertrand, Dolly Chugh, and Sendhil Mullainathan. Implicit discrimination. *American Economic Review*, 95(2):94–98, 2005.

[17] Marianne Bertrand and Sendhil Mullainathan. Are emily and greg more employable than lakisha and jamal? a field experiment on labor market discrimination. *American economic review*, 94(4):991–1013, 2004.

[18] John Beshears, James J Choi, Christopher Harris, David Laibson, Brigitte C Madrian, and Jung Sakong. Self control and commitment: can decreasing the liquidity of a savings account increase deposits? 2015.

[19] Jeff E Biddle and Daniel S Hamermesh. Wage discrimination over the business cycle. *IZA Journal of Labor Policy*, 2(1):7, 2013.

[20] Alberto Bisin and Jess Benhabib. Skewed wealth distributions: Theory and empirics. *Journal of Economic Literature*, 2017.

[21] Lawrence Bobo. Group conflict, prejudice, and the paradox of contemporary racial attitudes. In *Eliminating racism*, pages 85–114. Springer, 1988.

[22] Fernanda Brollo and Ugo Troiano. What happens when a woman wins an election? evidence from close races in brazil. *Journal of Development Economics*, 122:28–45, 2016.

[23] Eric Brunner, Stephen L Ross, and Ebonya Washington. Economics and policy preferences: causal evidence of the impact of economic conditions on support for redistribution and other ballot proposals. *Review of Economics and Statistics*, 93(3):888–906, 2011.

[24] Sebastian Calonico, Matias D Cattaneo, and Rocio Titiunik. Robust nonparametric confidence intervals for regression-discontinuity designs. *Econometrica*, 82(6):2295–2326, 2014.

[25] Laurent E Calvet, John Y Campbell, and Paolo Sodini. Fight or flight? portfolio rebalancing by individual investors. *The Quarterly journal of economics*, 124(1):301–348, 2009.

[26] John Y Campbell. Household finance. *The journal of finance*, 61(4):1553–1604, 2006.

[27] John Y Campbell, Tarun Ramadorai, and Benjamin Ranish. Do the rich get richer in the stock market? evidence from india. 2018.

[28] John Y Campbell, Luis M Viceira, Luis M Viceira, et al. *Strategic asset allocation: portfolio choice for long-term investors*. Clarendon Lectures in Economic, 2002.

[29] Sean D Campbell, Morris A Davis, Joshua Gallin, and Robert F Martin. What moves housing markets: A variance decomposition of the rent–price ratio. *Journal of Urban Economics*, 66(2):90–102, 2009.

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

[31] Christopher D Carroll. Why do the rich save so much? *Does Atlas Shrug?: The Economic Consequences of Taxing the Rich*, page 465, 2000.

[32] Francesco Caselli and Wilbur John Coleman. On the theory of ethnic conflict. *Journal of the European Economic Association*, 11(suppl_1):161–192, 2013.

[33] Southern Poverty Law Center. Whose heritage? public symbols of the confederacy, 2016.

[34] Kerwin Kofi Charles and Jonathan Guryan. Prejudice and wages: an empirical assessment of becker's the economics of discrimination. *Journal of Political Economy*, 116(5):773–809, 2008.

[35] Kerwin Kofi Charles and Jonathan Guryan. Studying discrimination: Fundamental challenges and recent progress. *Annu. Rev. Econ.*, 3(1):479–511, 2011.

[36] Kerwin Kofi Charles and Erik Hurst. The transition to home ownership and the black-white wealth gap. *The Review of Economics and Statistics*, 84(2):281–297, 2002.

[37] Kerwin Kofi Charles, Erik Hurst, and Matthew J Notowidigdo. Housing booms and busts, labor market opportunities, and college attendance. Technical report, National Bureau of Economic Research, 2015.

[38] Raghabendra Chattopadhyay and Esther Duflo. Women as policy makers: Evidence from a randomized policy experiment in india. *Econometrica*, 72(5):1409–1443, 2004.

[39] Raj Chetty, Nathaniel Hendren, Maggie R Jones, and Sonya R Porter. Race and economic opportunity in the united states: An intergenerational perspective. Technical report, National Bureau of Economic Research, 2018.

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

[41] Joshua Comenetz. Frequently occurring surnames in the 2010 census, 2016.

[42] Morris A Davis and Stijn Van Nieuwerburgh. Housing, finance, and the macroeconomy. In *Handbook of regional and urban economics*, volume 5, pages 753–811. Elsevier, 2015.

[43] Mariacristina De Nardi and Giulio Fella. Saving and wealth inequality. *Review of Economic Dynamics*, 26:280–300, 2017.

[44] Stefano DellaVigna. The obama effect on economic outcomes: Evidence from event studies. Technical report, Citeseer, 2010.

[45] Denise DiPasquale and Edward L Glaeser. The los angeles riot and the economics of urban unrest. *Journal of Urban Economics*, 43(1):52–78, 1998.

[46] John Dollard, Neal E Miller, Leonard W Doob, Orval Hobart Mowrer, and Robert R Sears. *Frustration and aggression.* Yale University Press, 1939.

[47] John J Donohue III and Steven D Levitt. The impact of race on policing and arrests. *The Journal of Law and Economics*, 44(2):367–394, 2001.

[48] Anthony Downs. An economic theory of political action in a democracy. *Journal of political economy*, 65(2):135–150, 1957.

[49] Christian Dustmann, Francesca Fabbri, and Ian Preston. Racial harassment, ethnic concentration, and economic conditions. *The Scandinavian Journal of Economics*, 113(3):689–711, 2011.

[50] Daniel A Effron, Jessica S Cameron, and Benoit Monin. Endorsing obama licenses favoring whites. *Journal of experimental social psychology*, 45(3):590–593, 2009.

[51] Andrew C Eggers, Anthony Fowler, Jens Hainmueller, Andrew B Hall, and James M Snyder. On the validity of the regression discontinuity design for estimating electoral effects: New evidence from over 40,000 close races. *American Journal of Political Science*, 59(1):259–274, 2015.

[52] Ernst Engel. Die produktions-und konsumptionsverhältnisse des königreichs sachsen. *Zeitschrift des Statistischen Bureaus des Königlich Sächsischen Ministeriums des Innern*, 8:1–54, 1857.

[53] Andreas Fagereng, Luigi Guiso, Davide Malacrino, and Luigi Pistaferri. Heterogeneity and persistence in returns to wealth. Technical report, National Bureau of Economic Research, 2016.

[54] Fernando Ferreira and Joseph Gyourko. Do political parties matter? evidence from us cities. *The Quarterly journal of economics*, 124(1):399–422, 2009.

[55] Fernando Ferreira and Joseph Gyourko. Does gender matter for political leadership? the case of us mayors. *Journal of Public Economics*, 112:24–39, 2014.

[56] Benjamin M Friedman. *The Moral Consequences of Economic Growth.* Vintage, 2010.

[57] Paul Frijters. Discrimination and job-uncertainty. *Journal of economic behavior & organization*, 36(4):433–446, 1998.

[58] Siyao Fu, Haibo He, and Zeng-Guang Hou. Learning race from face: A survey. *IEEE transactions on pattern analysis and machine intelligence*, 36(12):2483–2509, 2014.

[59] Xavier Gabaix, Jean-Michel Lasry, Pierre-Louis Lions, and Benjamin Moll. The dynamics of inequality. *Econometrica*, 84(6):2071–2111, 2016.

[60] Bertrand Garbinti, Jonathan Goupille-Lebret, and Thomas Piketty. Accounting for wealth inequality dynamics: Methods, estimates and simulations for france (1800-2014). 2017.

[61] Eric Ghysels, Alberto Plazzi, Rossen Valkanov, and Walter Torous. Forecasting real estate prices. In *Handbook of economic forecasting*, volume 2, pages 509–580. Elsevier, 2013.

[62] Laura Giuliano, David I Levine, and Jonathan Leonard. Racial bias in the manager-employee relationship an analysis of quits, dismissals, and promotions at a large retail firm. *Journal of Human Resources*, 46(1):26–52, 2011.

[63] Malcolm Gladwell. *Blink: The power of thinking without thinking*. Back Bay Books, 2007.

[64] Edward L Glaeser and Charles G Nathanson. Housing bubbles. In *Handbook of regional and urban economics*, volume 5, pages 701–751. Elsevier, 2015.

[65] Seth K Goldman and Diana C Mutz. *The Obama effect: How the 2008 campaign changed white racial attitudes*. Russell Sage Foundation, 2014.

[66] Todd A Gormley and David A Matsa. Growing out of trouble? corporate responses to liability risk. *The Review of Financial Studies*, 24(8):2781–2821, 2011.

[67] Donald P Green, Jack Glaser, and Andrew Rich. From lynching to gay bashing: the elusive connection between economic conditions and hate crime. *Journal of Personality and Social Psychology*, 75(1):82, 1998.

[68] Anthony G Greenwald, Debbie E McGhee, and Jordan LK Schwartz. Measuring individual differences in implicit cognition: the implicit association test. *Journal of personality and social psychology*, 74(6):1464, 1998.

[69] Anthony G Greenwald, T Andrew Poehlman, Eric Luis Uhlmann, and Mahzarin R Banaji. Understanding and using the implicit association test: Iii. meta-analysis of predictive validity. *Journal of personality and social psychology*, 97(1):17, 2009.

[70] Luigi Guiso, Paola Sapienza, and Luigi Zingales. Civic capital as the missing link. In *Handbook of social economics*, volume 1, pages 417–480. Elsevier, 2011.

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

[72] Fatih Guvenen, Serdar Ozkan, and Jae Song. The nature of countercyclical income risk. *Journal of Political Economy*, 122(3):621–660, 2014.

[73] Jonathan Heathcote, Kjetil Storesletten, and Giovanni L Violante. Quantitative macroeconomics with heterogeneous households. *Annu. Rev. Econ.*, 1(1):319–354, 2009.

[74] Lancelot Henry de Frahan and Jung Sakong. Intergenerational elasticity of consumption. 2018.

[75] Gordon Hodson and Michael A Busseri. Bright minds and dark attitudes: Lower cognitive ability predicts greater prejudice through right-wing ideology and low intergroup contact. *Psychological science*, 23(2):187–195, 2012.

[76] Michael A Hogg. Subjective uncertainty reduction through self-categorization: A motivational theory of social identity processes. *European Review of Social Psychology*, 11(1):223–255, 2000.

[77] Michael A Hogg and Dominic Abrams. Social motivation, self-esteem and social identity. *Social identity theory: Constructive and critical advances*, 28:47, 1990.

[78] Michael A Hogg and Dominic Abrams. Towards a single-process uncertainty-reduction model of social motivation in groups. 1993.

[79] Jeffrey Hoopes, Patrick Langetieg, Stefan Nagel, Daniel Reck, Joel Slemrod, and Bryan Stuart. 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, 2016.

[80] Daniel J Hopkins and Katherine T McCabe. After it's too late: estimating the policy impacts of black mayoralties in us cities. *American Politics Research*, 40(4):665–700, 2012.

[81] Carl Iver Hovland and Robert R Sears. Minor studies of aggression: Vi. correlation of lynchings with economic indices. *The Journal of Psychology*, 9(2):301–310, 1940.

[82] Jon Hurwitz and Mark Peffley. Public perceptions of race and crime: The role of racial stereotypes. *American journal of political science*, pages 375–401, 1997.

[83] Guido Imbens and Karthik Kalyanaraman. Optimal bandwidth choice for the regression discontinuity estimator. *The Review of economic studies*, 79(3):933–959, 2012.

[84] Guido W Imbens and Thomas Lemieux. Regression discontinuity designs: A guide to practice. *Journal of econometrics*, 142(2):615–635, 2008.

[85] David W Johnston and Grace Lordan. Racial prejudice and labour market penalties during economic downturns. *European Economic Review*, 84:57–75, 2016.

[86] Greg Kaplan, Kurt Mitman, and Giovanni L Violante. The housing boom and bust: Model meets evidence. Technical report, National Bureau of Economic Research, 2017.

[87] Edmond J Keller. The impact of black mayors on urban policy. *The Annals of the American Academy of Political and Social Science*, 439(1):40–52, 1978.

[88] Donald Kinder and Jennifer Chudy. After obama. In *The Forum*, volume 14, pages 3–15. De Gruyter, 2016.

[89] Maryam Kouchaki. Vicarious moral licensing: the influence of others' past moral actions on moral behavior. *Journal of personality and social psychology*, 101(4):702, 2011.

[90] Alan B Krueger and Jörn-Steffen Pischke. A statistical analysis of crime against foreigners in unified germany. *Journal of Human Resources*, pages 182–209, 1997.

[91] Moritz Kuhn and José-Victor Rios-Rull. 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, 2016.

[92] Benedek Kurdi, Allison E Seitchik, Jordan Axt, Timothy Carroll, Arpi Karapetyan, Neela Kaushik, Diana Tomezsko, Anthony G Greenwald, and Mahzarin R Banaji. Relationship between the implicit association test and intergroup behavior: A meta-analysis. 2018.

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

[94] Kevin Lang, Jee Lehmann, and K Yeon. Racial discrimination in the labor market: Theory and empirics. *Journal of Economic Literature*, 50(4):959–1006, 2012.

[95] David S Lee, Enrico Moretti, and Matthew J Butler. Do voters affect or elect policies? evidence from the us house. *The Quarterly Journal of Economics*, 119(3):807–859, 2004.

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

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

[98] Melissa J Marschall and Anirudh VS Ruhil. The pomp of power: Black mayoralties in urban america. *Social Science Quarterly*, 87(4):828–850, 2006.

[99] Atif Mian, Kamalesh Rao, and Amir Sufi. Household balance sheets, consumption, and the economic slump. *The Quarterly Journal of Economics*, 128(4):1687–1726, 2013.

[100] Atif Mian and Amir Sufi. House prices, home equity–based borrowing, and the us household leverage crisis. *The American Economic Review*, 101(5):2132–2156, 2011.

[101] Atif Mian and Amir Sufi. Who bears the cost of recessions? the role of house prices and household debt. In *Handbook of Macroeconomics*, volume 2, pages 255–296. Elsevier, 2016.

[102] Atif Mian, Amir Sufi, and Emil Verner. Household debt and business cycles worldwide. *The Quarterly Journal of Economics*, 132(4):1755–1817, 2017.

[103] Atif Mian, Amir Sufi, and Emil Verner. How do credit supply shocks affect the real economy? evidence from the united states in the 1980s. Technical report, National Bureau of Economic Research, 2017.

[104] Neal E Miller. I. the frustration-aggression hypothesis. *Psychological review*, 48(4):337, 1941.

[105] Benoit Monin and Dale T Miller. Moral credentials and the expression of prejudice. *Journal of personality and social psychology*, 81(1):33, 2001.

[106] Yusuf Neggers. Enfranchising your own? experimental evidence on bureaucrat diversity and election bias in india. *American Economic Review*, 108(6):1288–1321, 2018.

[107] Nathan Nunn. The importance of history for economic development. *Annu. Rev. Econ.*, 1(1):65–92, 2009.

[108] John VC Nye, Ilia Rainer, and Thomas Stratmann. Do black mayors improve black relative to white employment outcomes? evidence from large us cities. *The Journal of Law, Economics, & Organization*, 31(2):383–430, 2014.

[109] National Association of Realtors. 2016 profile of international activity in u.s. residential real estate. Technical report, 2016.

[110] Melvin Oliver and Thomas Shapiro. *Black wealth/white wealth: A new perspective on racial inequality.* Routledge, 2013.

[111] Susan Olzak. The political context of competition: Lynching and urban racial violence, 1882–1914. *Social Forces*, 69(2):395–421, 1990.

[112] Susan Olzak and Suzanne Shanahan. Deprivation and race riots: An extension of spilerman's analysis. *Social Forces*, 74(3):931–961, 1996.

[113] Frederick L Oswald, Gregory Mitchell, Hart Blanton, James Jaccard, and Philip E Tetlock. Predicting ethnic and racial discrimination: A meta-analysis of iat criterion studies. *Journal of personality and social psychology*, 105(2):171, 2013.

[114] Christopher Sebastian Parker. Race and politics in the age of obama. *Annual Review of Sociology*, 42:217–230, 2016.

[115] Fabian T Pfeffer, Robert F Schoeni, Arthur Kennickell, and Patricia Andreski. Measuring wealth and wealth inequality: Comparing two us surveys. *Journal of economic and social measurement*, 41(2):103–120, 2016.

[116] Monika Piazzesi and Martin Schneider. Housing and macroeconomics. In *Handbook of Macroeconomics*, volume 2, pages 1547–1640. Elsevier, 2016.

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

[118] Joseph Price and Justin Wolfers. Racial discrimination among nba referees. *The Quarterly journal of economics*, 125(4):1859–1887, 2010.

[119] Vincenzo Quadrini and José-Víctor Ríos-Rull. Inequality in macroeconomics. In *Handbook of Income Distribution*, volume 2, pages 1229–1302. Elsevier, 2015.

[120] Lincoln Quillian. Prejudice as a response to perceived group threat: Population composition and anti-immigrant and racial prejudice in europe. *American sociological review*, pages 586–611, 1995.

[121] Raghuram G Rajan. *Fault lines: How hidden fractures still threaten the world economy.* princeton University press, 2011.

[122] Ernesto Reuben, Paola Sapienza, and Luigi Zingales. How stereotypes impair women's careers in science. *Proceedings of the National Academy of Sciences*, 111(12):4403–4408, 2014.

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

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

[125] Jung Sakong. Cross-border gentrification: Chinese buyers' housing price impact and xenophobia. 2018.

[126] Jung Sakong. The obama effect: Effect of black electoral victory on racial prejudice and inequality. 2018.

[127] Jung Sakong. Racial prejudice is not normal: A collage of empirical evidence. 2018.

[128] Grace Hall Saltzstein. Black mayors and police policies. *The journal of Politics*, 51(3):525–544, 1989.

[129] Jim Sidanius, Erik Devereux, and Felicia Pratto. A comparison of symbolic racism theory and social dominance theory as explanations for racial policy attitudes. *The Journal of Social Psychology*, 132(3):377–395, 1992.

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

[131] John Smith. Reputation, social identity and social conflict. *Journal of Public Economic Theory*, 14(4):677–709, 2012.

[132] Seth Stephens-Davidowitz. The cost of racial animus on a black candidate: Evidence using google search data. *Journal of Public Economics*, 118:26–40, 2014.

[133] Michael A Stoll, Steven Raphael, and Harry J Holzer. Black job applicants and the hiring officer's race. *ILR Review*, 57(2):267–287, 2004.

[134] Guido Tabellini. Culture and institutions: economic development in the regions of europe. *Journal of the European Economic association*, 8(4):677–716, 2010.

[135] Henri Tajfel and John C Turner. An integrative theory of intergroup conflict. *The social psychology of intergroup relations*, 33(47):74, 1979.

[136] Michael Tesler. *Post-racial or most-racial?: Race and politics in the Obama Era.* University of Chicago Press, 2016.

[137] Michael Tesler and David O Sears. *Obama's race: The 2008 election and the dream of a post-racial America.* University of Chicago Press, 2010.

[138] Konstantinos Tzioumis. Demographic aspects of first names. *Scientific data*, 5:180025, 2018.

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

[140] Ebonya Washington. How black candidates affect voter turnout. *The Quarterly Journal of Economics*, 121(3):973–998, 2006.

[141] Iván Werning. Incomplete markets and aggregate demand. Technical report, National Bureau of Economic Research, 2015.

[142] Christopher H Wheeler and Luke M Olson. Racial differences in mortgage denials over the housing cycle: Evidence from us metropolitan areas. *Journal of Housing Economics*, 30:33–49, 2015.

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

[144] David L Word, Charles D Coleman, Robert Nunziata, and Robert Kominski. Demographic aspects of surnames from census 2000. *Unpublished manuscript, Retrieved from http://citeseerx. ist. psu. edu/viewdoc/download*, 2008.

[145] Kaiyuan Xu, Brian Nosek, and Anthony Greenwald. Psychology data from the race implicit association test on the project implicit demo website. *Journal of Open Psychology Data*, 2(1), 2014.

APPENDIX A

CYCLICAL HOUSING TRANSACTIONS AND WEALTH INEQUALITY

A.1 Figures and Tables

Figure A.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

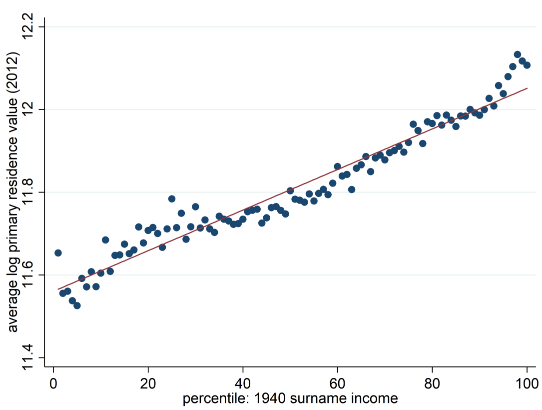

(b) By surname: Per capita holding vs. 1998

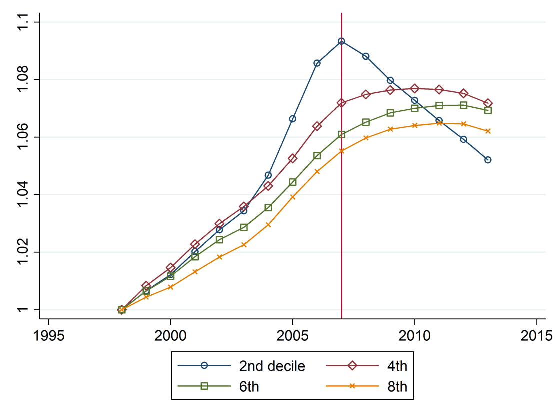

(c) By surname: Detrended log residual

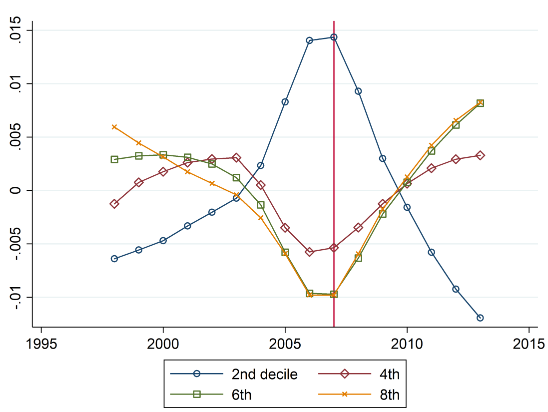

(d) National house price: Detrended log residual

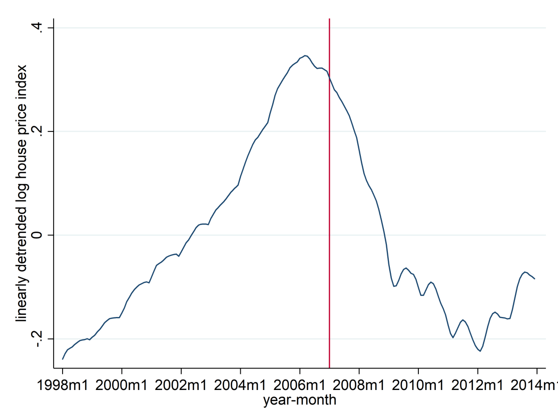

Figure A.2: Estimated quantity elasticity versus wealth level

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

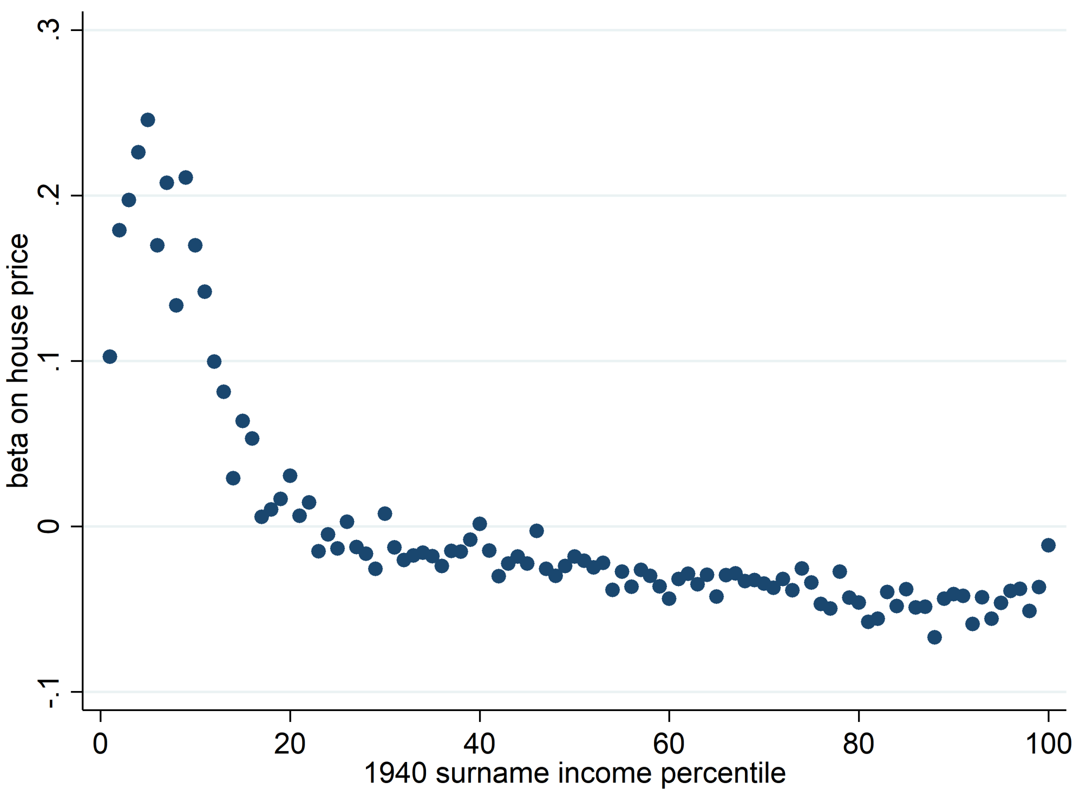

Figure A.2: Estimated quantity elasticity versus wealth level (continued)

