Racial Differences in the Total Rate of Return on Owner-Occupied Housing

Rebecca Diamond
Stanford GSB and NBER

William Diamond
Wharton

August 29, 2024

Abstract

We quantify racial differences in the total rate of return on housing in a nationally representative sample of homeowners from 1974-2021. We develop a new method to estimate the rental value of each owner-occupied house, using a house's resale value to proxy for unobservable quality. Black and Hispanic homeowners earn higher but more volatile rates of return than White homeowners, due in part to higher rental yields. These differences are largely explained by household income and education differences. Including returns on properties that do not transact in sample is crucial to accurately measure racial differences in total rates of return.

*First draft: November 18, 2023. We thank our research assistants Alex Almeida, Sam Asher, Marie Kaul, Weldon Lin, Alexis Payne, and Gustav Yang for outstanding work. We also thank Amir Kermani, Matthijs Korevaar, Mark Garmaise, and Kelly Shue for excellent discussions. Rebecca Diamond acknowledges support from the Stanford Graduate School of Business, the National Science Foundation (CAREER Grant 1848036).

The largest single component of most households' wealth is their home [Campbell, 2006]. The rate of return earned on an investment in housing is therefore a key determinant of a household's ability to accumulate wealth. Unlike purely financial assets, an investment in housing provides a flow of consumption services (the ability to live in or rent out a home), financial costs such as taxes, and potential capital gains from the increase in a home's price. This paper develops a new method to infer the rental value of owner-occupied housing and uses this method to measure the total rate of return homeowners get from a housing investment, inclusive of a home's rental value, property taxes, and maintenance and home improvement costs. We apply this method to studying disparities in the total returns on housing earned by homeowners of different racial groups. We find that Black and Hispanic homeowners earn higher but more volatile returns than White homeowners, primarily due to higher rental yields on the homes they buy.

Racial minorities have a disproportionate share of their wealth in housing (59% for Black vs 38% for White), so the return on housing plays a key role in the evolution of the racial wealth gap [Derenoncourt et al., 2023, Barsky et al., 2002, Kuhn et al., 2020]. In addition, the Black homeownership rate is roughly half the White homeownership rate [Gupta et al., 2023, Charles and Hurst, 2002], motivating interest from policymakers in increasing minority homeownership. Consistent with previous work [Kermani and Wong, 2022, Akbar et al., 2022, Wolff, 2022, Kahn, 2024], we find a lower rate of house price growth for Black homeowners than White homeowners. However, we also find that Black homeowners earn higher rental yields. On net, average racial disparities in the total return on owner-occupied housing push towards narrowing the racial wealth gap.

Using data on a panel of homes from the American Housing Survey (AHS) from 1974 to 2021, we observe a subset of homes switching between the rented and owner-occupied segments of the housing market. This allows us to observe the rental value and sale price of the same home and estimate an empirical relationship between rents and house prices. We therefore use a house's resale price as an additional predictor of its rental value beyond the information contained in the house's observable characteristics. We develop a method that uses the information in house prices to accurately predict the rents of both low-value and high-value homes. We also apply a version of our method to a much larger administrative database of housing transactions and rentals from Corelogic, finding similar results.

Our rental imputation method integrates two sub-literatures on measuring the return on housing. One branch attempts to measure the total return to landlords by jointly observing rents and prices on the same housing units [Chambers et al., 2021, Eichholtz et al., 2021], without attempting to extrapolate to the owner-occupied housing stock. Another branch [Demers and Eisfeldt, 2001, Gorback and Schubert, 2024, Halket et al., 2024, Gilbukh et al., 2023, Colonnello et al., 2024] imputes rental yields using hedonic house characteristics and/or aggregates data at a local level. Our innovation is to use information on an individual house's price to more accurately impute its rental value. As well as providing a new approach to measuring the total return on owner-occupied housing, our method could also be applied in other rent imputation contexts, such as the role of owners' imputed rent in computing CPI and GDP statistics [Crone et al., 2009, Garner and Short, 2009].