(b) For 1988-2013 sample: Full period

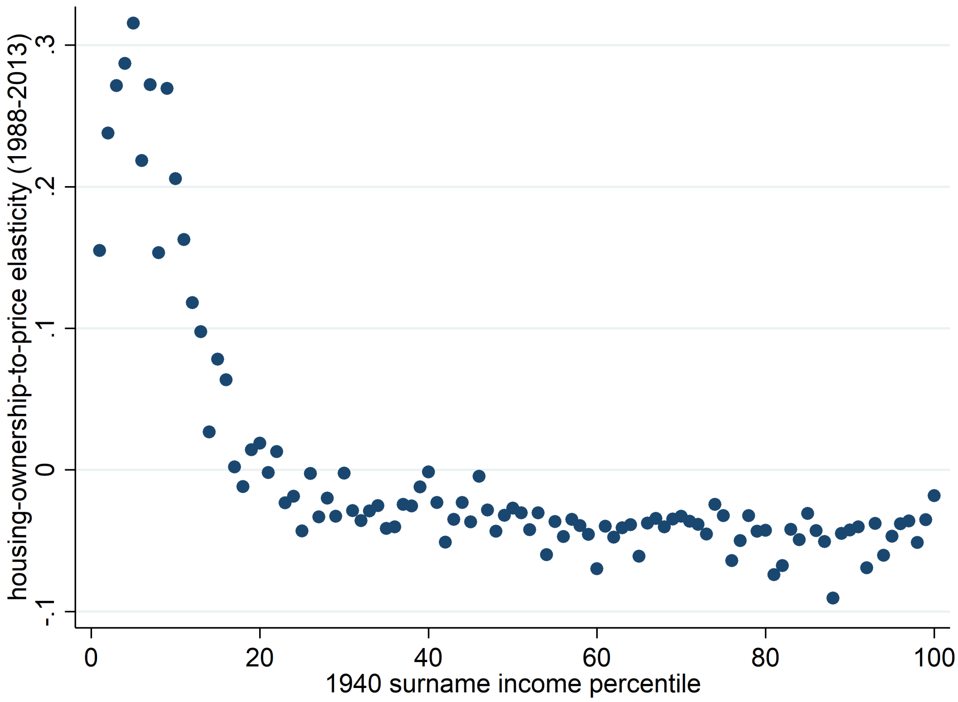

(c) For 1988-2013 sample: Only 1988-2002

Figure A.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

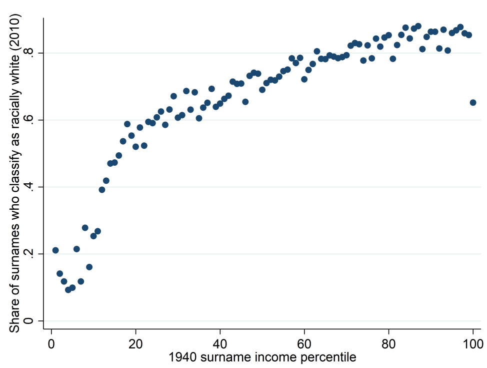

(b) 1998-2013 sample: White share (residual)

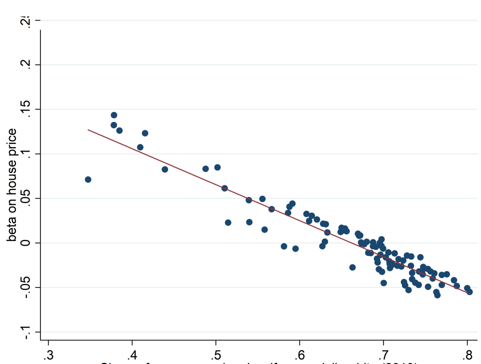

(c) 1998-2013 sample: 1940 income (residual)

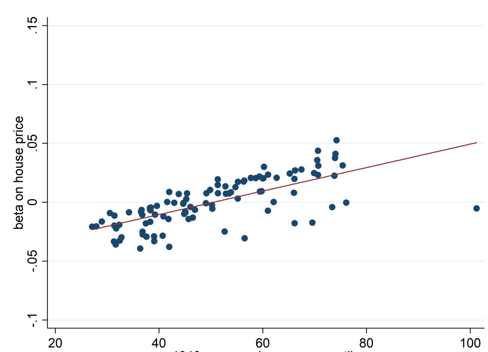

Figure A.3: "Beta" between- and within-racial share (continued)

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

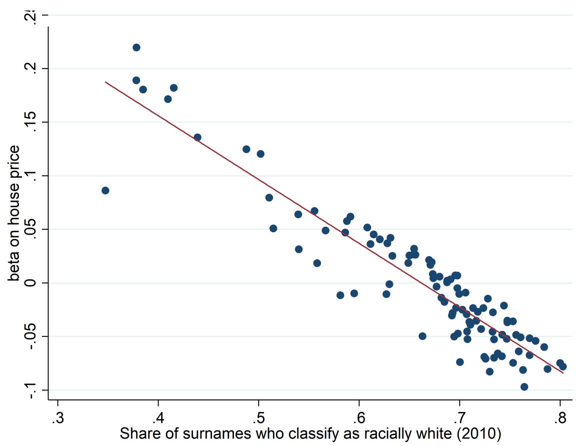

(e) 1988-2013 sample (full): 1940 income (residual)

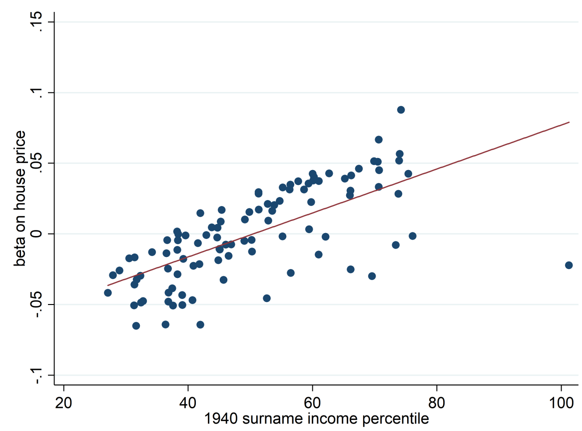

(f) 1988-2013 sample (1988-2002): White share (residual)

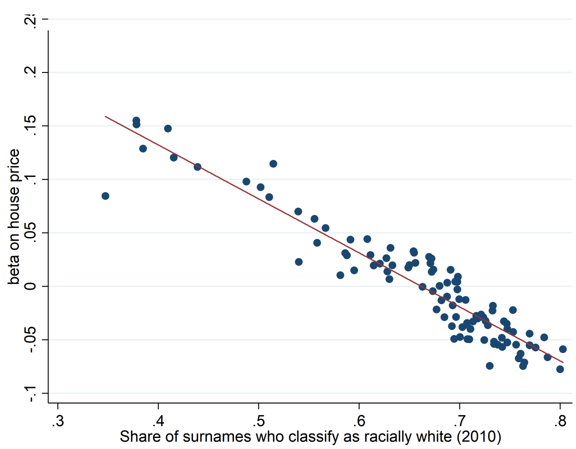

(g) 1988-2013 sample (1988-2002): 1940 income (residual)

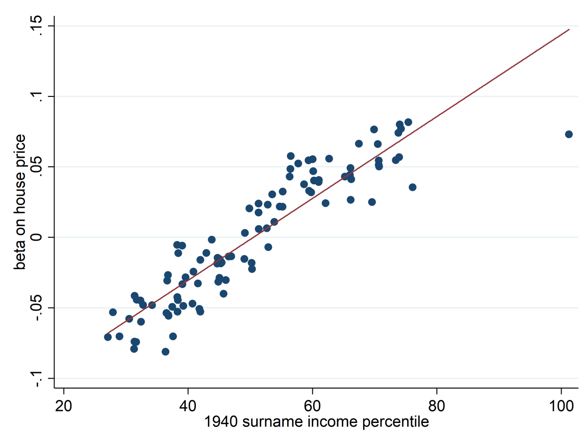

Figure A.4: Conversion to return differential

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 [40]: $var\left(E_t dR_t\right) \approx (0.0546)^2$, $\tilde{b} \approx 3.8$. 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

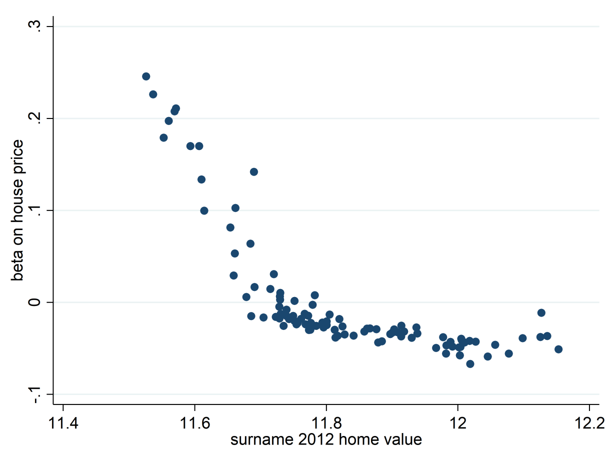

Figure A.4: Conversion to return differential (continued)

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

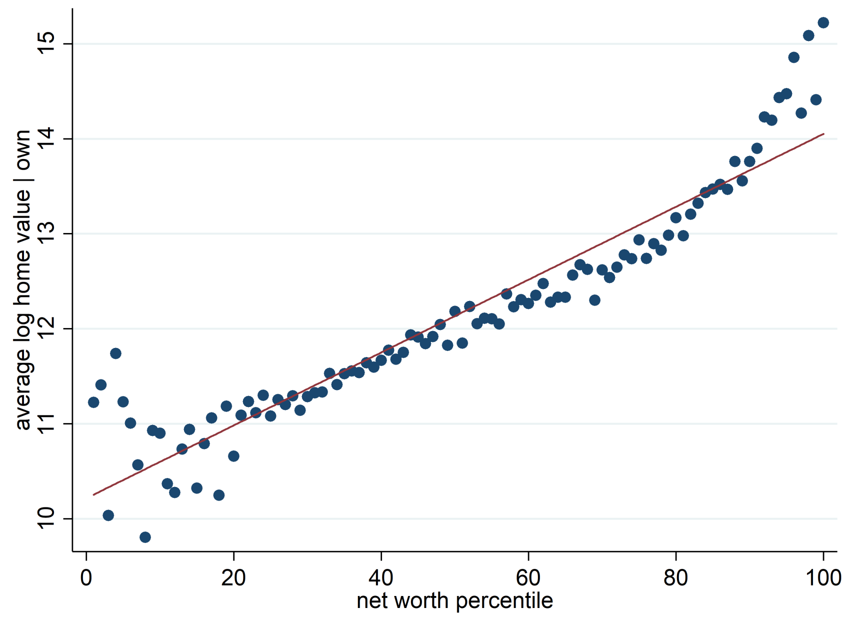

(c) Return differential vs. wealth percentile

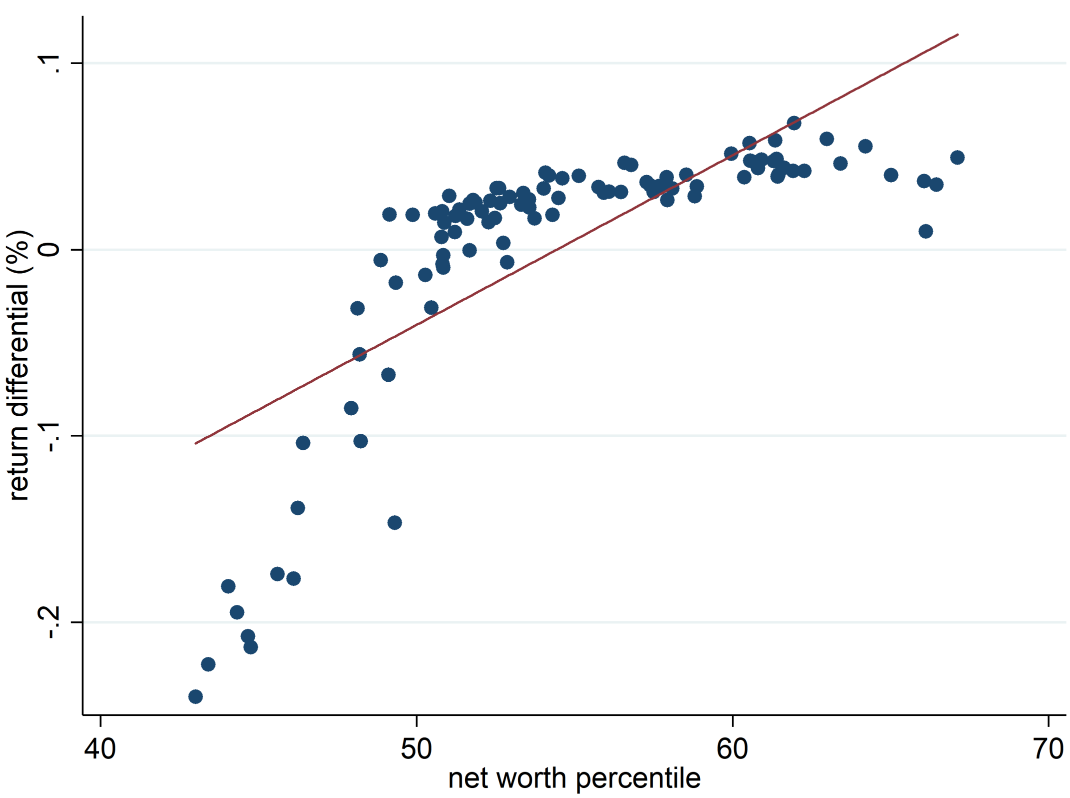

Figure A.5: Elasticity gradient by local cyclicality

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.

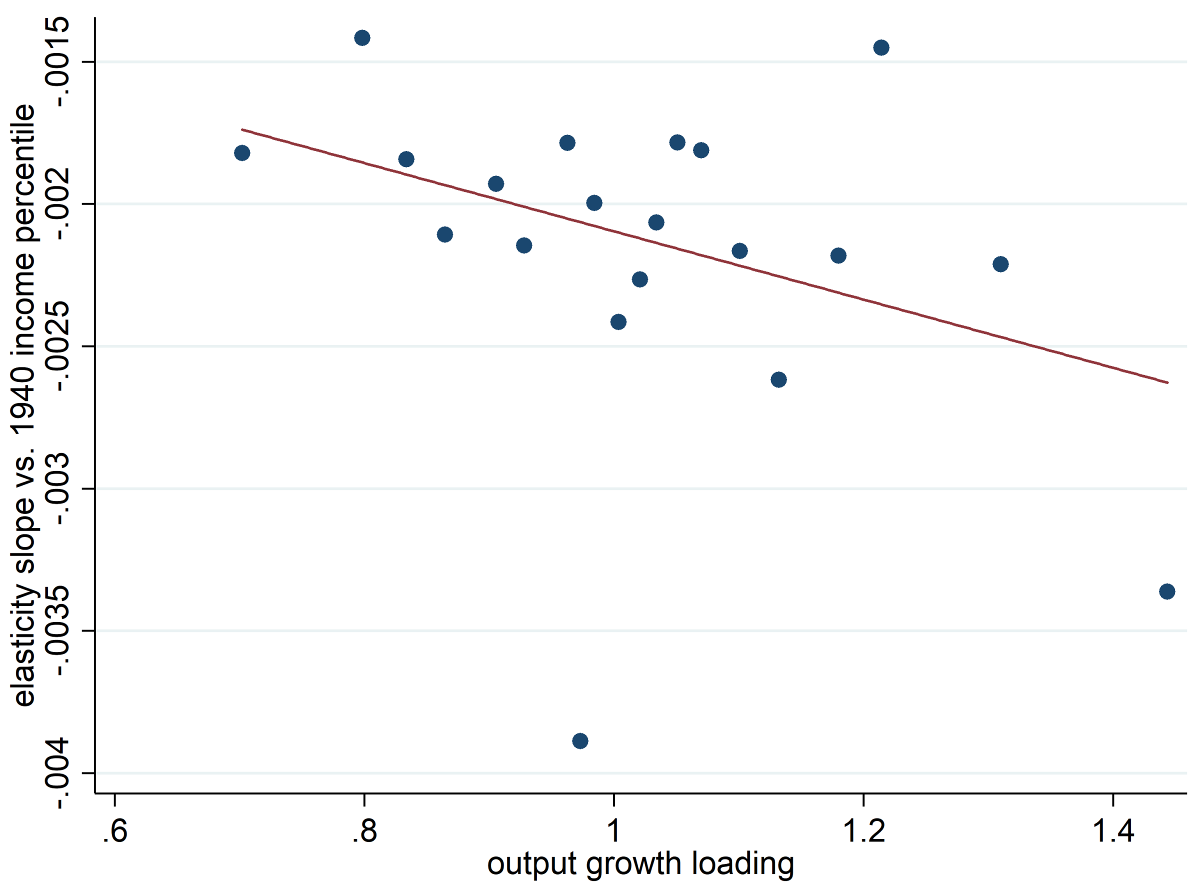

Figure A.6: Wealth inequality level: Coefficient of variation from zip code data

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 [123], 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)

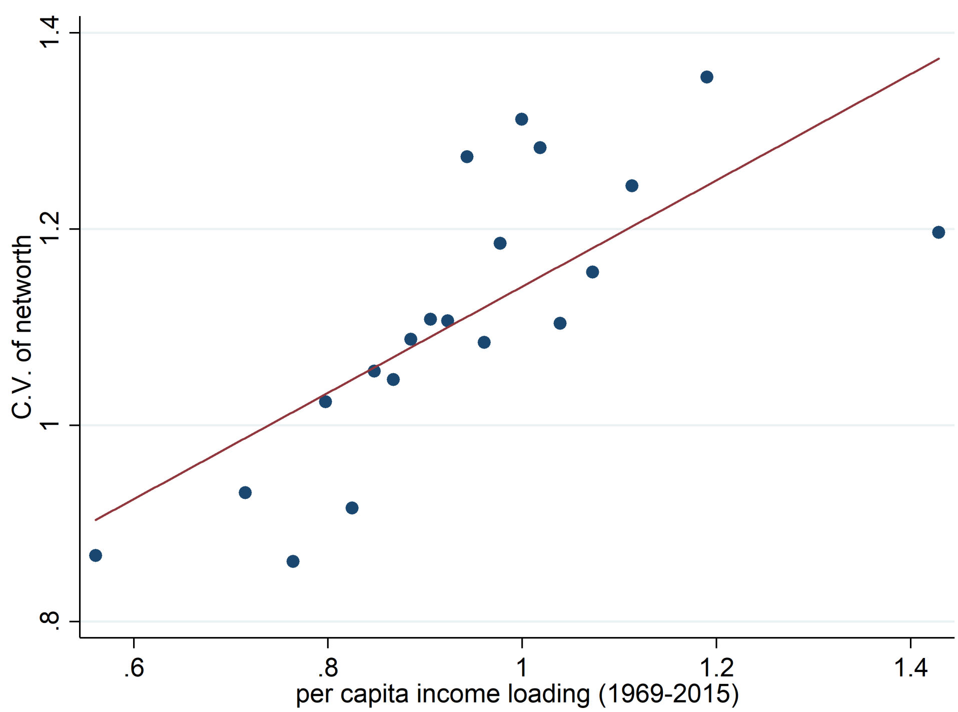

Figure A.6: Wealth inequality level: Coefficient of variation from zip code data (continued)

(b) Asset coefficient of variation (CBSA)

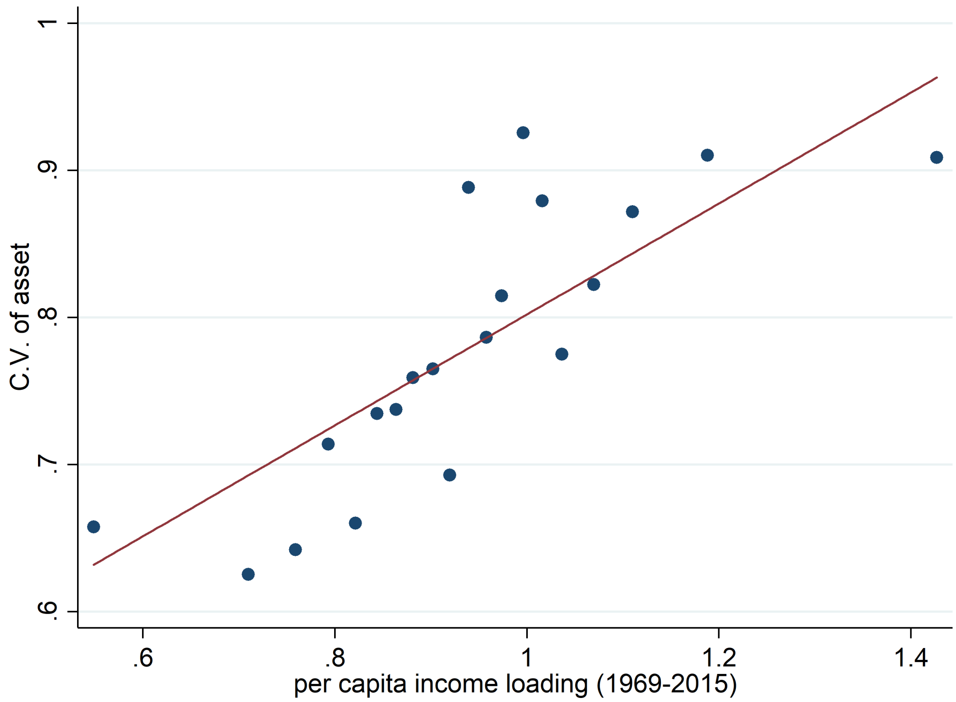

(c) Household income coefficient of variation: IPUMS vs. IRS

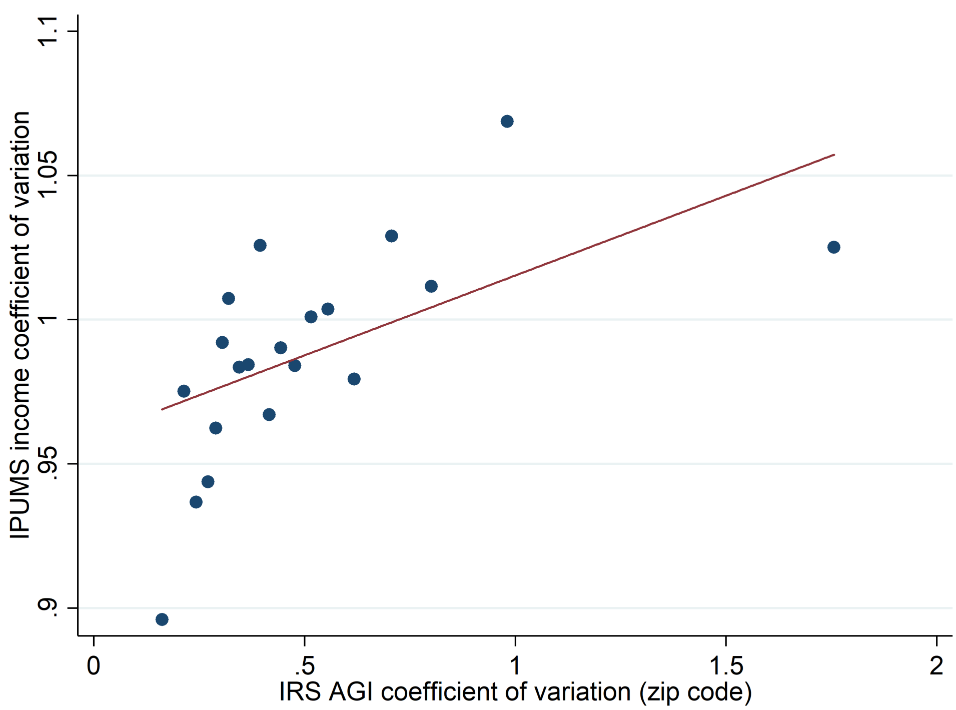

Table A.1: Validation of surname-based historical income against Census 2000

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. 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. 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 (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 A.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 A.3: Level of wealth inequality

For county c in CBSA m,

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

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 [123], 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 A.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 [123], 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

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

A.2 Additional Theoretical Results

A.2.1 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.[1]

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:

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

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

1. 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 ([72]), (2) credit-supply fluctuations disproportionately affect lower-income households ([102]; [103]), (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.

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 $Q_t > 0$ 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 [139] and [28],[2] 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.[3]

A reduction in ξ_t from the relaxed borrowing constraint acts like a reduction in the risk-free rate for portfolio choice.[4] 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

2. 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 $W > 0$.

3. 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.

4. 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, $\phi < 1$.

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 ([141]; [13]). 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

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

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 [139]. With this simplification, household consumption c_t and portfolio share in risky asset θ_t are given by

$$\begin{aligned} c_t - y_t &= b_0 + b_1\,(w_t - y_t) \\ \theta_t &= \frac{\mu - \left(r_f + \xi\right) + \frac{1}{2}\sigma_u^2 - (1-b_1)\,\sigma_{yu}}{b_1 \sigma_u^2} \end{aligned}$$

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.[5]

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

5. 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 ([97]). This force is specific to housing as an illiquid asset (i.e., asset with large transaction costs).

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.[6][7]

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.

6. 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.

7. This implication is also seen in [86], 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.

A.3 Derivations

A.3.1 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

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

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.

Back-of-envelope

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

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

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.

A.3.2 Micro-foundation

Then, log Euler equations for the two assets,

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

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

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

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

Euler with portfolio return:

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

Simplifications: Assume a fixed $\bar{\xi}$ and stationary environment with fixed expected return as in [139]. Comparative static on $\bar{\xi}$. (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:

$$c_t - y_t = b_0 + b_1 (w_t - y_t)$$

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

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

$$\begin{aligned}&b_0+b_1k+(1-b_1)\,g+b_1E\left[r_{p,t+1}\right]-(b_1\rho_c+1)\,b_0+\left(b_1\rho_w-b_1^2\rho_c-b_1\right)(w_t-y_t)\\&=\log\beta+(1-\theta)\,\xi+E\left[r_{p,t+1}\right]+\frac{1}{2}Var\left(r_{p,t+1}-(c_{t+1}-c_t)\right)\end{aligned}$$

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.

A.4 Additional Figures and Tables

Figure A.7: 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

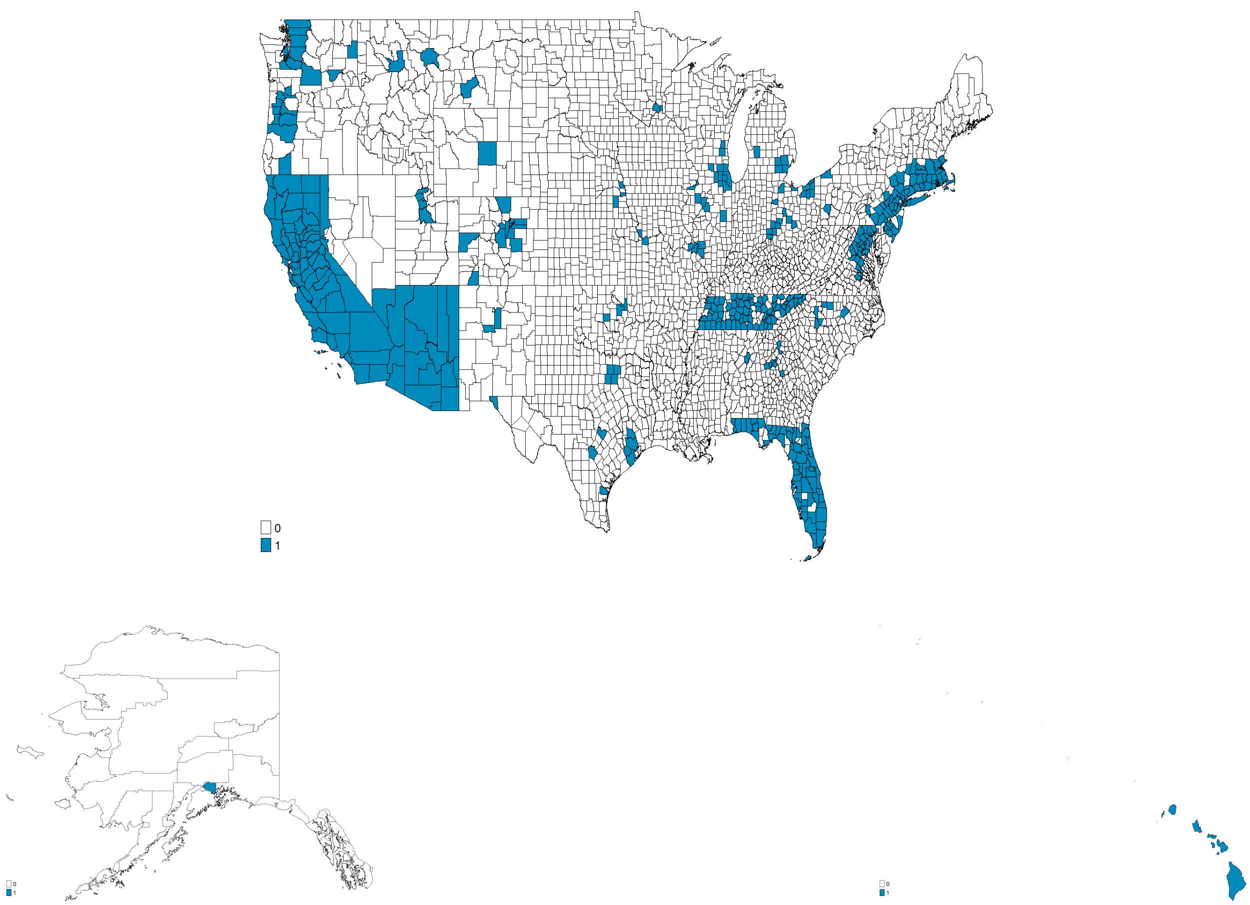

Figure A.7: CoreLogic samples (continued)

(b) Counties in the 1988-2013 sample

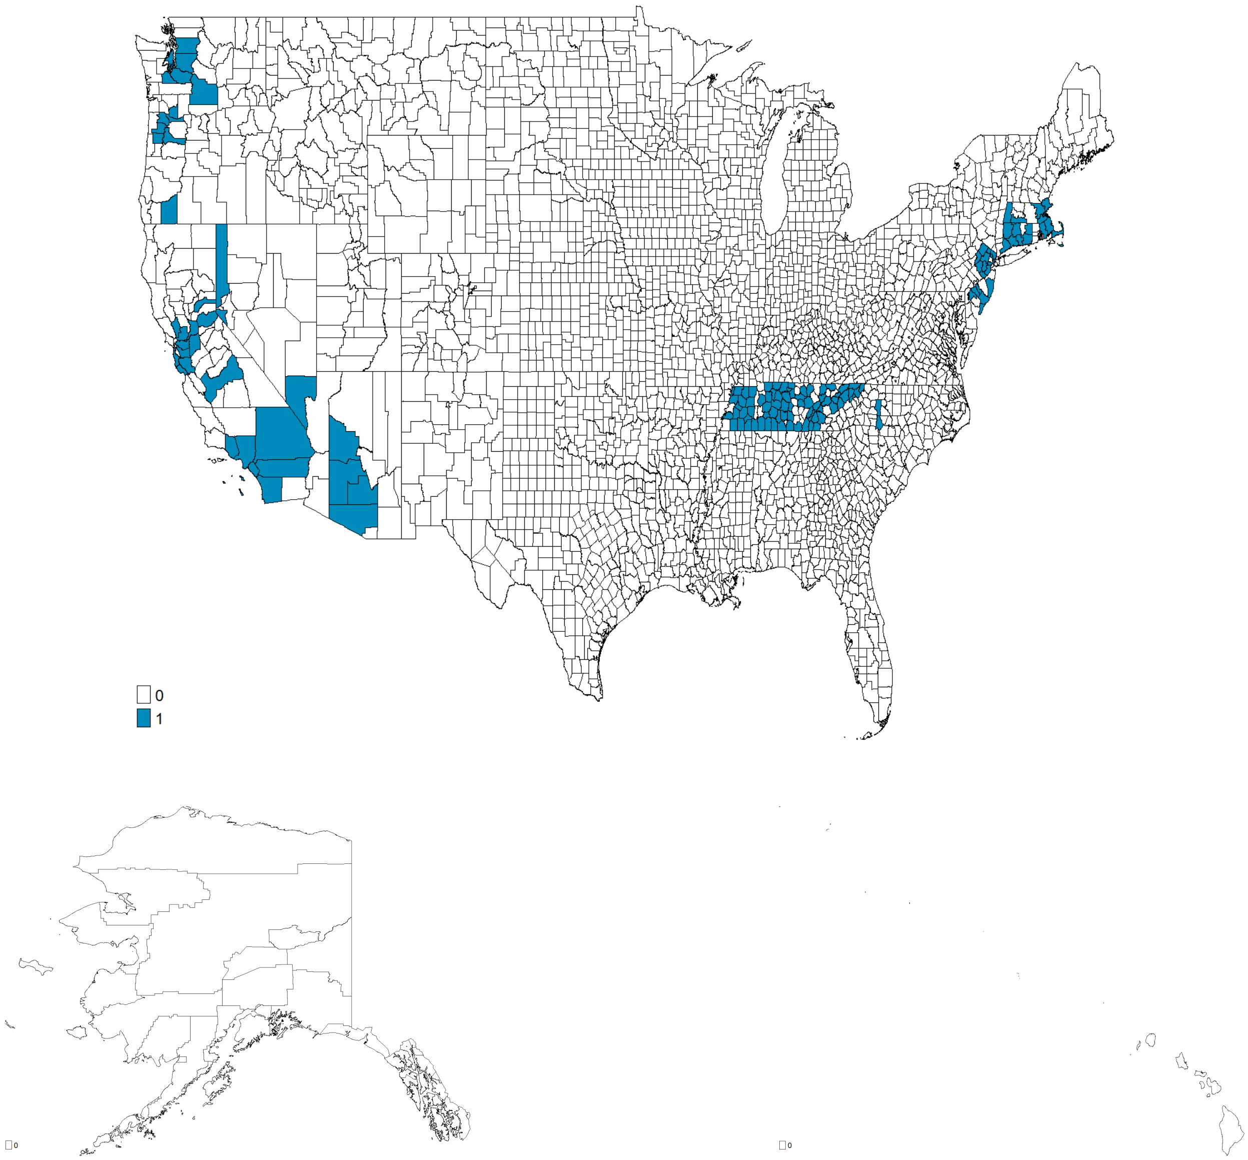

Figure A.7: CoreLogic samples (continued)

(c) Fraction of the US population included

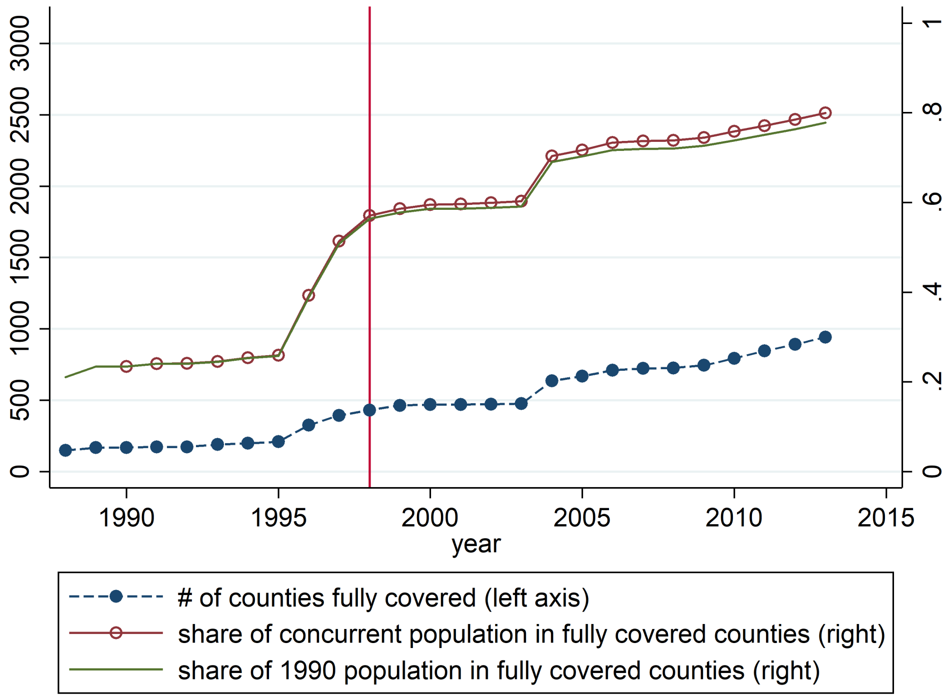

(d) House price of in-sample counties for 1998-2013

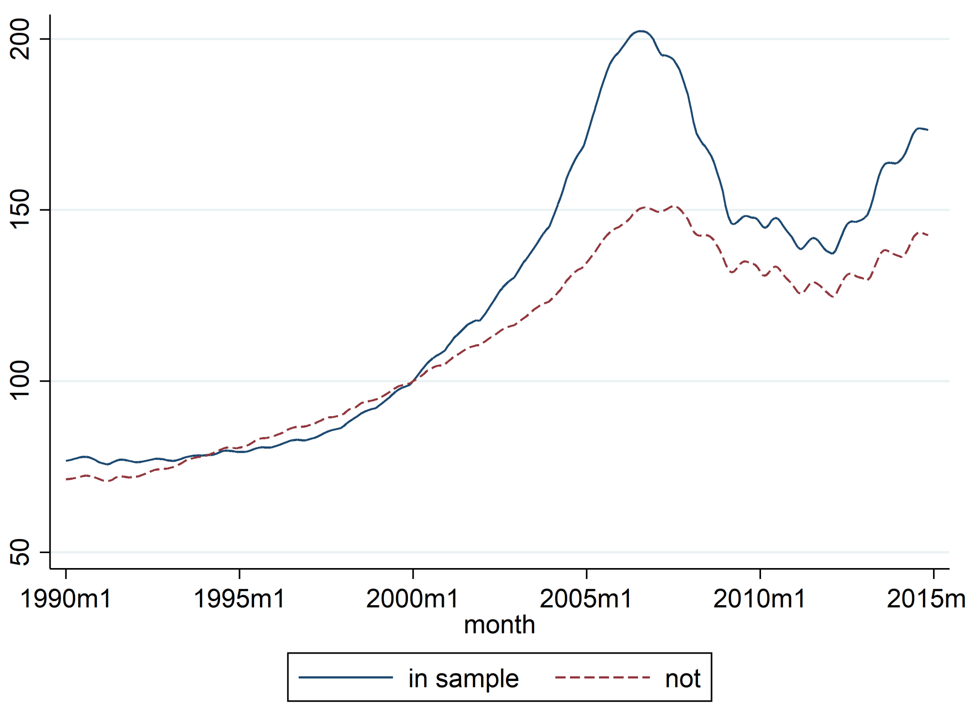

Figure A.7: CoreLogic samples (continued)

(e) Aggregate (Fed)

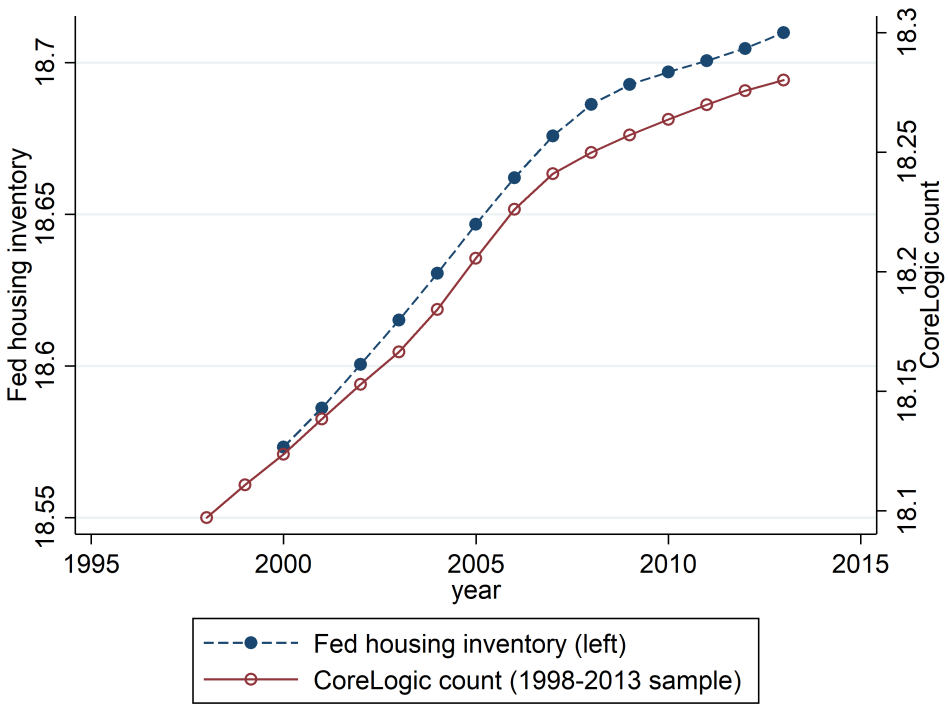

(f) Aggregate (Census)

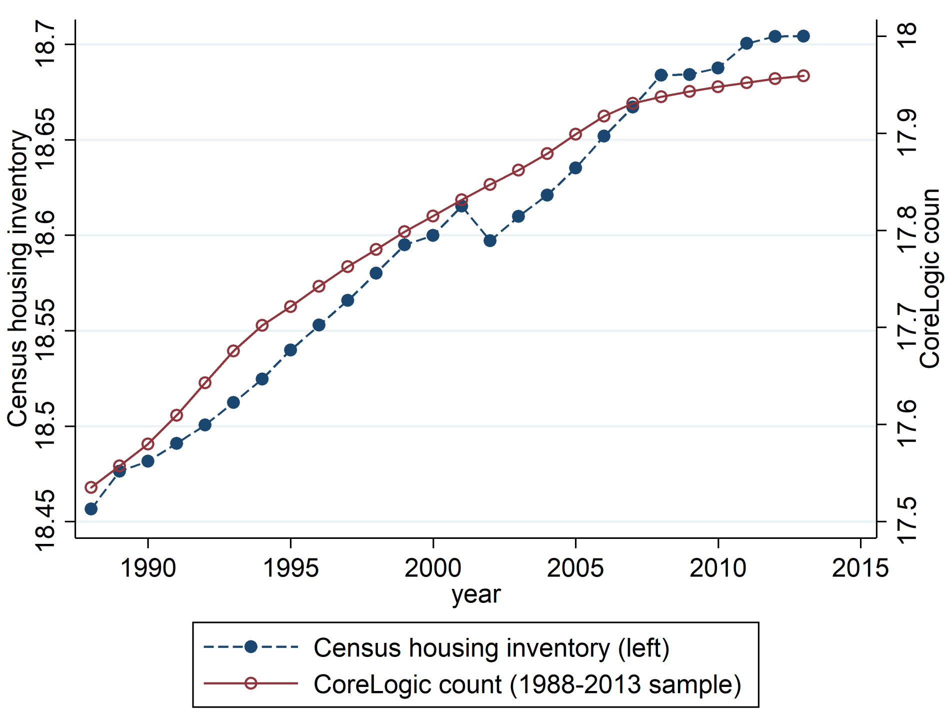

Figure A.8: Wealth inequality vs. income inequality

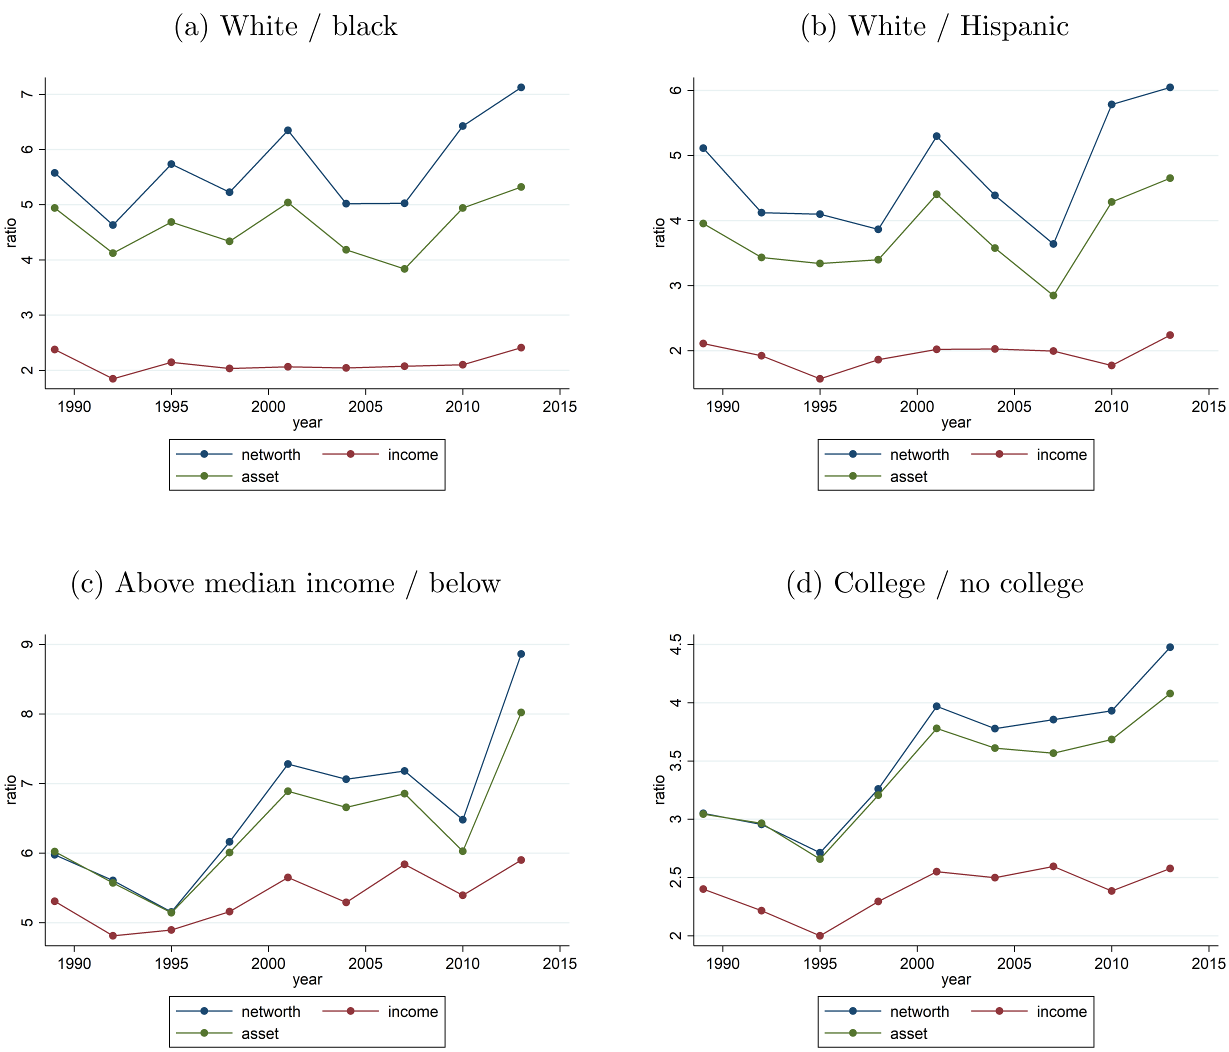

Figure A.9: Wealth inequality level: Geographical variation

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 [123], 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)