We apply our method to measuring differences in the rental yield of owner-occupied housing across racial groups using our AHS data. We find Black homeowners' average rent-to-price ratio is 8.86%, Hispanic homeowners' is 8.36%, and White homeowners' rent-to-price ratio is 7.83%. In a regression with year fixed effects, we find a Black-White rent-to-price ratio gap of 1.2 percentage points(pp) and Hispanic-White gap of 0.79 pp. These rental yields are inferred from a model fit to AHS data on houses that switch between being rented and owner-occupied. In the raw data on these "switchers," we find an average rent-to-price ratio of 9.91% for Black homeowners, 9.42% for Hispanic homeowners, and 9.04% for White homeowners. The lower yields we find for the overall owner-occupied housing stock than these switcher homes in particular reflects the fact that a house's rent-to-price ratio tends to decrease with its price. Because owner-occupied homes tend to be more expensive than observably equivalent rental homes, our rent imputation procedure is ideal for capturing these quality differences reflected in a home's price.

We find independent evidence of higher rental yields for Black and Hispanic homeowners than White homeowners in two additional data sources. First, we examine direct evidence in the Consumer Expenditure Survey (CEX), which surveys homeowners both about the value of their homes and about the rent they believe they could earn on their home. We find qualitatively similar but somewhat larger racial disparities in rental yields in the CEX than the AHS. In addition, we independently estimate the rental yield of owner-occupied housing using administrative data from Corelogic, covering 2007-2017. Here, we jointly use data on housing purchases and sales matched to rents from Multiple Listing Services (MLS) to observe a population of "switchers" with both rent and price observations in the subset of regions where the data is available. After imputing rental yields for owner-occupied homes in this much larger dataset, we find similar racial disparities in rental yields as a subsample from the AHS that matches the geography and time period of our administrative data. In regressions with equivalent year/time fixed effects in our two datasets, we find a 1.67 pp rental yield gap between Black and White homeowners in Corelogic and a 1.47 pp gap in the AHS. We similarly find a 1.0 pp Hispanic-White gap in Corelogic and 0.92 pp Hispanic-White gap in the AHS. In addition, we show that roughly 30% of these racial disparities cannot be accounted for by zipcode-year fixed effects and would be missed if we aggregated our rental yield estimates at a zipcode level [Demers and Eisfeldt, 2001].

Next, we use our rental yield estimates to analyze racial disparities in total rates of return on housing. We find in our benchmark result with AHS data that the average log return on owner-occupied housing is 7.2% for Black homeowners, 7.9% for Hispanic homeowners, and 6.5% for White homeowners from 1974 to 2021. This gap is not driven by differences in house price appreciation. We find White homeowners earn 0.9% annual real appreciation rates, while Black homeowners earn 0.7%, consistent with prior results of a negative Black-White house price appreciation gap. We find Hispanics earn a higher average appreciation rate of 1.6%. The overall rate of return for Black homeowners is higher due to a 0.7 pp higher dividend yield than for White homeowners. Hispanic homeowners earn a 0.6 pp higher dividend yield than Whites. These dividend yield differences are lower than the rental yield differences since we find that Black and Hispanic homeowners spend more on home maintenance relative to the value of their homes than White homeowners. We verify that the racial disparities in property taxes we measure are consistent with Avenancio-Leon and Howard [2022], and these differences have a small impact on the total rate of return on housing. Our average rate of return is similar to the housing return in Jorda et al. [2019].

We use our rich AHS microdata to analyze the determinants of these rate of return differences at the household level. We regress the log rate of return on each house on the homeowner's race together with a growing list of control variables. After controlling for homeowner education and income, we are able to fully explain the average racial return differences with observable characteristics. We find that income/education is the most useful explainer of the Black-White total return gap, and both geography and income/education help to explain the Hispanic-White total return gap. Our evidence is consistent with lower income homeowners buying lower value homes, which in turn have higher average rates of return due to their higher rental yields. The racial disparities in housing returns we measure can mostly be accounted for by disparities in economic status across racial groups.