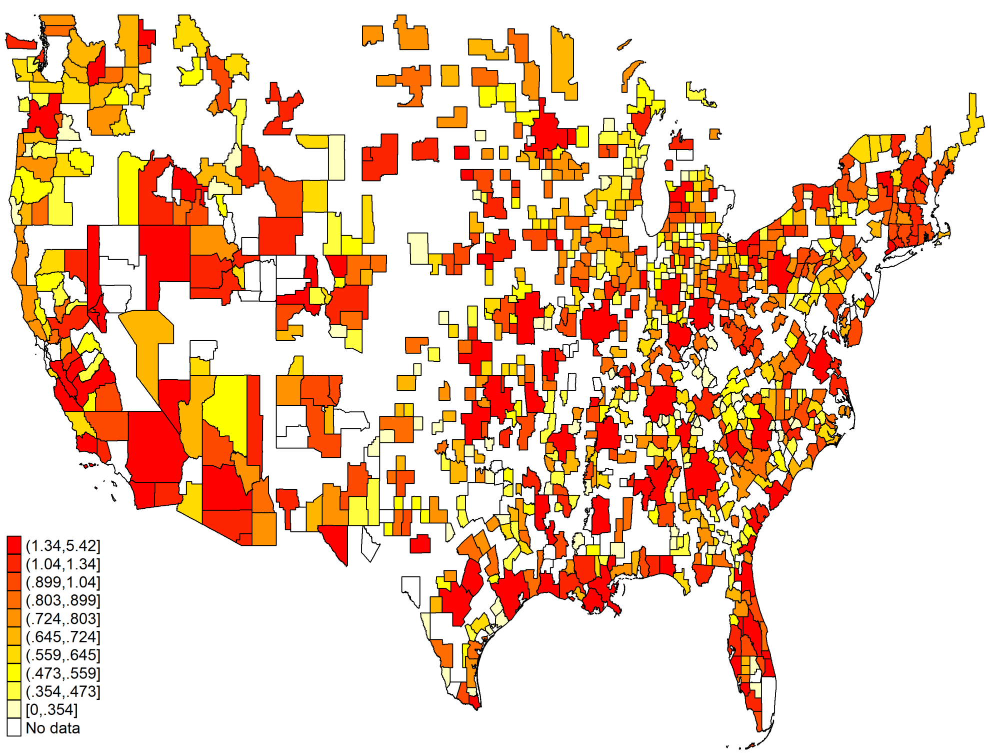

Figure A.9: Wealth inequality level: Geographical variation (continued)

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

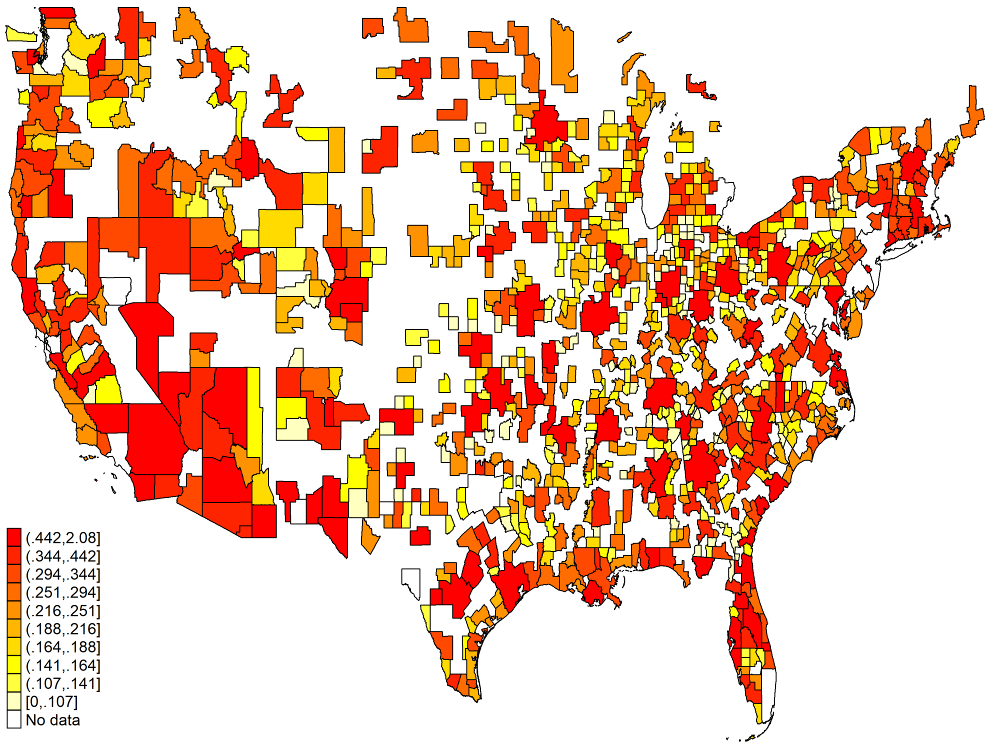

(c) Business cycle loading by county (1969-2015)

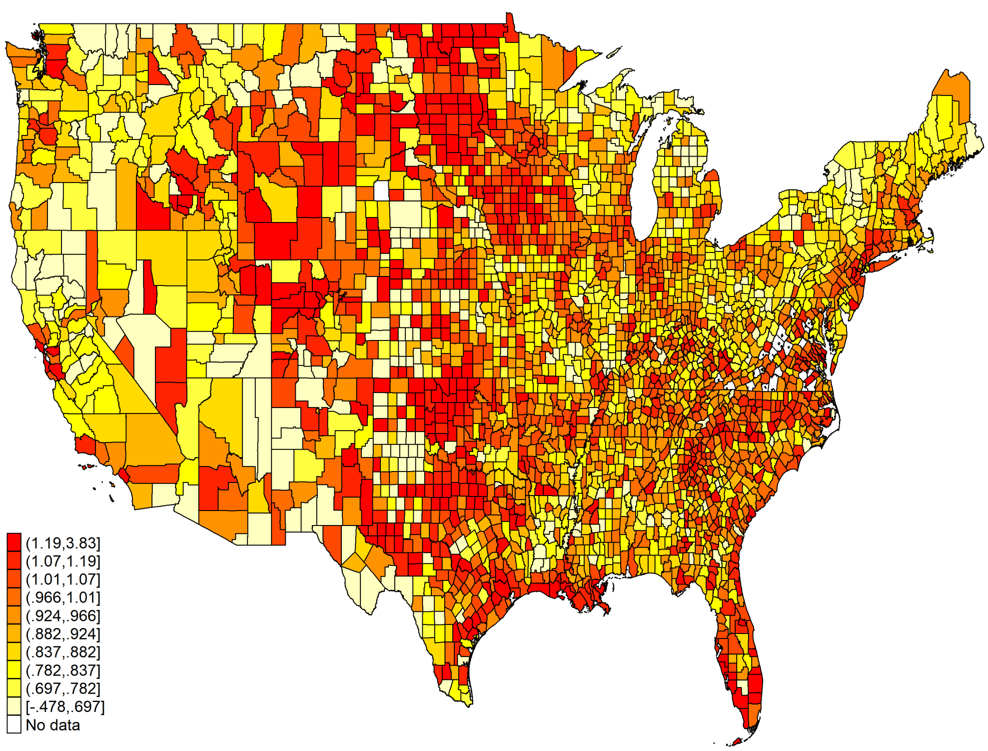

Figure A.10: Wealth vs. income (SCF)

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)

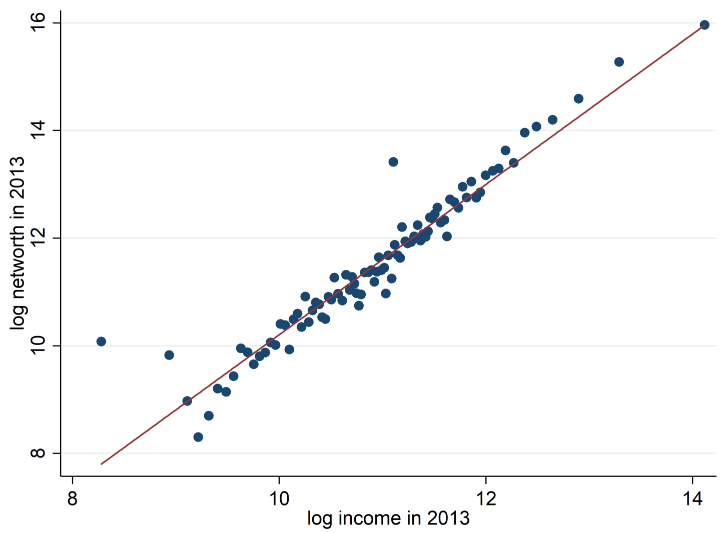

(b) Log asset vs. log income (2013)

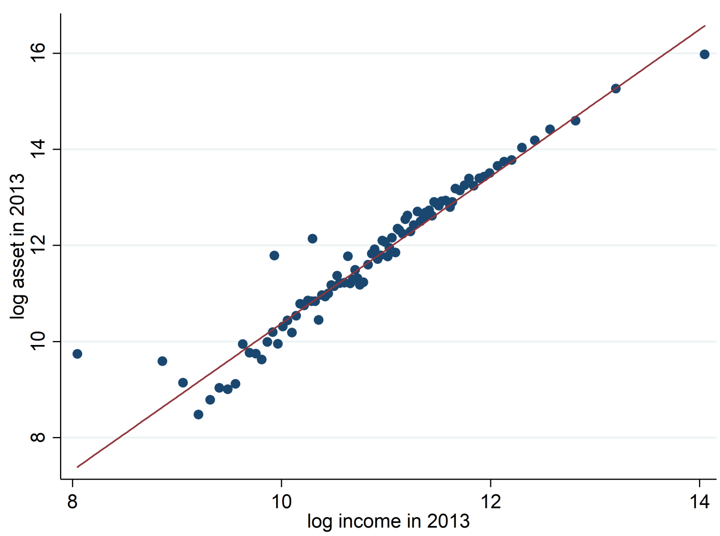

Table A.4: Level of wealth inequality: CBSA coefficient of variation

(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.5: Level of wealth inequality: Zip code wealth-wage elasticity

(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

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

APPENDIX B

RACIAL PREJUDICE IS NOT NORMAL: A COLLAGE OF EMPIRICAL EVIDENCE

B.1 Figures and Tables

Figure B.1: Google searches map

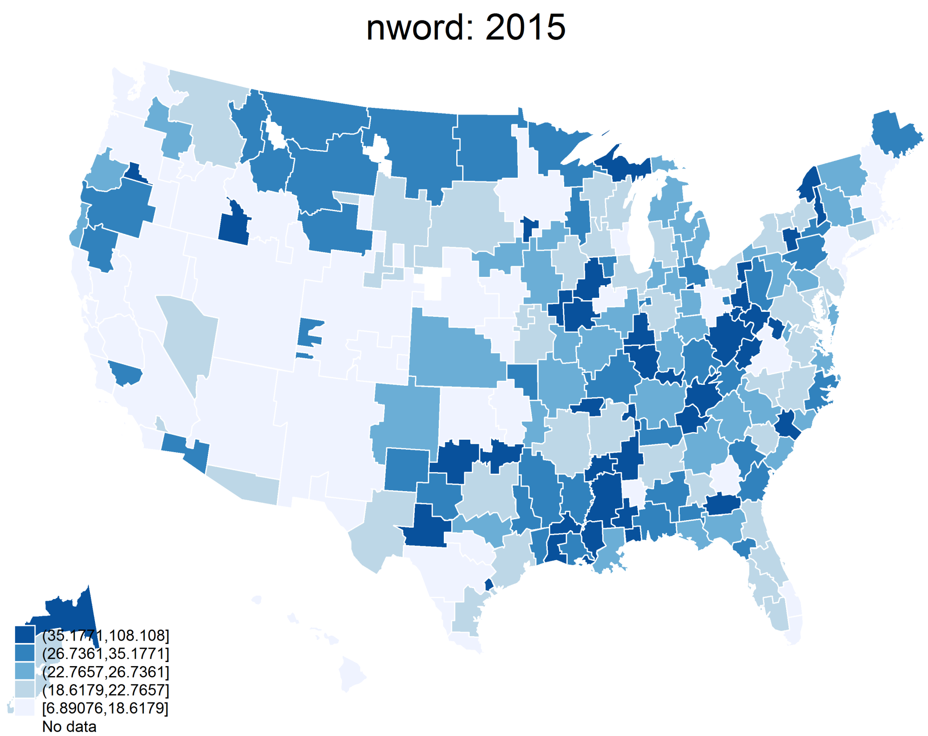

(a) N-word

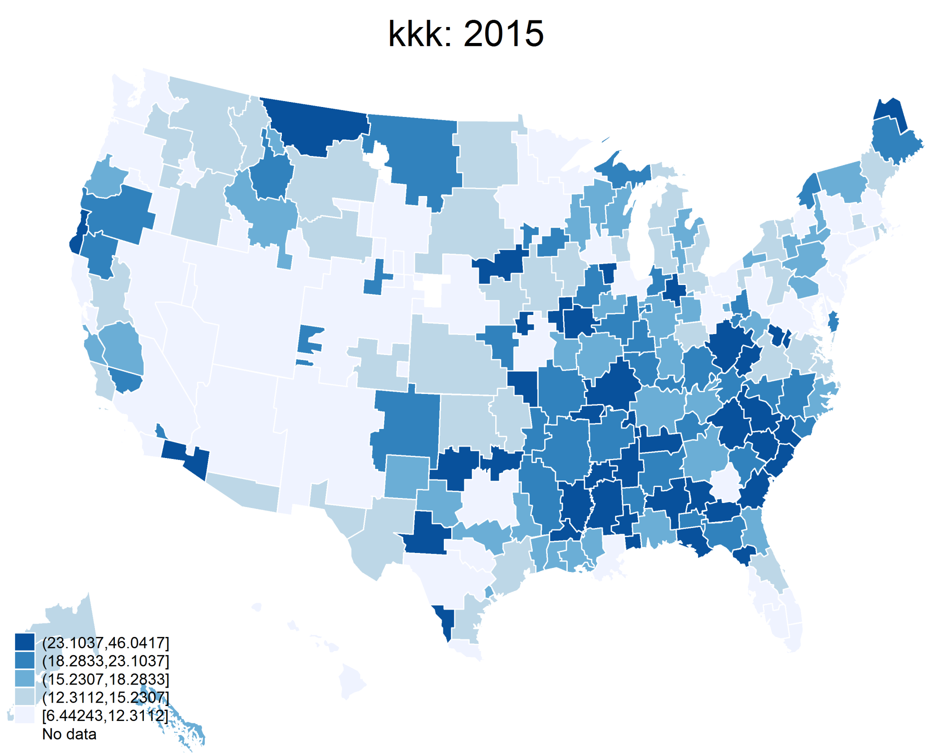

(b) "KKK"

Figure B.2: Google searches

In subplots (a) and (b), points represent within-bin averages (across DMA x year observations) of x and y using bins defined by 100 quantiles sorted on x. The underlying x and y variables are demeaned of their DMA means and year means. Observations are weighted by the size of DMAs.

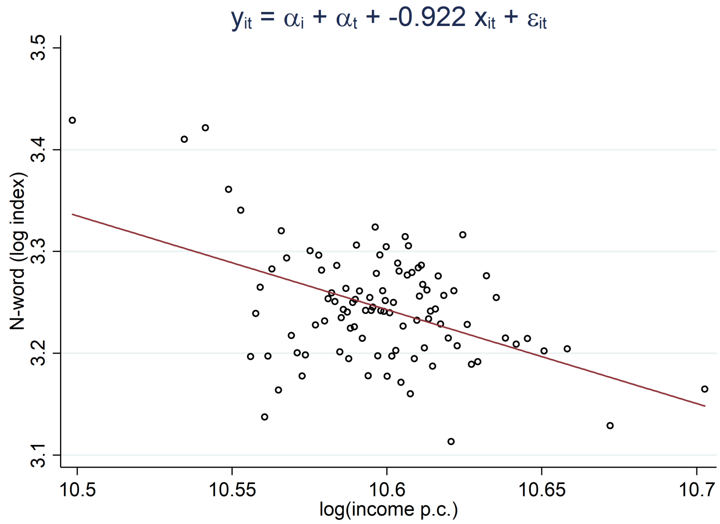

(a) N-word

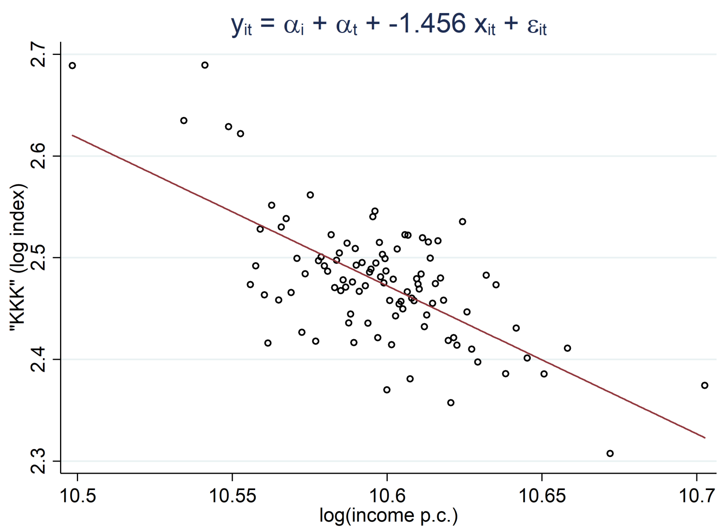

(b) ”KKK”

Figure B.3: Countercyclical prejudice: White-on-black crime vs. white-on-white

The white-on-black crime measures is a ratio of white-on-black crime count against white-on-white crime, in categories of intimidation, simple assault and vandalism. Subplot (b) plots the year fixed effects from a DMA-level panel, weighted by size of DMAs. Fixed effects are presented because the panel is highly unbalanced. In subplot (c), points represent within-bin averages (across DMA x year observations) of x and y using bins defined by 100 quantiles sorted on x. The underlying x and y variables are demeaned of their DMA means and year means. . Observations are weighted by the size of DMAs.

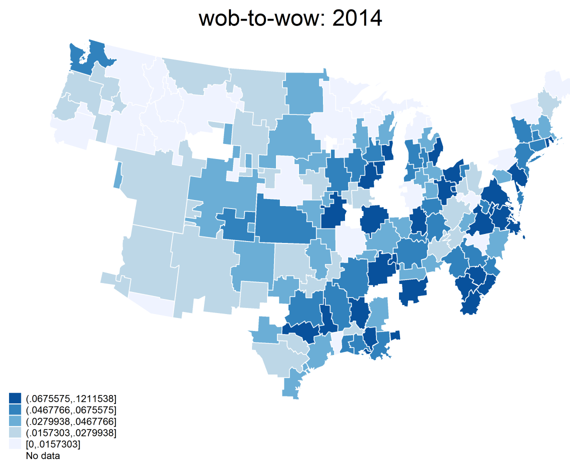

(a) Map

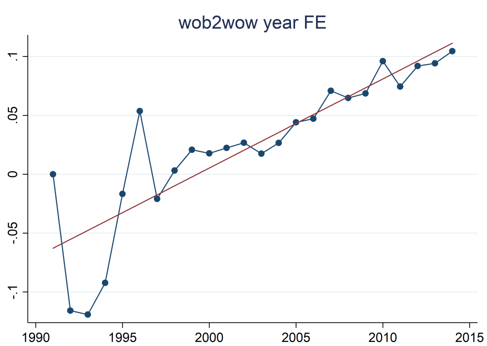

(b) Time series

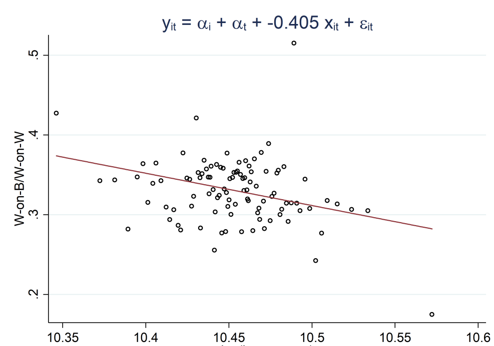

(c) Diff-in-diff

Figure B.4: Countercyclical prejudice: Corporal punishment

Subplot (a) plots the ratio of corporal punishment rate between two races. Subplot (b) plots the change in black-to-white measure (red line in subplot (a)) against the income change of the state from 2006 to 2009. Size of the circle represents the size of the state by population.

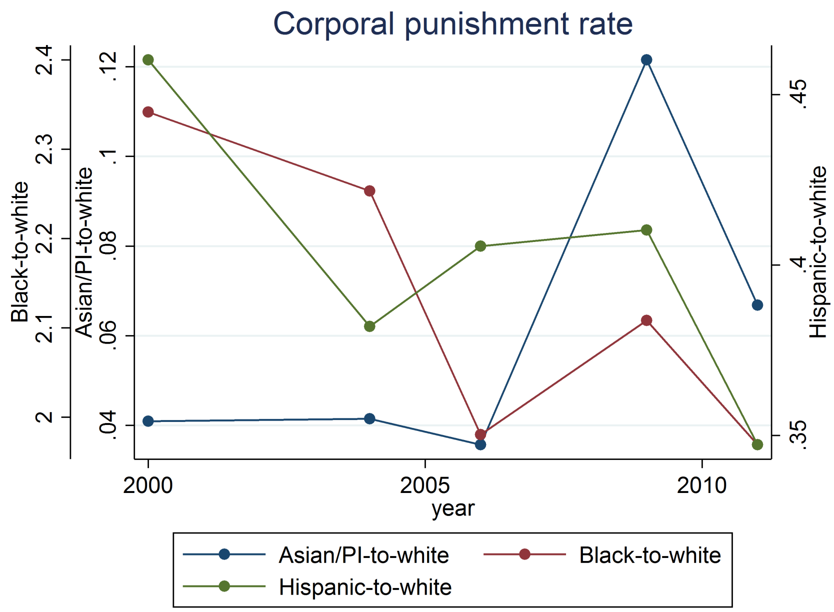

(a) Times series

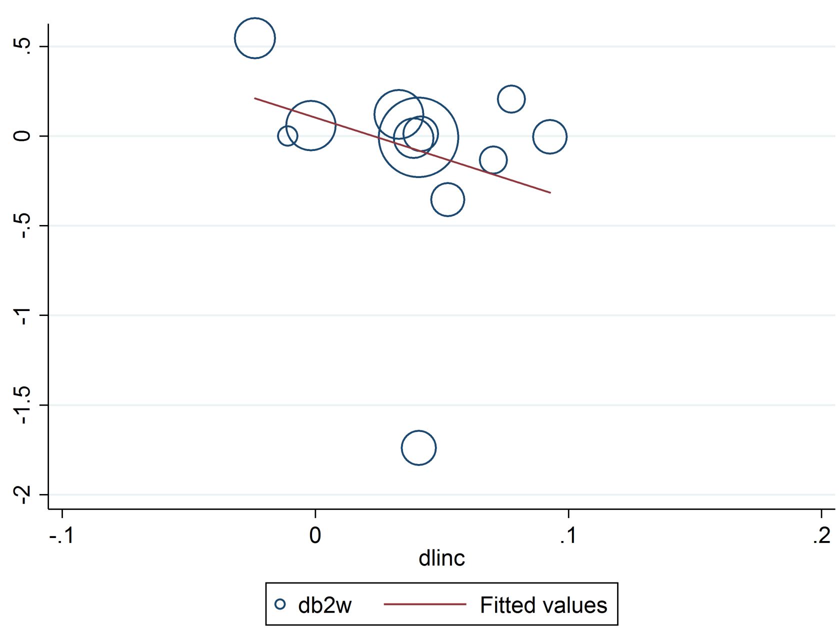

(b) State change 2006-2009

Figure B.5: Countercyclical prejudice: GSS

The census division cross-sectional subplot plots bin scatter of stacked question-division-year level panel, after taking out year FE and question-division FE. The questions used in the second subplot are: (1) How strongly would you object if a family member brought a black friend home for dinner? (2) Do you think blacks should have as good a chance as anyone to get any kind of job, or do you think white people should have the first chance at any kind of job? (3) Do you think there should be laws against marriages between blacks and whites? (4) How would it make you feel if a close relative of yours were planning to marry a black? (5) If a black with the same income and education as you have moved in to your block, would it make any difference to you? (6) Would you vote for a law that says a homeowner can refuse to sell to blacks, or one that says homeowners cannot refuse to sell based on skin color? (7) If your party nominated a black for president, would you vote for him if he were qualified for the job?

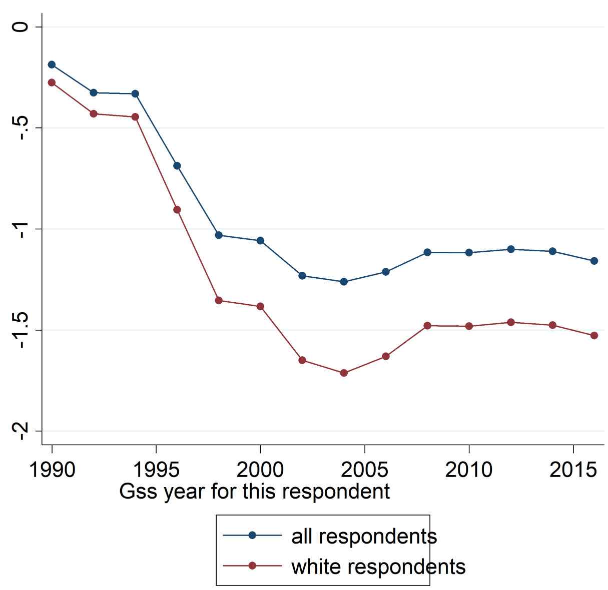

(a) Times series

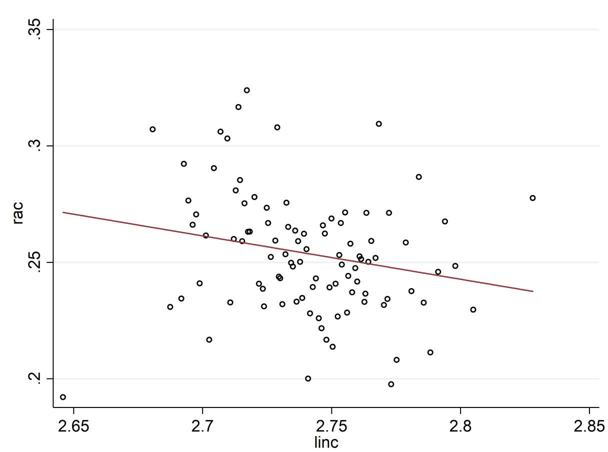

(b) Census division cross-section

Figure B.6: Prejudice IV: Legacy of slavery

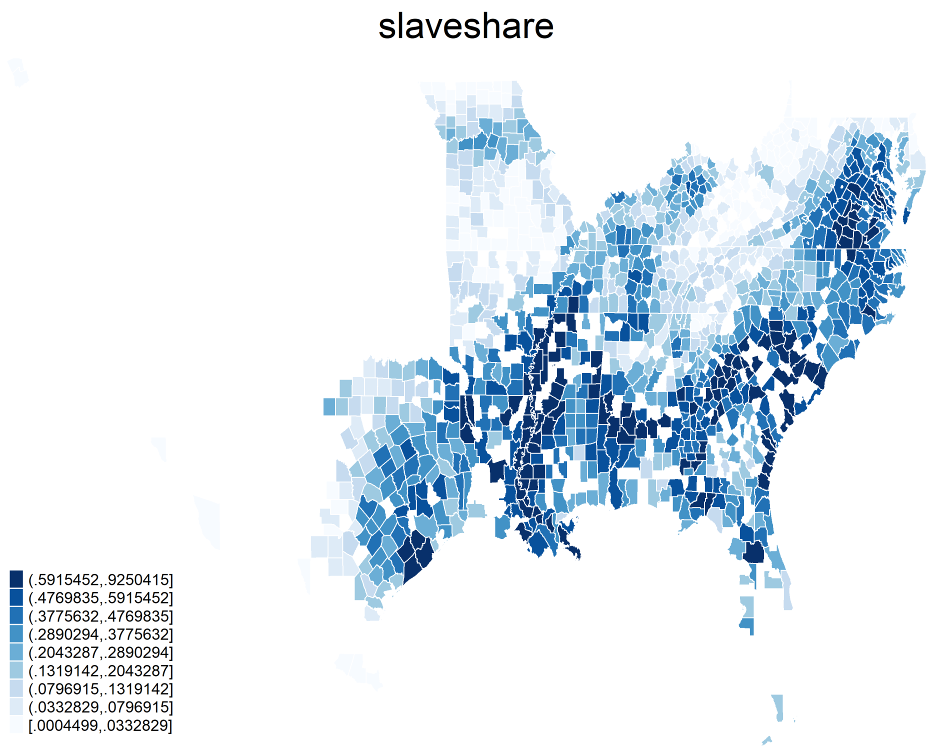

(a) Slave share of population

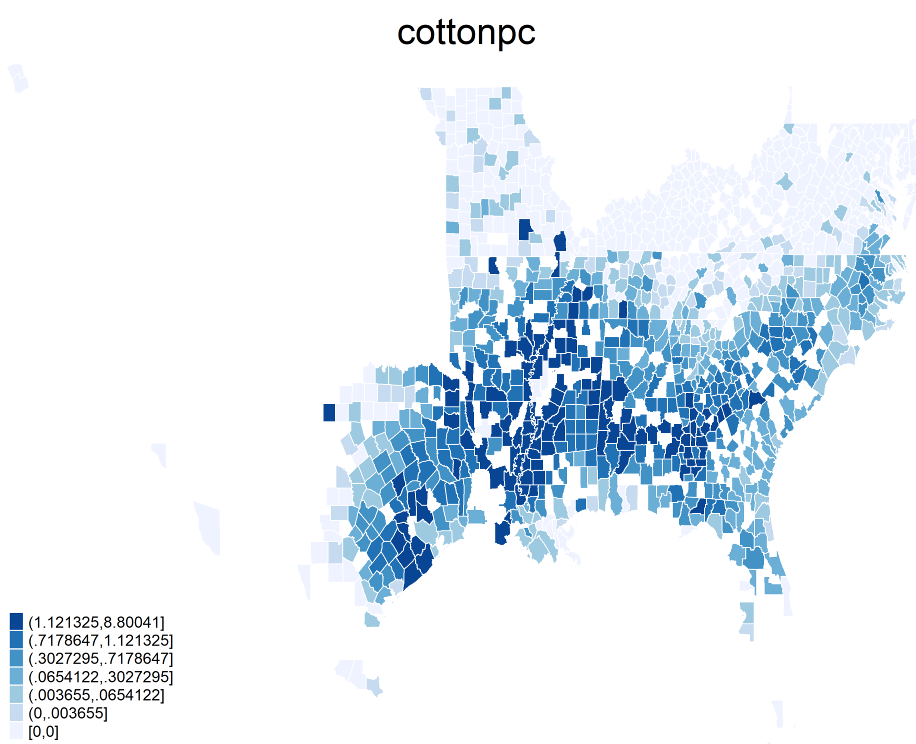

(b) Cotton production per capita

Figure B.7: Real estate stock & net change

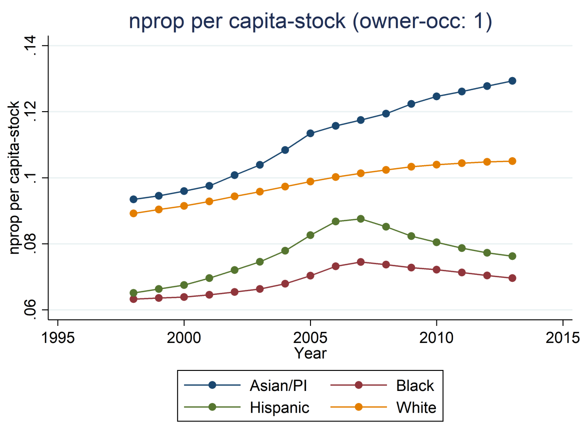

(a) Stock: owner-occupied

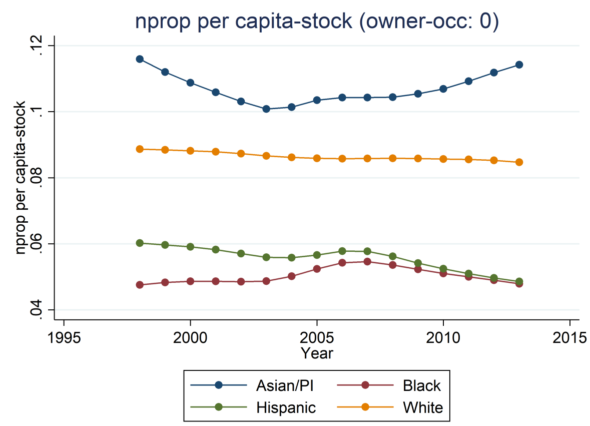

(b) Stock: not owner-occupied

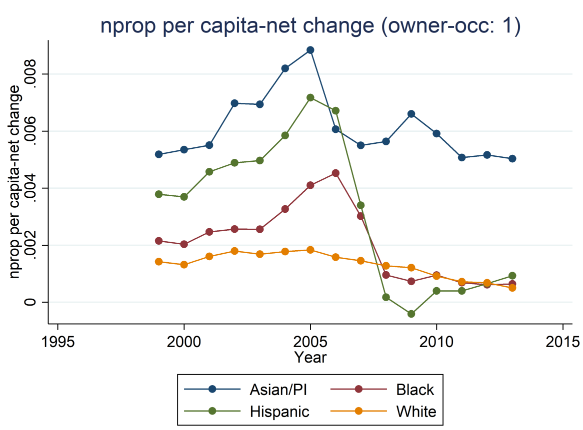

(c) Net change: owner-occupied

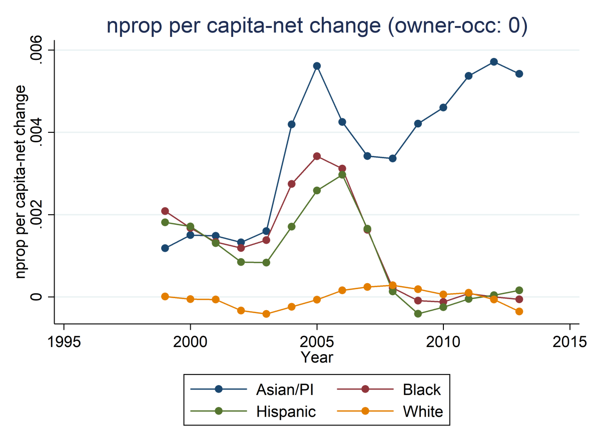

(d) Net change: not owner-occupied

Figure B.8: Percentage change by share of last name

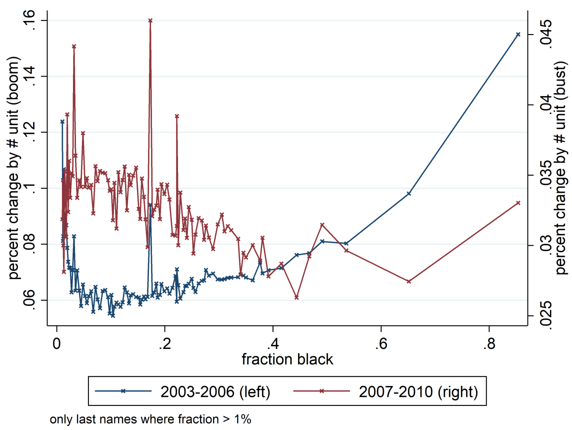

(a) Black: owner-occupied

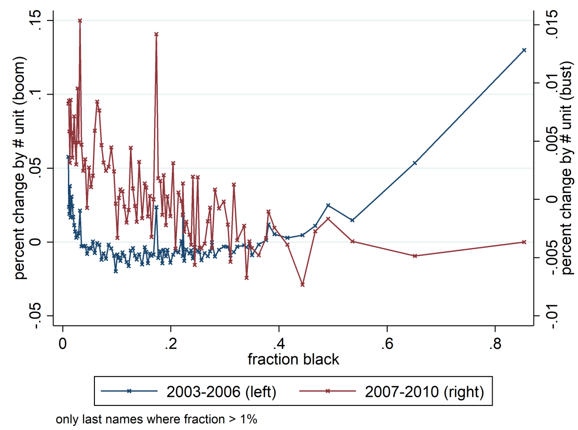

(b) Black: not owner-occupied

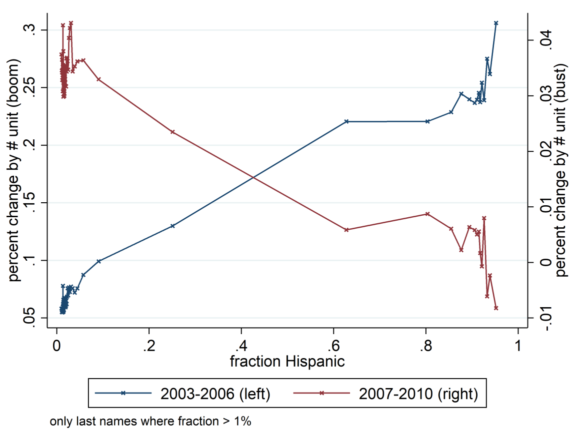

(c) Hispanic: owner-occupied

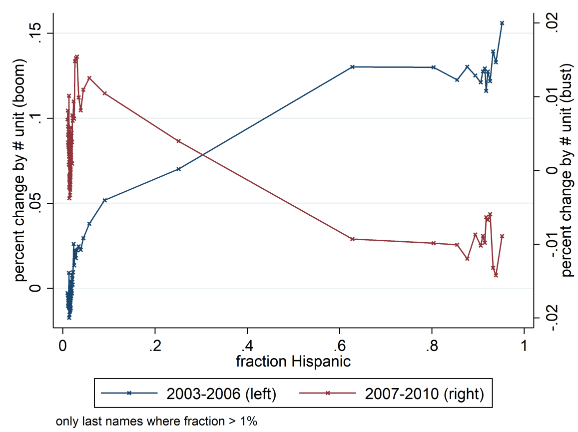

(d) Hispanic: not owner-occupied

Figure B.9: Real estate gross purchase & propensity to sell

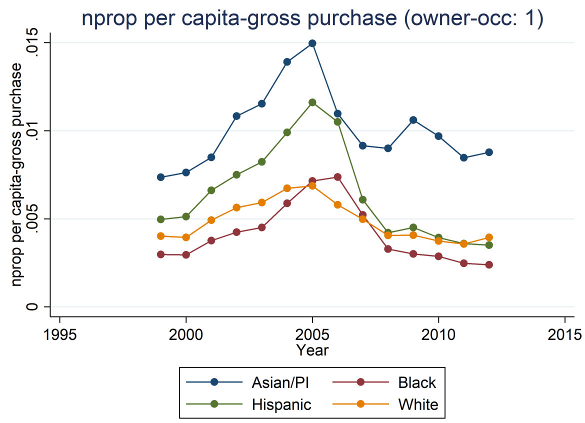

(a) Gross purchase: owner-occupied

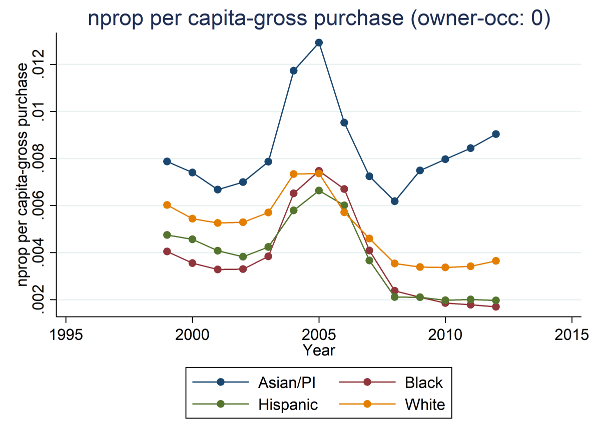

(b) Gross purchase: not owner-occupied

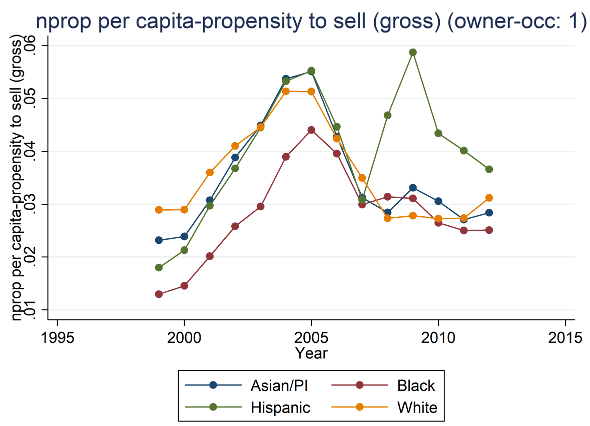

(c) Propensity to sell: owner-occupied

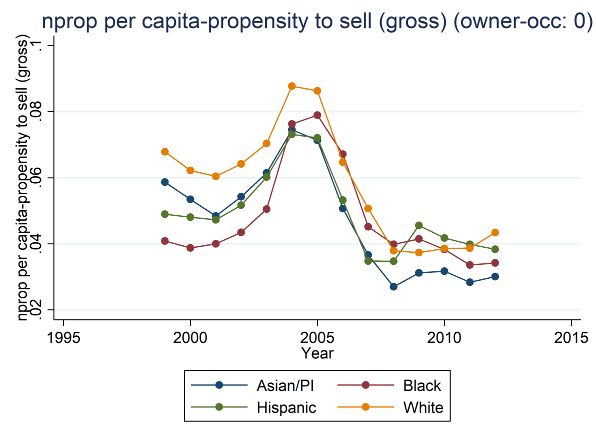

(d) Propensity to sell: not owner-occupied

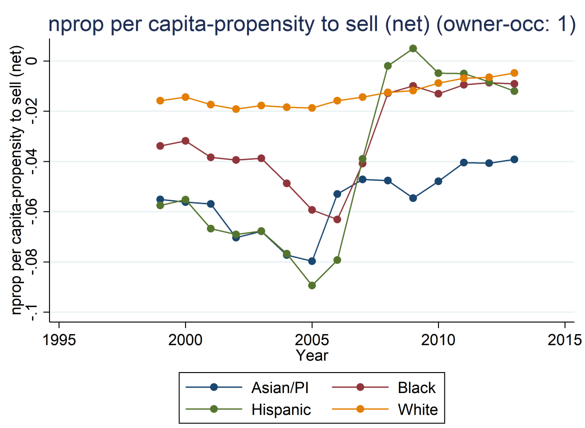

(e) Propensity to decumulate (net): owner-occupied

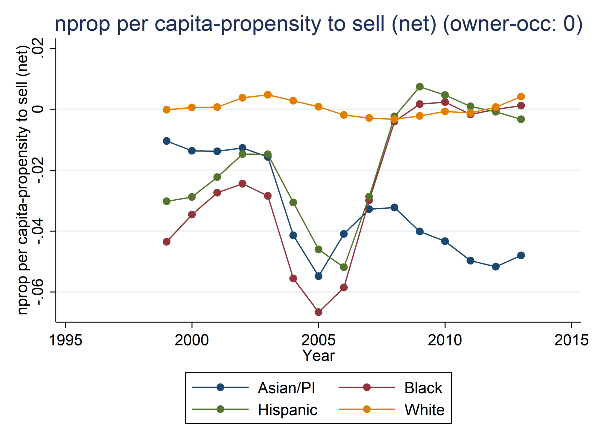

(f) Propensity to decumulate (net): not owner-occupied

Figure B.10: Net change in real estate: Cross-section

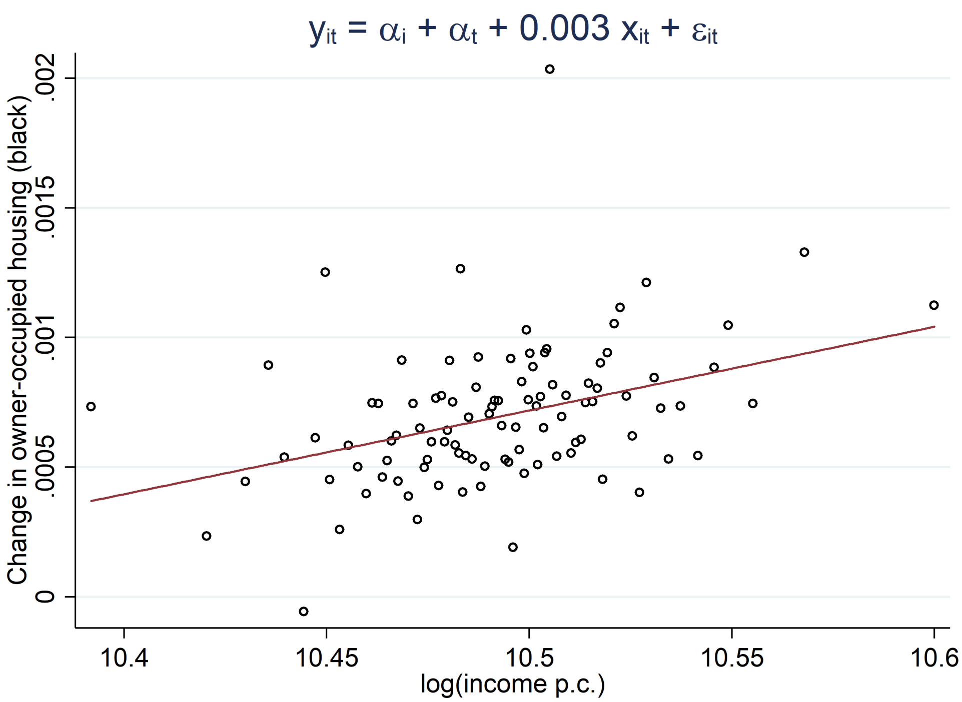

(a) Owner-occupied

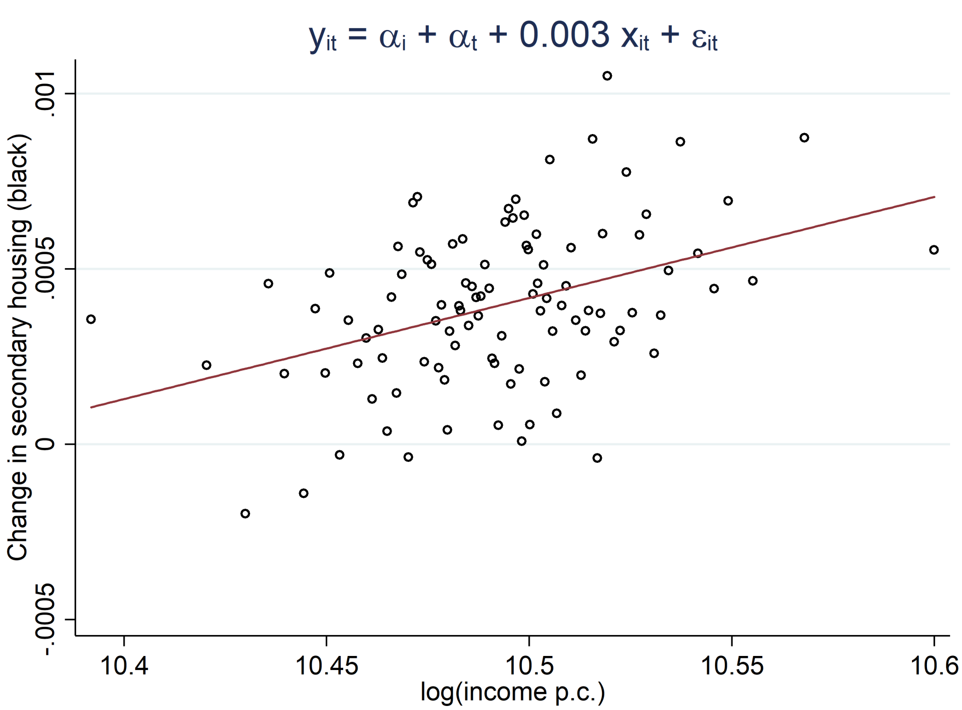

(b) Secondary housing

Figure B.11: Employment: Time series

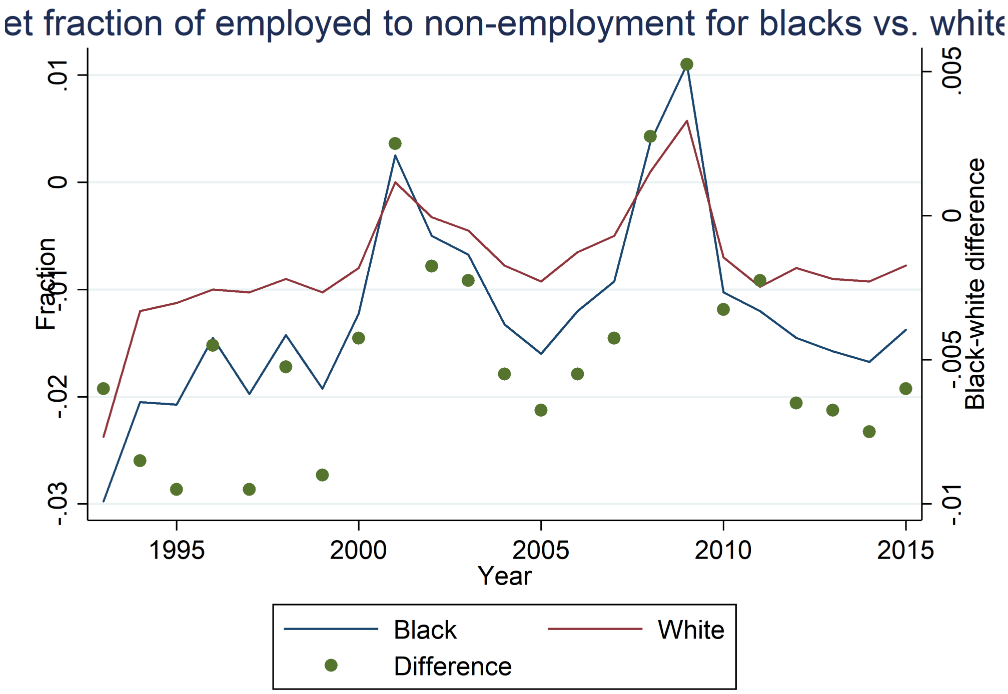

(a) Transition to non-employment

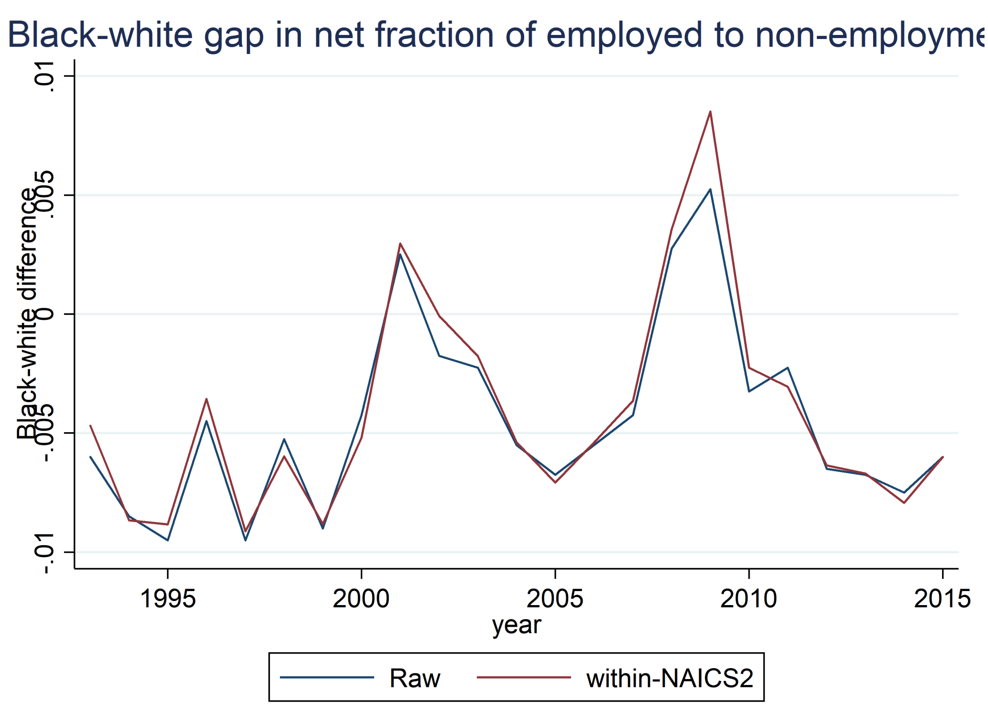

(b) Gap

Figure B.12: Employment: Cross-section

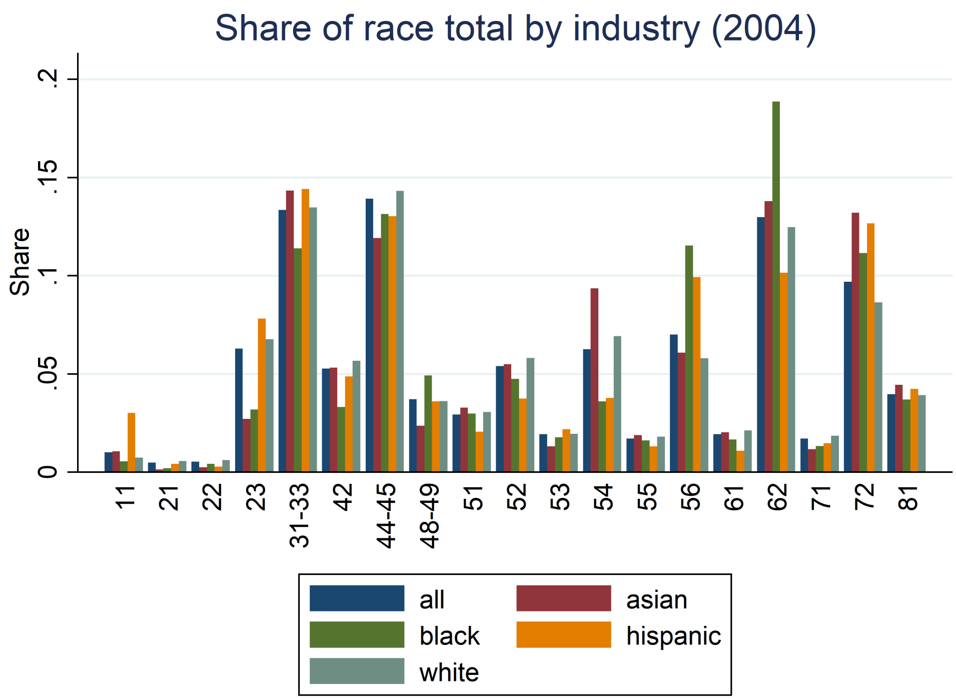

(a) Industry (2-digit NAICS)

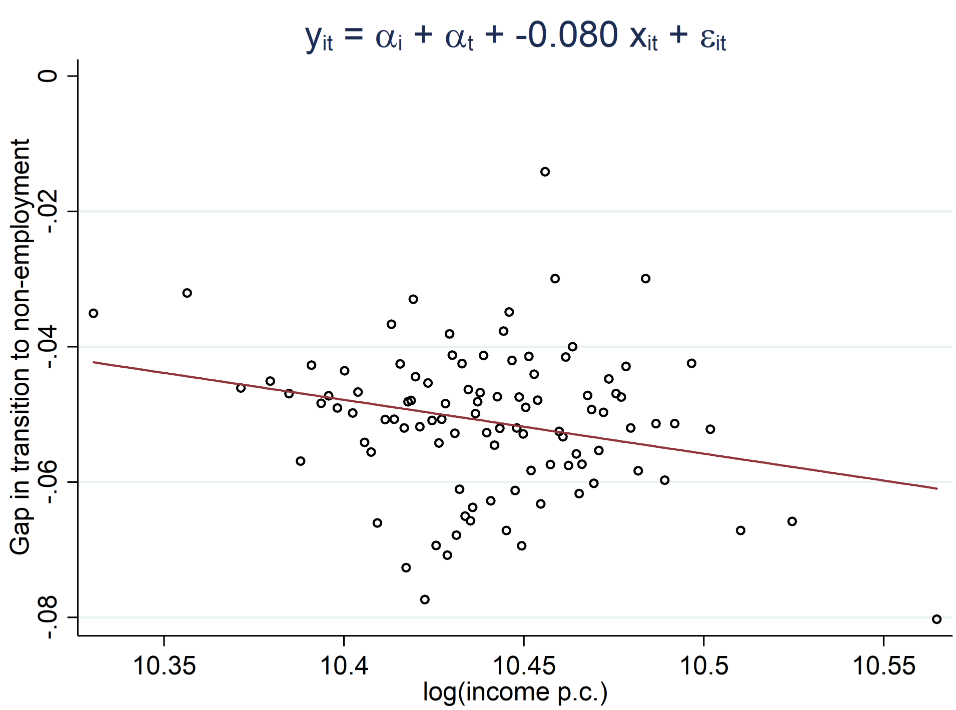

(b) Diff-in-diff

Figure B.13: Mortgage rejection

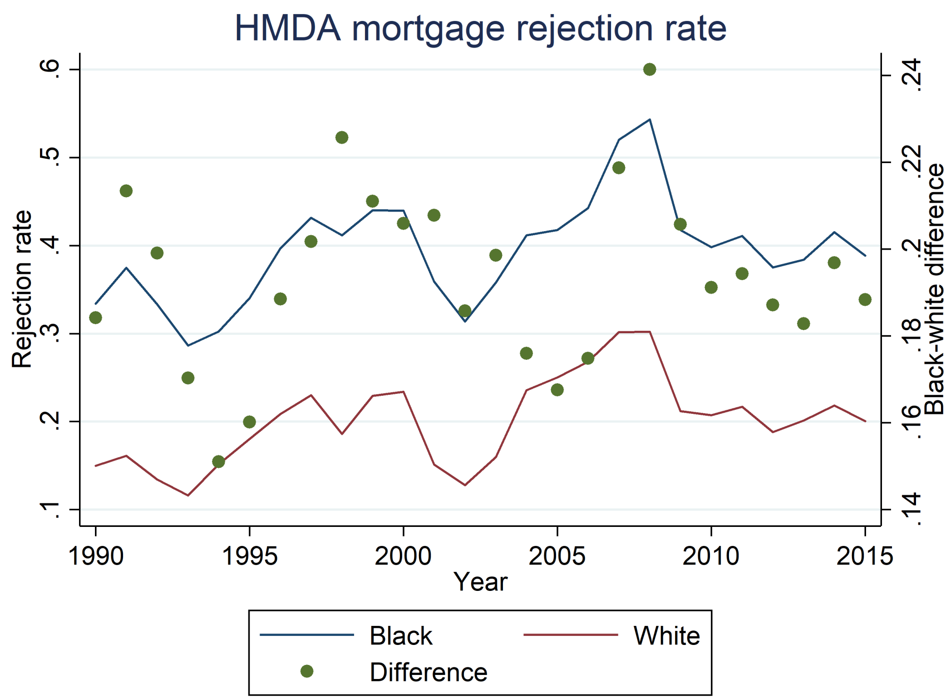

(a) Rejection rates

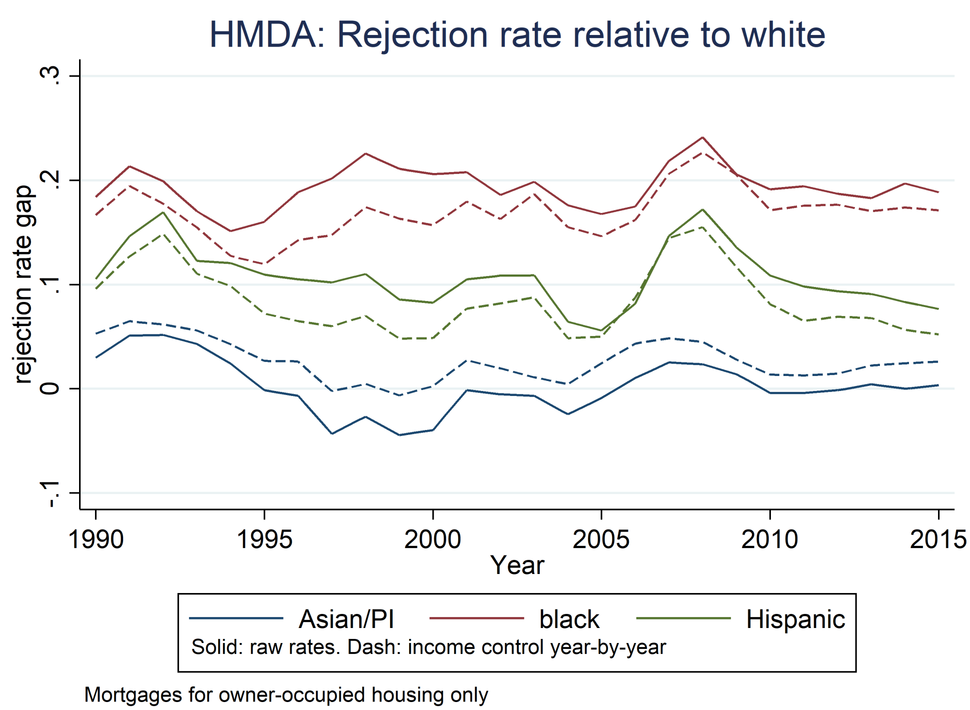

(b) Rejection rate gap

Figure B.14: Mortgage origination

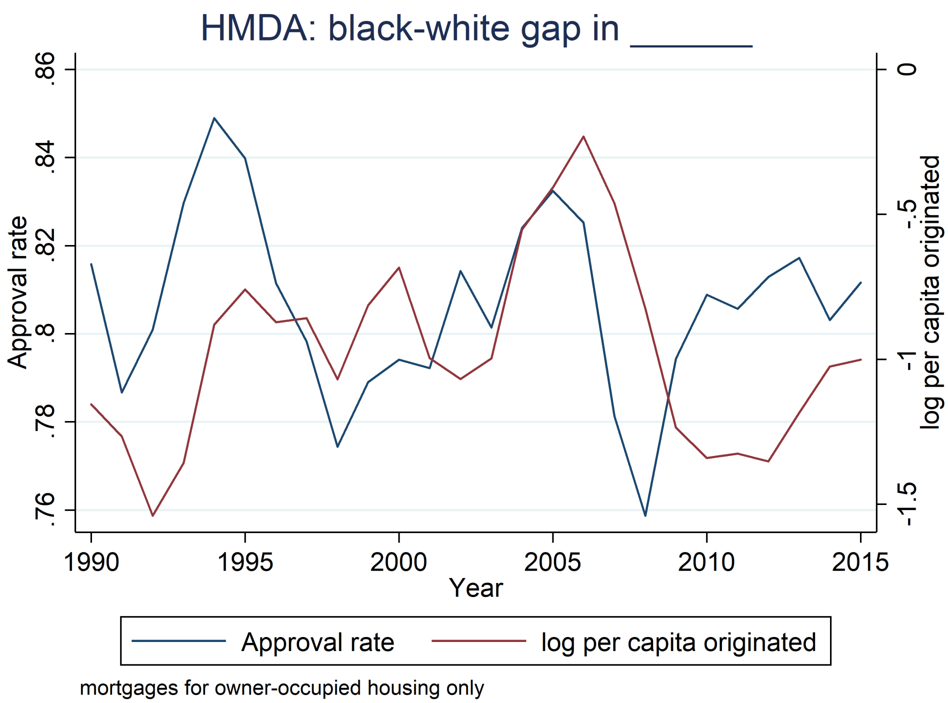

(a) Time series

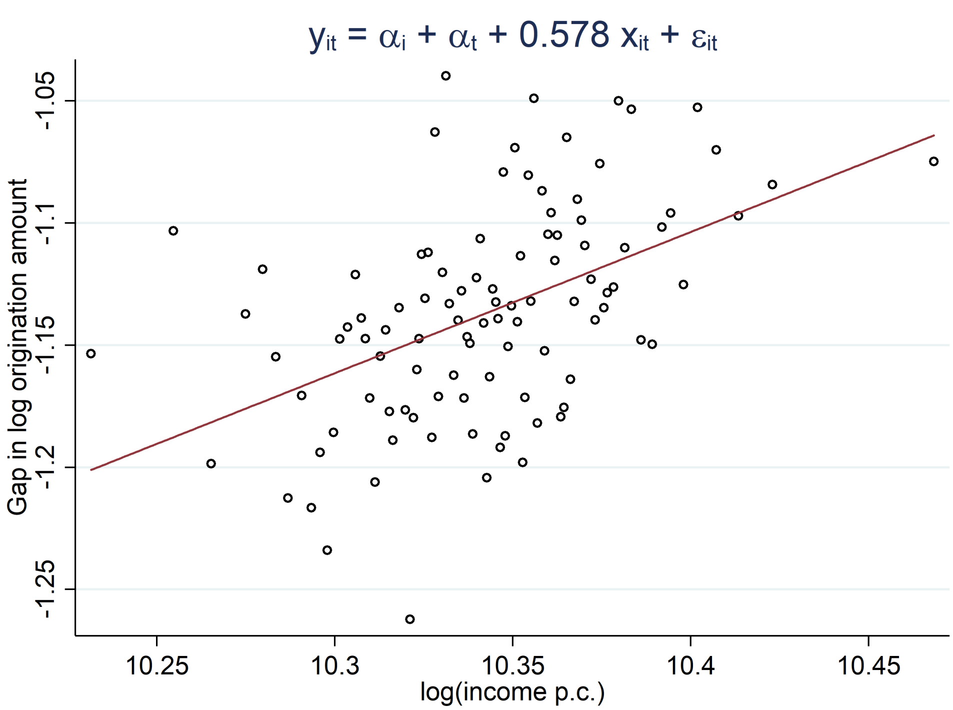

(b) Diff-in-diff

Table B.1: Validation of prejudice measures

All regressions are weighted by DMA population.

(a) Prejudice measures vs. 1860 slave share

	(1) N-word (log index)	(2) "KKK" (log index)	(3) crime	(4) spanking	(5) GSS index
Slave share (1860)	0.909*	1.713***	0.773	0.013	0.109
	(2.61)	(3.99)	(0.17)	(1.13)	(0.48)
Black pop share 2010	-2.653***	-3.427***	-2.064	0.053	-0.457
	(-6.46)	(-5.72)	(-1.00)	(1.76)	(-1.61)
Observations	163	165	125	83	126
Adjusted R^2	0.379	0.418	0.477	0.510	0.355
State FE	O	O	O	O	O

t statistics in parentheses

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

(b) Slave share instrumented with 1860 cotton production

	(1) Slave share (1860)	(2) N-word (log index)	(3) "KKK" (log index)	(4) crime	(5) spanking	(6) GSS index
Cotton per capita (1860)	0.273***					
	(7.58)					
Slave share (1860)		1.336*	3.312***	1.137	0.092**	0.725
		(2.33)	(4.57)	(0.28)	(3.39)	(1.63)
Black pop share 2010		-2.902***	-4.359***	-2.268	-0.021	-0.799*
		(-6.01)	(-6.83)	(-0.74)	(-0.55)	(-2.53)
Observations	166	163	165	125	83	126
Adjusted R^2	0.877	0.373	0.358	0.477	0.188	0.293
State FE	O	O	O	O	O	O

t statistics in parentheses

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

Table B.1: Validation of prejudice measures (continued)

(c) Prejudice measures vs. GSS (average level)

	(1) N-word (log index)	(2) "KKK" (log index)	(3) crime	(4) spanking
GSS index	0.858***	0.880***	-0.058	0.034***
	(5.25)	(4.15)	(-0.05)	(4.67)
Black pop share 2010	-1.637***	-1.926***	-0.673	0.070**
	(-5.00)	(-3.76)	(-0.15)	(3.05)
Observations	143	143	107	73
Adjusted R^2	0.418	0.419	0.473	0.647
State FE	O	O	O	O

t statistics in parentheses

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

(d) Prejudice measures vs. GSS (diff-in-diff)

	(1) N-word (log index)	(2) KKK (log index)	(3) crime	(4) spanking
GSS index	0.036	0.044*	0.121	0.001*
	(1.73)	(2.02)	(0.51)	(2.29)
Observations	1378	1378	1237	281
Adjusted R^2	0.841	0.886	0.822	0.829
DMA FE, year FE	O	O	O	O

t statistics in parentheses

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

Table B.2: Countercyclical prejudice cross-section

All regressions are weighted by DMA population.

(a) Panel 2004-2014

	Log income	N-word		"KKK"		crime		spanking		GSS	
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)	(11)
	1st stage	OLS	IV	OLS	IV	OLS	IV	OLS	IV	OLS	IV
Saiz x nat HP	-0.039***										
	(-6.06)										
bartik	0.855***										
	(4.03)										
log(income p.c.)		-0.928*	-1.719**	-1.459***	-2.744***	-3.592*	2.102	-0.015*	0.003	0.117	4.070
		(-2.36)	(-2.67)	(-3.74)	(-4.82)	(-2.52)	(0.36)	(-2.12)	(0.14)	(0.56)	(1.76)
Observations	2226	2442	1840	2482	1861	2441	1494	1729	1087	2899	1306
Adjusted R^2	0.980	0.842	0.848	0.862	0.884	0.796	0.896	0.784	0.783	0.387	0.131
DMA FE, year FE	O	O	O	O	O	O	O	O	O	O	O
Hansen's J			2.761		4.974		1.174		0.127		0.001
p-value			(0.10)		(0.03)		(0.28)		(0.72)		(0.98)

t statistics in parentheses

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

Table B.2: Countercyclical prejudice cross-section (continued)

(b) Single cross-section 2006-2009

	Log income	N-word		"KKK"		crime		spanking		GSS	
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)	(11)
	1st stage	OLS	IV	OLS	IV	OLS	IV	OLS	IV	OLS	IV
Saiz el	0.021***										
	(7.75)										
bartik	1.886***										
	(5.39)										
log(income)		-0.937*	-1.883**	-1.229***	-2.584***	-3.599*	-14.842**	-0.022*	-0.026	0.433	0.863
		(-2.13)	(-2.74)	(-4.07)	(-5.27)	(-2.10)	(-2.65)	(-2.17)	(-1.33)	(0.32)	(0.70)
Constant	0.032**	-0.326***		-0.238***		0.069		-0.001**		0.146*	
	(2.70)	(-10.49)		(-12.21)		(0.71)		(-3.34)		(2.21)	
Observations	186	204	184	207	186	137	125	105	95	106	102
Adjusted R^2	0.329	0.048	-0.017	0.092	-0.040	0.016	-0.272	0.011	0.027	-0.007	-0.009
Hansen's J			0.025		3.099		0.173		0.394		0.495
p-value			(0.87)		(0.08)		(0.68)		(0.53)		(0.48)

t statistics in parentheses

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

Table B.2: Countercyclical prejudice cross-section (continued)

(c) Single cross-section 2000-2006

	Log income	crime		spanking		GSS	
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
	1st stage	OLS	IV	OLS	IV	OLS	IV
ADH import exposure	-0.008*						
	(-2.14)						
housing demand	0.369***						
	(3.90)						
log(income)		-0.155	10.282	-0.011	-0.055	-0.159	-0.358
		(-0.04)	(1.10)	(-1.27)	(-1.80)	(-0.22)	(-0.31)
Constant	0.228***	-0.601		0.000		-0.197	
	(15.78)	(-0.73)		(0.11)		(-1.28)	
Observations	180	92	82	105	95	79	79
Adjusted R^2	0.268	-0.011	-0.088	-0.006	-0.070	-0.012	-0.013
Hansen's J			0.088		3.685		1.196
p-value			(0.77)		(0.05)		(0.27)

t statistics in parentheses

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

APPENDIX C

THE OBAMA EFFECT: EFFECT OF BLACK ELECTORAL VICTORY ON RACIAL PREJUDICE AND INEQUALITY

C.1 Figures and Tables

Figure C.1: Screenshots from Our Campaigns

The figures below are screenshots from the Our Campaigns website (www.ourcampaigns.com/). Panels (a) and (b) are for Chicago's fifth city ward, with panel (a) showing the political position's main page with the current officeholder and the map of the jurisdiction along with other details, and panel (b) showing the electoral race history, with the date, type, candidates and their count and shares of votes received for each electoral race. Panel (c) is the candidate page for Harold Washington, who was the first black mayor of Chicago. Among other personal details, "Tags" contains his race.

(a) Political office example

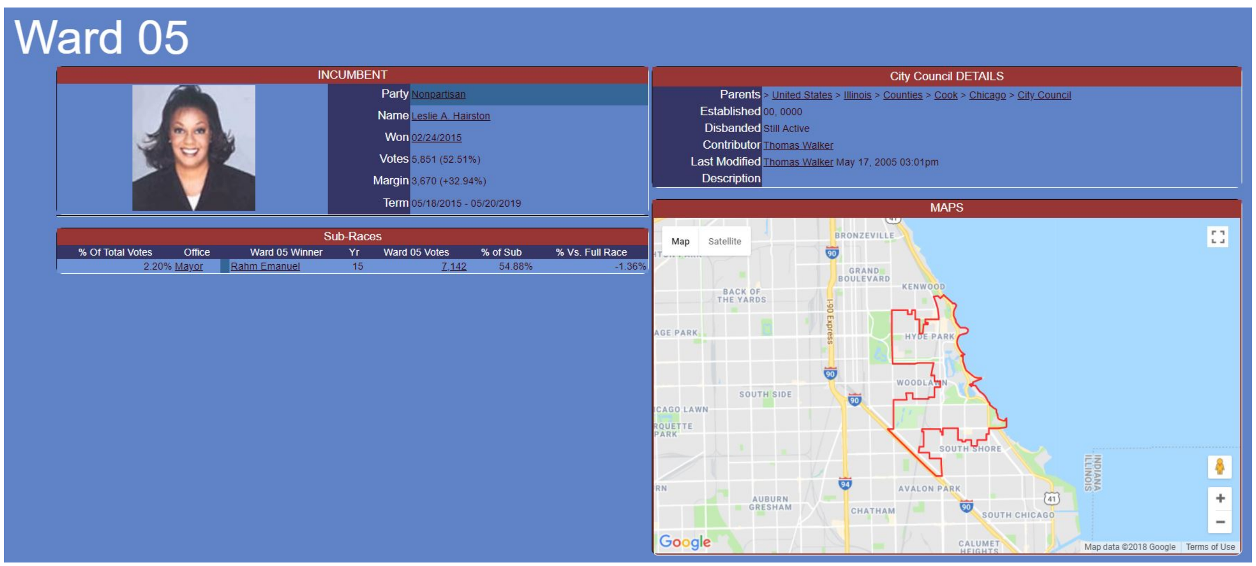

Figure C.1: Screenshots from Our Campaigns (continued)

(b) Electoral race history example

RACES [Show Primaries]

Date	Type	Results						
Feb 24, 2015	General Election	Leslie A. Hairston(I) 5,851 52.51%	Anne Marie Miles 2,181 19.57%	Tiffany N. Brooks 891 8.00%	Jocelyn Hare 821 7.37%	Jedidiah L. Brown 792 7.11%	Robin Boyd Clark 599 5.38%	Write-In (W) 8 0.07%

Date	Type	Results				
Feb 22, 2011	General Election	Leslie A. Hairston(I) 7,217 61.77%	Anne Marie Miles 2,489 21.30%	Glenn Ross 826 7.07%	Carolyn Hightower Chalmers 701 6.00%	Michele A. Tankersley 451 3.86%
Feb 27, 2007	General Election	Leslie A. Hairston(I) 6,748 74.67%	Oscar Worrill 1,769 19.58%	Sylvester "Junebug" Hendricks 520 5.75%		
Feb 25, 2003	General Election	Leslie A. Hairston(I) 6,355 71.93%	Oscar Worrill 1,073 12.14%	Carolyn Hightower Chalmers 713 8.07%	Anthony T. Blair 694 7.86%	

(c) Candidate example

Washington, Harold

CANDIDATE DETAILS

Affiliation	Democratic
Name	Harold Washington
Address	Chicago, Illinois , United States
Email	None
Website	None
Born	April 15, 1922
Died	November 25, 1987 (65 years)
Contributor	Wishful Thinking
Last Modifed	RBH Jan 31, 2016 04:33am
Tags	Black - Divorced - Methodist -
	Harold Washington (1922-1987) was the first African-American mayor of Chicago.

Figure C.2: IAT data from Project Implicit Database

Panel (a) plots the average Race Implicit Association Test (IAT) D score for white respondents from Project Implicit Database. The average has been taken after projecting out dummies for age (each whole number), nine education buckets and gender. Panel (b) plots the distribution of raw Race IAT D score from Project Implicit Database for 2002-2017. The D score has a possible range of -2 to +2, where higher number indicates bias against black Americans. The vertical red lines indicate break points for a common description of pro-white bias. [CHECK which ranges indicate which colloquial descriptor of bias]

(a) Composition-adjusted, for white respondents

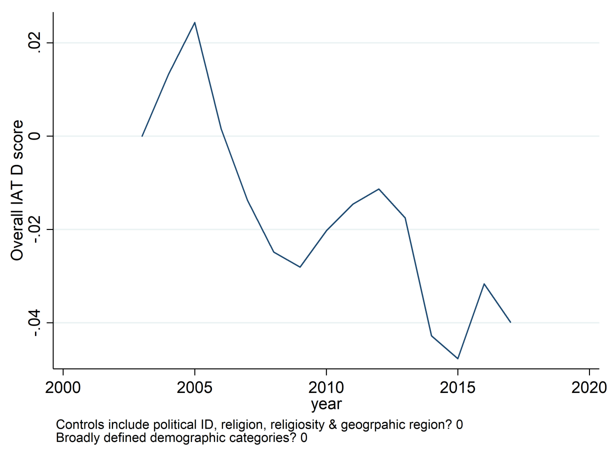

(b) Distribution (raw)

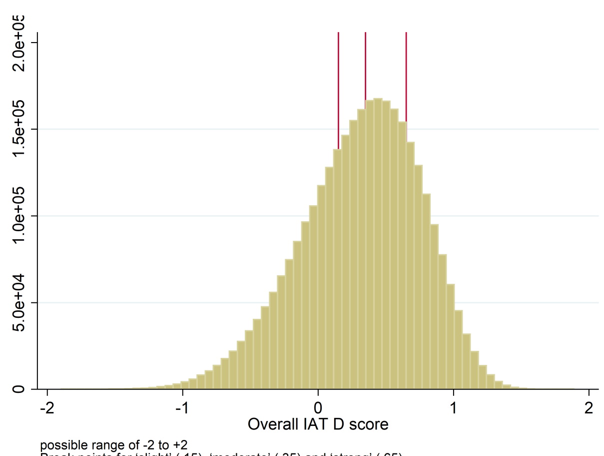

Figure C.3: IAT distribution

Panel (a) plots the distribution of raw Race IAT D score from Project Implicit Database for 2002-2017. The D score has a possible range of -2 to +2, where higher number indicates bias against black Americans. The vertical red lines indicate break points for a common description of pro-white bias. [CHECK which ranges indicate which colloquial descriptor of bias] The next five panels plot the raw Race IAT D score by demographic sub-groups for 2002-2017. Panel (b) splits the Race IAT scores by the response to the question, "What brought you to this website." "Voluntary" responses in green line include "Recommendation of a friend or co-worker," "Mention or link at a non-news Internet site," "Mention in a news story (any medium)" or "My Internet search for this or a related topic." "Mandatory" responses in blue line include "Assignment for work" or "Assignment for school." Responses of "other" and those without a response to this question were classified as "unknown/other" in the red line. Panel (c) splits the Race IAT scores by education and plots two groups: the green histogram is for those with up to a high school diploma, while the white histogram is for those with bachelor's degrees or higher. Panel (d) splits the Race IAT scores by gender: the green histogram is for females, while the white histogram is for males. Panel (e) splits the Race IAT scores by self-stated political ideology along a 7-point scale: the red histogram is for those who responded with any degree of "conservative", the blue histogram is for those who responded with any degree of "liberal", while the white histogram is for those who responded "neutral." Panel (f) splits the Race IAT scores by self-stated race: the blue line is for Asians, the green line is for black Americans, the red line is for those of Hispanic origin, and the black line is for white Americans.

(a) By race

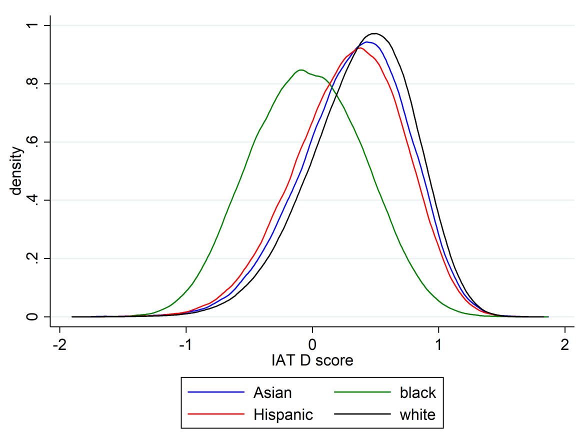

(b) By where respondents came from

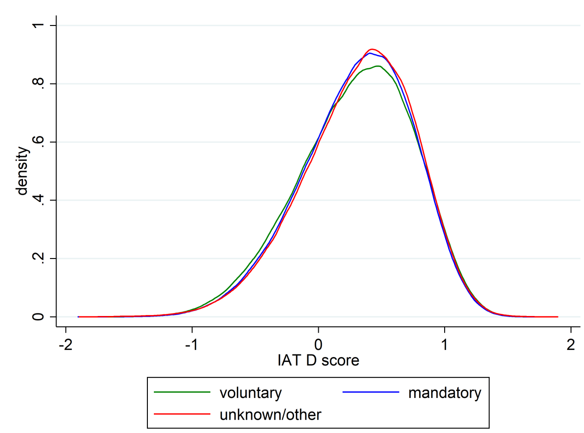

Figure C.3: IAT distribution (continued)

(c) By education

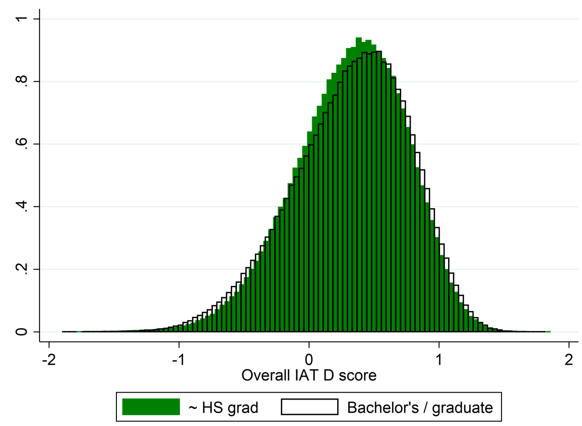

(d) By gender

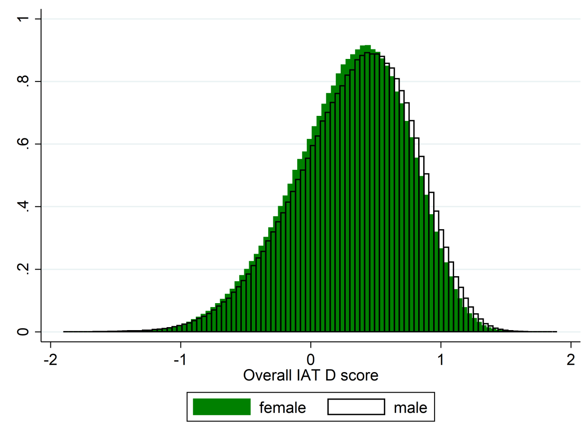

(e) By political ideology

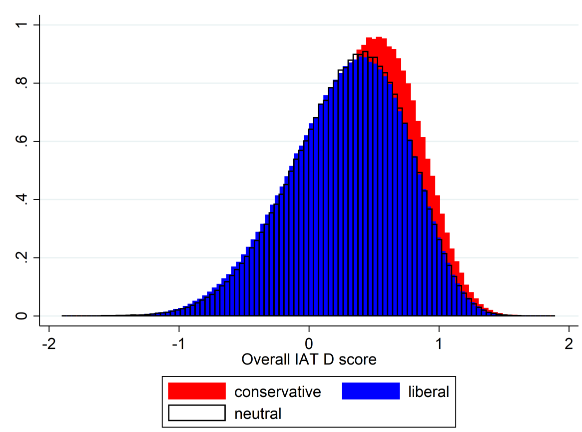

(f) By region (white only)

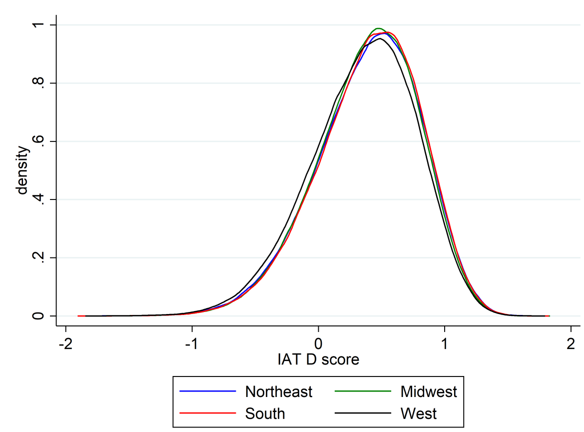

Figure C.4: IAT re-weighted using ACS

Both figures plot the distribution of Race IAT D scores, using no weight (blue line) and using weights imputed using the 2008-2012 American Community Survey (ACS) to make the sample representative along the demographic variables of gender × age × education (red line). Panel (a) plots the distribution for all observation; panel (b) plots the distribution only for white respondents.

(a) All races

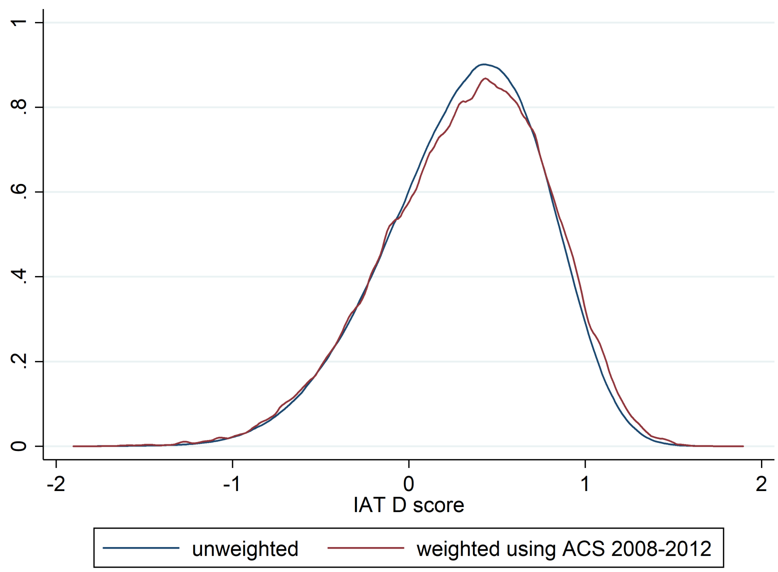

(b) White only

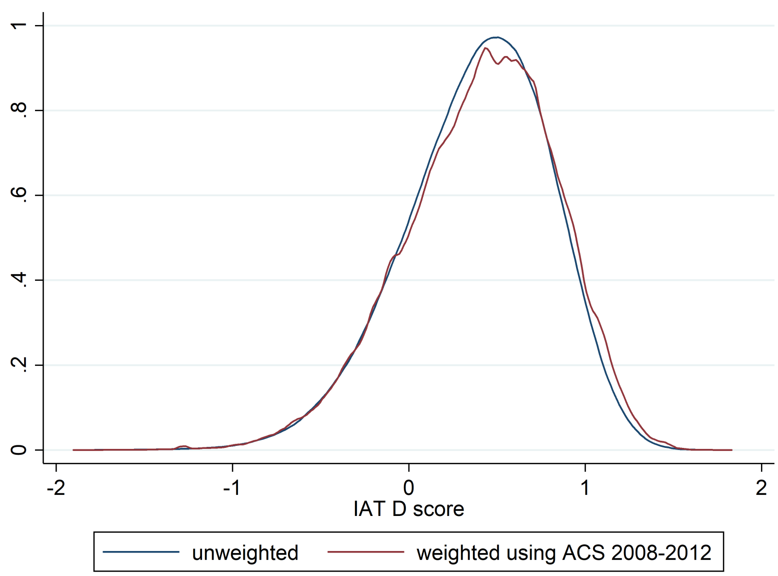

Figure C.5: IAT time series by source

Panels (a), (b), (c), (d) plot average Race IAT D scores for sub-groups from the Project Implicit Database; panels (e) and (f) on the bottom row plot counts of responses for sub-groups. Panels (a), (c) and (e) on the left column are for all responses; panels (b), (d) and (f) on the right column are for responses with self-reported race of non-Hispanic white. Within each plot, separate lines are sub-groups based on the responses to the question, “What brought you to this website.” “Voluntary” responses in green line include “Recommendation of a friend or co-worker,” “Mention or link at a non-news Internet site,” “Mention in a news story (any medium)” or “My Internet search for this or a related topic.” “Mandatory” responses in red line include “Assignment for work” or “Assignment for school.” The blue lines are for both of these sub-groups. Plots (a) and (b) took sub-group averages of the raw Race IAT D score; plots (c) and (d) took the averages after projecting out dummies for age (each whole number), nine education buckets and gender.

(a) All: Raw

(b) White only: Raw

Figure C.5: IAT time series by source (continued)

(c) All: Ex demo

(d) White only: Ex demo

(e) All: Count

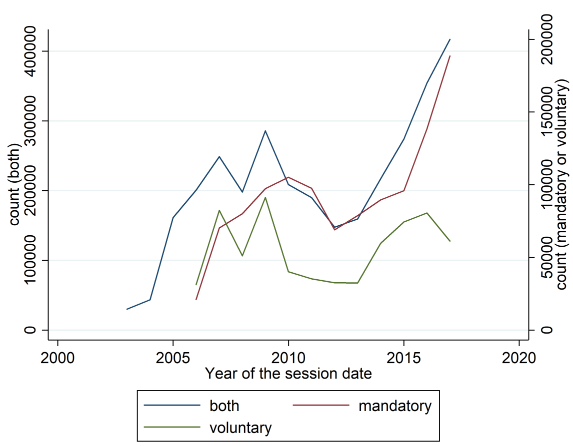

(f) White only: Count

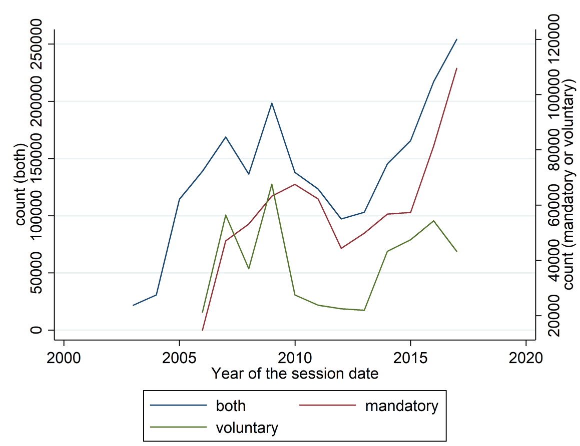

Figure C.6: Difference-in-difference

All panels plot the average Race IAT D score for white respondents living in the areas affected by a set of elections in connected red dots, against surrounding areas in the same state in solid blue line. Panels (a) and (c) on the left are for elections where a black candidate won against a white runner-up; panels (b) an (d) on the right are for elections where a black runner-up lost to a white candidate. In the top panels (a) and (b), raw Race IAT D scores are used; in the bottom panels (c) and (d), the Race IAT D scores have been first residualized by dummies for age (each whole number), nine education buckets and gender, and then re-weighted by the same three demographic variables to be representative of the US population using the American Community Survey for 2008-2012. The dotted red lines and and light blue shaded area show standard errors for the areas affected and the surrounding areas, respectively.

(a) Black winner: Raw

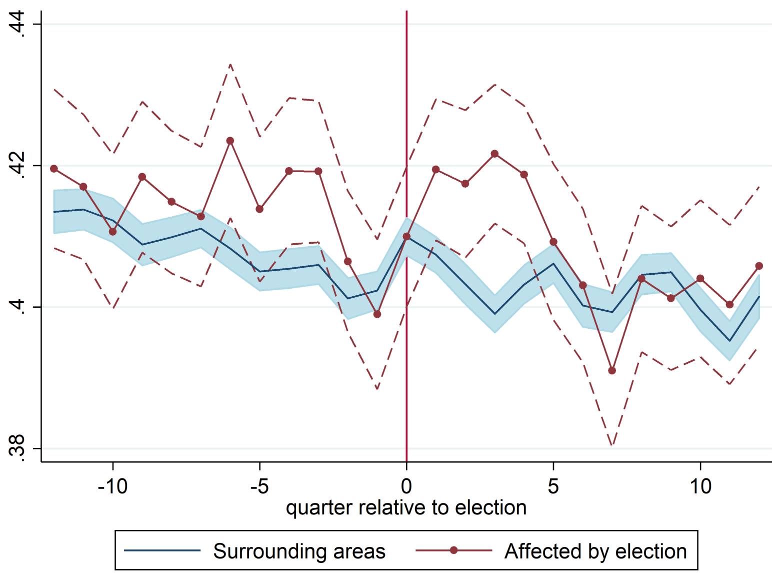

(b) Black runner-up: Raw

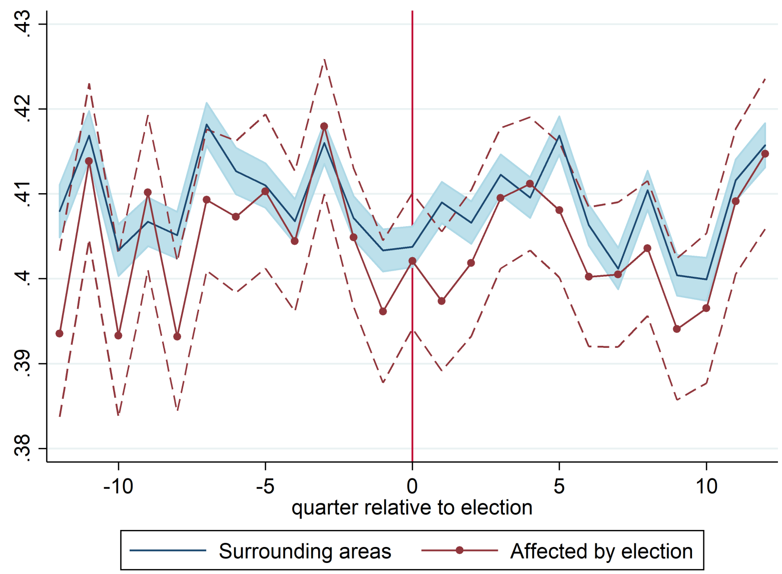

Figure C.6: Difference-in-difference (continued)

(c) Black winner: Composition-adjusted

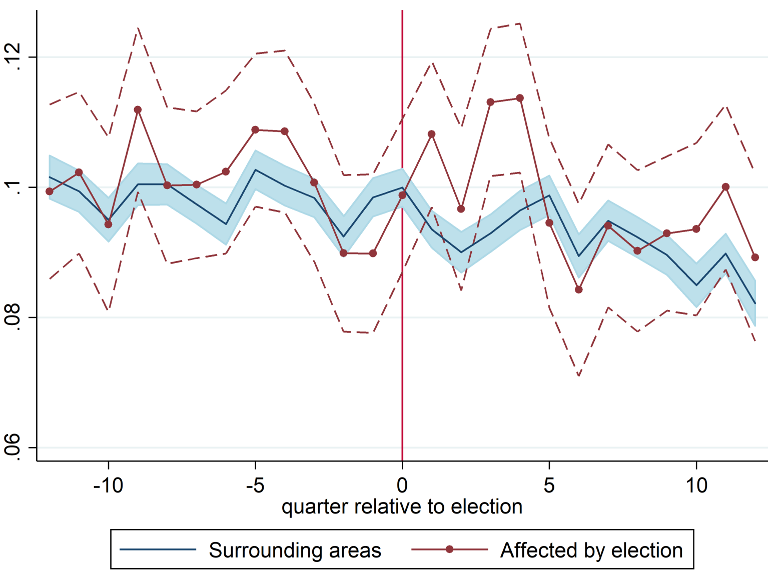

(d) Black runner-up: Composition-adjusted

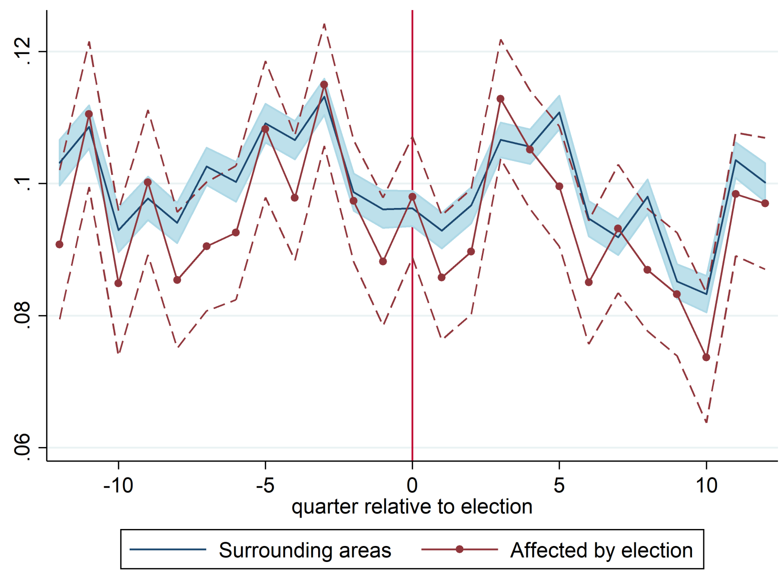

Figure C.7: Regression discontinuity

Both panels plot bin scatters of average Race IAT D score among white respondents from Project Implicit Database: panel (a) is for the 3-year window before an election; panel (b) is for the 3-year window after the election (i.e. placebo period). Each point represents within-bin averages of Race IAT D scores across election-county-month level observations), where the bins are defined by centiles sorted on the vote margin in the corresponding election. The sample has been restricted to elections where the winner and the runner-up include one black and one white candidate; vote margin is defined as the percentage point difference between the black candidate and the white candidate (i.e. positive vote margins represent elections where the black candidate won, while negative vote margins represent elections where the black candidate lost to the white candidate). The domain has been restricted to elections where the vote margin was within 10%. The averages were formed using Race IAT D scores after projecting out dummies for age (each whole number), nine education buckets and gender. The vertical dotted line divides black candidates' losses on the left to their victories on the right. The red lines are best linear fit on each side. The vertical distance along the vertical dotted line between the two red lines is a rough estimate of the regression discontinuity estimate.

(a) Before the election

(b) After the election

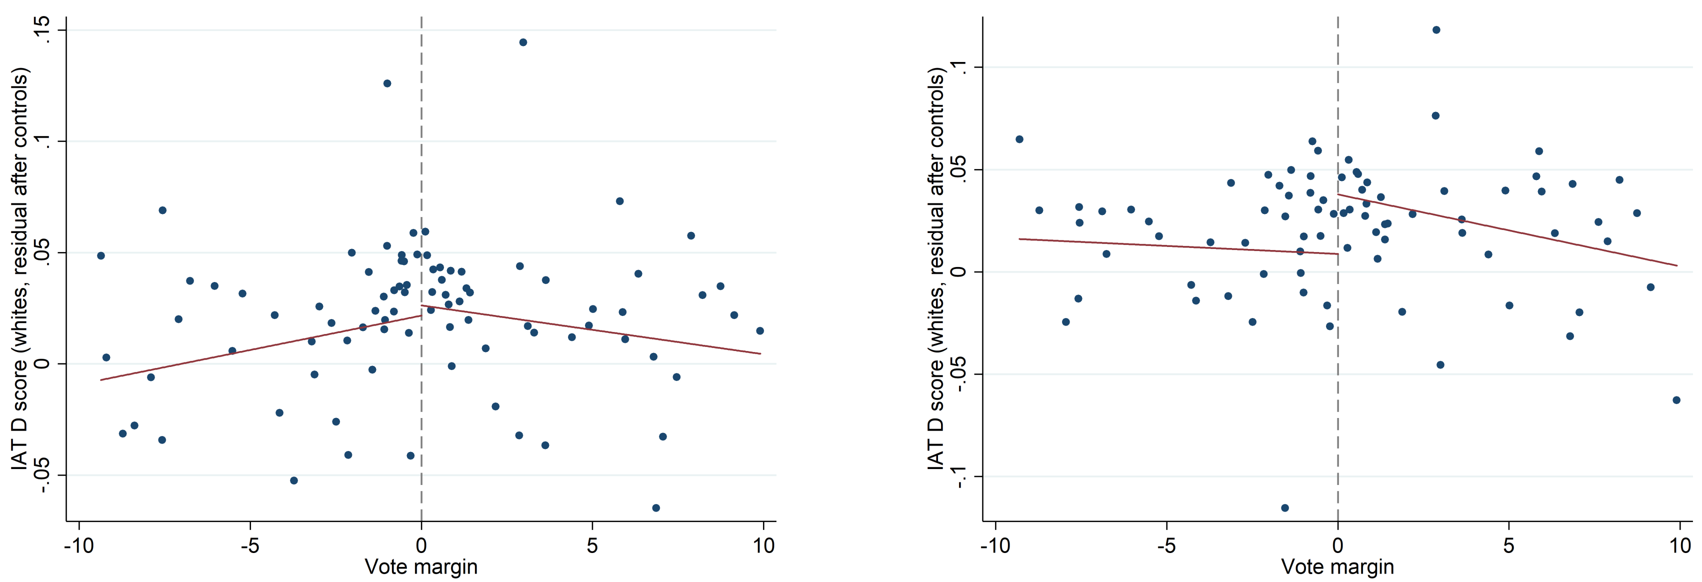

Figure C.8: Time series of discontinuity

Both panels plot time series of regression discontinuity estimates of the Race IAT D score between areas with a close black winner and a close black loser, for non-overlapping 3-month windows relative to the election. Each red point is the estimate, with 95% confidence intervals shaded in gray. The vertical red line at 0 indicates the 3-month period starting from the month of election; points to the right of the vertical line represents non-overlapping 3-month periods after the election. The regression discontinuity estimates have been estimated from by using elections where the vote margin between the winner and the runner-up was within 10%, with linear controls on either side of the discontinuity. Regressions have been run on observations at the election-county-month level, with each observation weighted by the fraction of the county affected by the eleciton.

(a) Raw

(b) Controlling for demographic composition

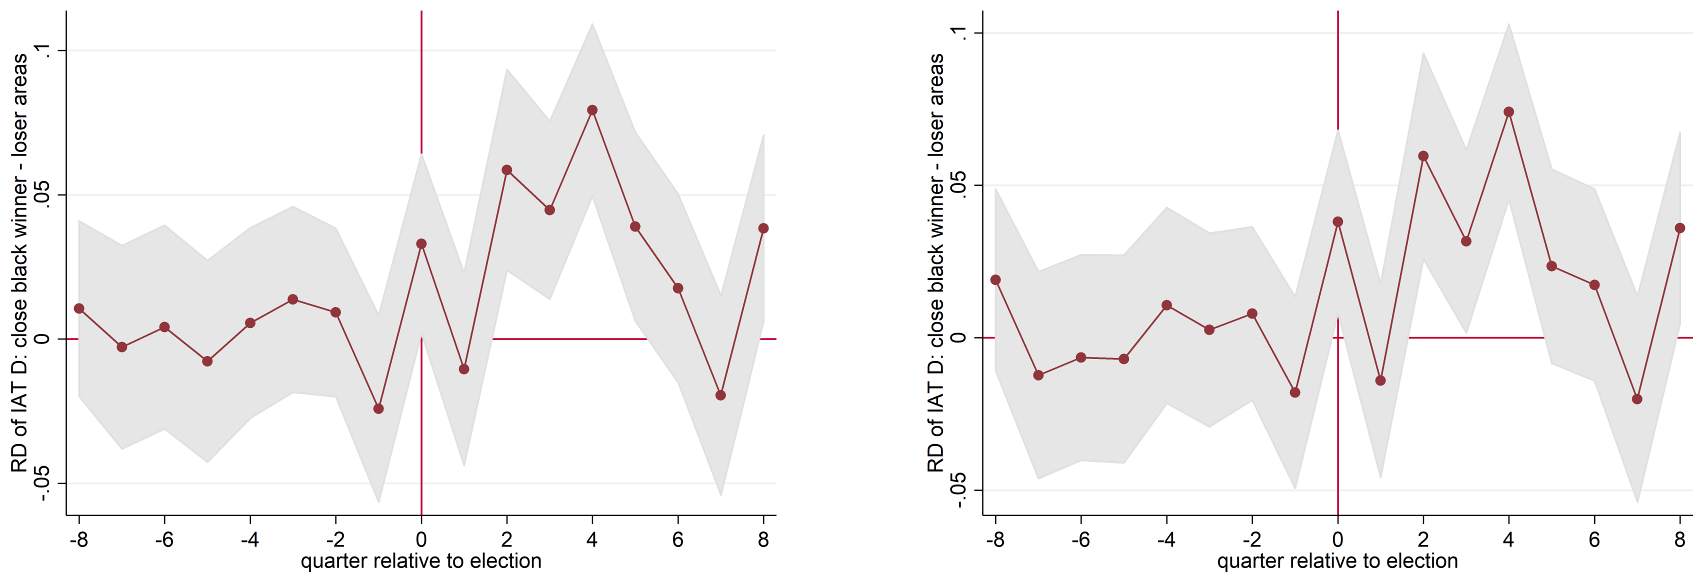

Table C.1: Our Campaigns

Panel (a) shows the number of elections falling into each category, for years 2003 and onward. Row titles indicate which type of office the elections in that row are for. The first two columns are for all electoral races; the third and fourth columns are for electoral races in which the winner and the runner-up include one white and one black candidate; the last two columns are such electoral races between black and white candidates that are close. Close election is defined here as those with a vote margin between the winner and the runner-up of 10% or less of votes. Odd columns show counts of all such elections; even columns show the count of elections in which the winner is a black candidate. The last row with the total counts show that for close elections between a black and a white candidate, the probability that the black candidate wins is roughly half.
Panel (b) shows. For all candidates who make it among the top two candidates in the electoral races in panel (a), after 2002.

(a) Offices

		all races		black-white races		close black-white races	
Gov level	Office type	all	black winner	all	black winner	all	black winner
federal	house	2,933	202	111	78	11	4
state	governor	308	12	26	11	15	9
	senate	9,362	655	471	156	51	24
	house	33,046	2,426	1,518	486	145	78
	state other	1,428	92	147	49	88	37
local	county other	7,350	475	448	146	77	34
	city other	4,309	601	277	148	72	35
	mayor	2,755	215	166	74	42	26
Total		61,491	4,678	3,164	1,148	501	247

Table C.1: Our Campaigns (continued)

(b) Candidate race identification

Method	race	Our Campaigns tag: Asian	Black	Hispanic	White	no tag
	total count by tag	39	276	187	921	49,938
Facial recognition	Asian	0.67	0.06	0.10	0.03	0.03
	Black	0.15	0.68	0.06	0.01	0.04
	White	0.13	0.19	0.78	0.92	0.48
	unidentifiable	0.00	0.02	0.01	0.03	0.01
	no photo	0.05	0.05	0.04	0.01	0.45
Surname	Asian	0.49	0.00	0.00	0.00	0.01
	Black	0.00	0.02	0.00	0.00	0.00
	Hispanic	0.05	0.01	0.66	0.01	0.03
	White	0.00	0.12	0.11	0.58	0.53
	mixed	0.41	0.76	0.14	0.33	0.34
	uncommon/unmatched	0.05	0.08	0.09	0.08	0.09

Table C.2: IAT distribution

Panel (a) shows the demographic distribution in the 2008-2012 5-year American Community Survey and 2008-2012 Race IAT data from Project Implicit Database, along gender, age bins and educational attainment. The ACS statistics have been accessed via th e Integrated Public Use Microdata Series (IPUMS), in order to get counts interacting the three demographic variables. Weights to make the IAT sample more representative of the American population were imputed as $\frac{\text{ACS population share}}{\text{IAT sample share}}$ for gender $\times$ age $\times$ education bins, fully interacted. Single-year ages were used, and educational attainment has been grouped into the categories displayed in panel (a). The last column of panel (a) displays the average weight for the broader demographic subgroup weighted using the IAT sample share; these weights are equivalent to the ratio of the first two columns.

(a) Distribution vs ACS

demographic variable		share 2008-2012		average
		ACS	IAT	IAT weight
gender	male	0.49	0.41	1.2
	female	0.51	0.59	.87
age	unknown	0.00	0.01	0
	under 19	0.25	0.23	1.1
	19-22	0.05	0.26	.2
	23-29	0.09	0.20	.47
	30-39	0.13	0.15	.89
	40-49	0.15	0.09	1.7
	50-59	0.13	0.05	2.7
	60+	0.18	0.02	9.5
education	no HS grad	0.27	0.14	1.9
	HS grad	0.29	0.09	3.2
	some college / associate's	0.17	0.46	.38
	bachelor's	0.13	0.13	.99
	graduate	0.07	0.06	1.3
	N/A	0.08	0.12	.65

Table C.2: IAT distribution (continued)

Panel (b) show the counts and characteristics of survey responses, grouped by the response to the question, "What brought you to this website," from the Project Implicit Database. For each subgroup, I also report the average age and the average Race IAT D score, both raw and residualized by demographic fixed effects.

(b) By source of respondent

source		count	average age	IAT score	
grouping	actual response			raw	adjusted
mandatory	Assignment for work	142,367	37.6	0.27	-0.02
	Assignment for school	989,129	24.1	0.31	-0.00
	Sub-total	1,131,496	25.8	0.31	-0.01
voluntary	Recommendation of a friend or co-worker	304,153	30.2	0.30	-0.02
	Mention or link at a non-news Internet site	187,723	33.1	0.30	-0.01
	Mention in a news story (any medium)	132,086	37.6	0.30	-0.02
	My Internet search for this topic or a related topic	48,546	31.8	0.30	-0.01
	Sub-total	672,508	32.6	0.30	-0.02
unknown	.	1,238,011	26.1	0.33	0.02
	Other	93,791	32.3	0.30	-0.02
	Sub-total	1,331,802	26.5	0.33	0.01

Table C.3: IAT validation

These tables report the cross-sectional and panel relationship between the local average of the Race IAT D score and the variables in column headers. The cross-section is either a county for panels (a) and (b), or a Nielsen Designated Market Area for panel (c). See text for definition and source for each variable used here. In panel (a), the sample is restricted to slave states in 1860, and state fixed effects are included, following [2].

(a) 1860 slavery (county)

	IAT D			Thermology white-black			Prefer white/black		
	(1) OLS	(2) OLS	(3) IV	(4) OLS	(5) OLS	(6) IV	(7) OLS	(8) OLS	(9) IV
slave share 1860	0.020**	0.029***	0.041***	0.224***	0.247***	0.222*	0.140***	0.139***	0.123**
	(2.08)	(2.58)	(2.92)	(3.23)	(2.76)	(1.94)	(3.89)	(3.04)	(2.09)
pop black		-0.020	-0.026*		-0.022	-0.009		0.007	0.015
		(-1.51)	(-1.86)		(-0.19)	(-0.07)		(0.11)	(0.22)
Adjusted R^2	0.039	0.039	0.039	0.101	0.078	0.078	0.154	0.141	0.141
Observations	1114	1111	1111	1114	1111	1111	1114	1111	1111

t statistics in parentheses

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Table C.3: IAT validation (continued)

(b) Racial prejudice (explicit measures from Project Implicit Database)

	Thermology white-black		Prefer white/black	
	(1) cross-section	(2) panel	(3) cross-section	(4) panel
IAT D	1.299***	0.983***	0.705***	0.532***
	(26.52)	(22.79)	(26.67)	(20.91)
Adjusted R^2	0.135	0.208	0.184	0.252
Observations	3115	37806	3114	37357

t statistics in parentheses

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

(c) Racial prejudice (other measures)

	DMA-level			county-level	
	(1) Google n-word	(2) Google KKK	(3) Spanking (b-w)	(4) crime w.o.b.	(5) GSS
IAT D	0.551***	0.299	0.013**	-3.175	0.319
	(2.62)	(1.57)	(2.18)	(-1.19)	(1.20)
Adjusted R^2	0.178	0.224	0.009	-0.001	0.061
Observations	204	207	105	1615	350

t statistics in parentheses

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Table C.4: Difference-in-difference

Using only elections with a black winner, and for election i, county j and year t, I estimate

$$Y_{ijt} = \alpha_{ij} + \alpha_{it} + \beta \,\{\text{in jurisdiction of winner}\}_{ij} \times \{\text{after election}\}_{it} + \eta_{ijt}$$

For each election i, I include the 3-year window before and after the election, and include all counties in the same state as the jurisdiction associated with the election.

(a) IAT scores

	Raw		Demo-adjusted		Both FE	
	(1) all	(2) mand	(3) all	(4) mand	(5) all	(6) mand
Treatment	0.006***	0.003	0.002	0.002		
	(2.81)	(0.91)	(0.83)	(0.46)		
Post	-0.005***	-0.001**	-0.006***	-0.001		
	(-10.78)	(-2.36)	(-10.90)	(-1.39)		
Treatment x Post	-0.001	-0.003	0.004	0.000	0.003	-0.003
	(-0.20)	(-0.68)	(1.00)	(0.11)	(1.06)	(-0.77)
Constant	0.408***	0.410***	0.098***	0.101***		
	(1210.98)	(904.88)	(261.59)	(201.58)		
Observations	2045102	1241535	1838082	1113151	1835660	1107872
Adjusted R^2	0.000	0.000	0.000	-0.000	0.020	0.021
FE demographic			O	O	O	O
ACS weights			O	O	O	O
FE year					O	O
FE county					O	O
Subsample?		mandatory		mandatory		mandatory

t statistics in parentheses

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Table C.5: Regression discontinuity: IAT D scores

Both panels display the main results from the regression discontinuity design. The main specification imposes a 10% bandwidth. For observations at election i, geography (e.g. county) j, and event time (e.g. month or year) t, I run the following regression specification:

$$Y_{ijt} = \alpha + \beta_1 \mathbf{1}\left\{\text{vote margin} > 0\right\}_{it} + \gamma_0 \left[\text{vote margin}\right]^-_{it} + \gamma_1 \left[\text{vote margin}\right]^+_{it} + \varepsilon_{ijt}$$

where "vote margin" is the percentage point vote margin between the black candidate and the white candidate. The first indicator term denotes a dummy for elections in which the black candidate won. The next two terms are linear controls separately for elections where the white candidate won (i.e. negative vote margin) and where the black candidate won (i.e. positive vote margin). Observations are weighted by the fraction of the geography (e.g. county) affected by the election. Outcome variable is the Race IAT D scores from Project Implicit Database.

	Raw		Demo-adjusted		With FEs	
	all	mand	all	mand	all	mand
			Post			
Black winner	0.025**	0.025**	0.027**	0.037***	0.012	0.013
	(2.50)	(2.16)	(2.22)	(2.98)	(1.34)	(1.25)
Margin (winner)	-0.004***	-0.003*	-0.002	-0.002	0.002	0.003
	(-2.61)	(-1.95)	(-0.97)	(-1.35)	(0.87)	(1.12)
Margin (loser)	0.001	-0.001	-0.001	-0.002	-0.000	0.001
	(0.49)	(-0.56)	(-0.23)	(-0.99)	(-0.26)	(0.50)
Observations	23590	18177	21462	16481	21444	16452
Adjusted R^2	0.002	0.001	0.001	0.002	0.075	0.070
			Pre			
Black winner	0.003	0.012	0.011	0.022	-0.003	-0.001
	(0.33)	(0.97)	(1.08)	(1.58)	(-0.29)	(-0.08)
Observations	23437	15319	22271	14421	22250	14390
Adjusted R^2	0.001	0.000	0.001	0.002	0.065	0.075
			Model specifications:			
FE demographic			O	O	O	O
ACS weights			O	O	O	O
FE year-month					O	O
FE county					O	O
only mandatory		O		O		O

t statistics in parentheses

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Table C.6: Economic gaps

Using only elections with a black winner, and for election i, county j and year t, I estimate

$$Y_{ijt} = \alpha_{ij} + \alpha_{it} + \beta \left\{\text{in jurisdiction of winner}\right\}_{ij} \times \left\{\text{after election}\right\}_{it} + \eta_{ijt}$$

For each election i, I include the 3-year window before and after the election, and include all counties in the same state as the jurisdiction associated with the election.

(a) Difference-in-difference

	(1) unemployment transition	(2) employment to pop	(3) rejection rate	(4) log origination to pop
Treatment x Post	-0.002	-0.000	-0.003	0.015*
	(-0.69)	(-0.13)	(-1.63)	(1.76)
Observations	546705	503597	427763	459933
Adjusted R^2	0.064	0.930	0.347	0.638

t statistics in parentheses

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Table C.6: Economic gaps (continued)

Panel (b) displays results from the regression discontinuity design. The main specification imposes a 10% bandwidth (i.e. regression sample only includes elections for which the vote margin is at most 10% between the winner and the runner-up). Optimal bandwidths are typically wider, and results with optimal bandwidths are reported in the appendix. For observations at election i, geography (e.g. county) j, and event time (e.g. month or year) t, the following regression specification is run:

$$Y_{ijt} = \alpha + \beta_1 \mathbf{1}\left\{\text{vote margin} > 0\right\}_{it} + \gamma_0 \left[\text{vote margin}\right]^{-}_{it} + \gamma_1 \left[\text{vote margin}\right]^{+}_{it} + \varepsilon_{ijt}$$

where "vote margin" is the percentage point vote margin between the black candidate and the white candidate. The first indicator term denotes a dummy for elections in which the black candidate won. The next two terms are linear controls separately for elections where the white candidate won (i.e. negative vote margin) and where the black candidate won (i.e. positive vote margin). Observations are weighted by the fraction of the geography (e.g. county) affected by the election. Panel (a) displays the regression discontinuity results for Race IAT D scores from Project Implicit Database.

(b) Regression discontinuity

	labor		mortgage	
	(1) unemployment transition	(2) employment to pop	(3) log origination to pop	(4) rejection rate
Black winner	0.025**	-0.050	-0.256**	0.035**
	(2.24)	(-1.55)	(-2.32)	(2.38)
Margin (winner)	0.000	-0.008	-0.017	0.001
	(0.12)	(-1.58)	(-1.19)	(0.54)
Margin (loser)	-0.003**	0.012**	0.024*	-0.001
	(-2.24)	(2.27)	(1.83)	(-0.37)
Constant	-0.045***	0.025	-1.028***	0.178***
	(-4.96)	(1.00)	(-13.16)	(18.45)
Observations	4510	4376	4311	4166
Adjusted R^2	0.004	0.023	0.024	0.027

t statistics in parentheses

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Table C.7: Heterogeneity

Both panels display results from the regression discontinuity design, with heterogeneous effects. The specification imposes a 10% bandwidth (i.e. regression sample only includes elections for which the vote margin is at most 10% between the winner and the runner-up). For observations at election i, geography (e.g. county) j, and event time (e.g. month or year) t, the following regression specification is run and I report γ_1^k:

$$Y_{ijt} = \sum_k \left\{ \alpha^k + \gamma_1^k \mathbf{1}\left\{\text{vote margin} > 0\right\}_{it} + \delta_0^k \left[\text{vote margin}\right]_{it}^- + \delta_1^k \left[\text{vote margin}\right]_{it}^+ \right\} \mathbf{1}\left\{\text{in sub-group}\right\}_{ij} + \eta_{ijt}$$

where "vote margin" is the percentage point vote margin between the black candidate and the white candidate. The first indicator term denotes a dummy for elections in which the black candidate won. The next two terms are linear controls separately for elections where the white candidate won (i.e. negative vote margin) and where the black candidate won (i.e. positive vote margin). Observations are weighted by the fraction of the geography (e.g. county) affected by the election.
In panel (a), the outcome variable Y_{ijt} is the raw and composition-adjusted Race IAT D scores, and there are two groups k where the sample is split along the median of the variable in the column header. In panel (b), the groups k are for above-median and below-median Race IAT D score level averaged for 2003-2017, while column headers indicate the outcome variable Y_{ijt}.

(a) Heterogeneity by

	IAT		black pop		income	
	(1)	(2)	(3)	(4)	(5)	(6)
	ex-demo FE	raw	ex-demo FE	raw	ex-demo FE	raw
Black winner	-0.013	-0.007	-0.000	0.003	0.039**	0.041**
	(-1.05)	(-0.55)	(-0.03)	(0.24)	(2.07)	(2.13)
HIwinner	0.047***	0.043***	0.025*	0.024*	-0.026	-0.024
	(3.75)	(3.26)	(1.81)	(1.77)	(-1.46)	(-1.35)
Observations	23590	23590	23590	23590	23590	23590
Adjusted R^2	0.015	0.014	0.003	0.004	0.002	0.003

t statistics in parentheses

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Table C.7: Heterogeneity (continued)

(b) Heterogeneity by average IAT level

	labor		mortgage	
	(1) unemployment transition	(2) employment to pop	(3) log origination to pop	(4) rejection rate
Black winner	0.013	-0.055	-0.131	0.011
	(0.93)	(-1.33)	(-0.96)	(0.52)
Black winner x high ______	0.020	0.012	-0.173	0.030
	(1.53)	(0.28)	(-1.13)	(1.30)
Observations	4510	4376	4311	4166
Adjusted R^2	0.005	0.025	0.032	0.060

t statistics in parentheses

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Table C.8: IV estimates

This table displays the second-stage estimates of an instrumental variables regression, using the regression discontinuity (dummy for the black candidate's victory) as the instrument. The main specification imposes a 10% bandwidth. For observations at election i, geography (e.g. county) j, and event time (e.g. month or year) t, the following regression specification is run, and I report β_1 for different outcome variable Y_{ijt}:

$$Y_{ijt} = \beta_0 + \beta_1 \text{IAT}_{ijt} + \tilde{\delta}_0 \left[\text{vote margin}\right]^-_{it} + \tilde{\delta}_1 \left[\text{vote margin}\right]^+_{it} + \varepsilon_{ijt}$$
$$\text{IAT}_{ijt} = \alpha + \gamma_1 \mathbf{1}\left\{\text{vote margin} > 0\right\}_{it} + \delta_0 \left[\text{vote margin}\right]^-_{it} + \delta_1 \left[\text{vote margin}\right]^+_{it} + \eta_{ijt}$$

where "vote margin" is the percentage point vote margin between the black candidate and the white candidate. The first indicator term denotes a dummy for elections in which the black candidate won. The next two terms are linear controls separately for elections where the white candidate won (i.e. negative vote margin) and where the black candidate won (i.e. positive vote margin). Observations are weighted by the fraction of the geography (e.g. county) affected by the election.

	labor		mortgage	
	(1) unemployment transition	(2) employment to pop	(3) log origination to pop	(4) rejection rate
IAT (adj)	0.439*	-1.224	-4.367*	0.639*
	(1.88)	(-1.55)	(-1.70)	(1.80)
Margin (winner)	0.003	-0.012**	-0.032**	0.006***
	(1.39)	(-2.28)	(-2.25)	(2.62)
Margin (loser)	-0.002	0.010	0.024	-0.001
	(-1.33)	(1.52)	(1.33)	(-0.30)
Constant	-0.041***	0.015	-1.091***	0.174***
	(-3.57)	(0.47)	(-10.23)	(11.78)
Observations	3265	3188	3099	3015
Adjusted R^2	.	.	.	.

t statistics in parentheses

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$