A second key difference in the total return on housing for Black and Hispanic homeowners relative to White homeowners is greater volatility and greater sensitivity to the business cycle. We find that both Black and Hispanic homeowners have particularly volatile and cyclical returns. Unlike the disparities in average returns, we find that observable differences between racial groups do not entirely explain the excess exposure of minority housing returns to the business cycle. Controlling for geography and income/education largely explains the White-Hispanic difference in return cyclicality, but observables poorly account for lower returns earned by Black homeowners in recessions. We consider the contribution of both rental yields and price changes to the cyclicality of returns. Even after accounting for observables, we find that Black and Hispanic homeowners have both rental yields and price appreciation that are more cyclical than for White homeowners. This is broadly consistent with the finding [Kermani and Wong, 2022] that the risk in minority housing returns is largely due to distressed sales during downturns, based on more severe financial constraints [Gupta et al., 2023, Bhutta and Hizmo, 2021] for Black and Hispanic homeowners.

Our paper contributes to the literature on disparities in housing returns by the use of nationally representative, household-level data that allows us to measure the key determinants of housing returns over 45 years. Much of the previous literature on racial disparities in housing returns uses geographically aggregated data, with Kermani and Wong [2022] as a notable exception. Our AHS data has costs and benefits compared to the large transaction-level data used in Kermani and Wong [2022], and the final section of our paper reconciles our results. Their baseline analysis uses realized housing transactions, which more accurately measures the returns on transacted homes but does not include returns for the majority of houses that do not transact within their sample. Our data provides a housing value even if there is not a transaction, but it may miss some of the financial distress they observe in sales prices. We explore the robustness of our results to potential mismeasurement of distressed sale prices in section 7 and find these effects to be small. The fact that we can measure the return on non-transacted properties, instead of conditioning on a selected sample of transactions which must include all foreclosures, decreases the sensitivity of our results to foreclosures. When we attempt to fill in the return on non-transacted properties in Corelogic transaction data, we find similar racial disparities in house price appreciation as we do in the AHS. If we make the conservative assumption that our results in the AHS entirely miss the impact of foreclosures, we show that Black and Hispanic homeowners still earn a higher total rate of return than White homeowners after we subtract a measure of foreclosure discounts from Kermani and Wong [2022].

In total, our paper demonstrates that it is crucial to account for the total return on housing, not just changes in house prices, when measuring racial disparities in housing returns. Our new method for imputing the rental value of owner-occupied housing may therefore be a broadly useful tool for this literature. Because lower value homes tend to have the highest rental yields, any method of measuring rates of return that ignores rents will understate the returns on low-value homes and overstate the returns on high-value homes. Because the average value of homeowners' residences varies by racial group, an accurate rental yield calculation is crucial for understanding racial disparities in the rate of return on housing. Together with other evidence on racial disparities in liquidity constraints [Ganong et al., 2024], home-ownership Charles and Hurst [2002], and credit market outcomes [Indarte et al., 2023], an accurate total rate of return measure for housing is key for understanding the future of the racial wealth gap.

1 Conceptual Framework

Unlike financial assets such as stocks or bonds, housing is both an investment good and a consumption good. Because a homeowner can live in their house, investing in housing generates a flow of consumption services as well as a possible increase in the value of the investment. To calculate the financial return to owning a home, we compare the financial benefit of homeownership to renting an identical home in a given year t. If a household were to rent house j in year t, it would pay r_{jt}. If the household were to purchase the home, it would pay price P_{jt}, spend $main_{jt}$ on maintenance, and tax_{jt} on property taxes and then receive P_{jt+1}, ignoring transaction costs.

The log financial return to owning house j, relative to renting it is:[1]

$$log(P_{jt+1}) - log(P_{jt} - r_{jt} + tax_{jt} + maint_{jt}). \qquad (1)$$

This return quantifies the *financial* benefit a household receives by owning a given house, instead of renting it. It is possible that if a homeowner switched to renting, it would prefer to rent a different home. While it would be more complicated to explicitly quantify homeowners' utility or welfare, the financial return we measure in equation 1 tells us the impact of homeownership on a household's wealth, compared to a counterfactual of paying rent for an identical home.

Taxes and maintenance (for which we include both routine maintenance and home improvement) in equation 1 are cashflows that can directly be observed in our AHS data. However, the rent r_{jt} for an

[1] The "dividend" component of the housing return, $r_{jt} - tax_{jt} - maint_{jt}$, we will assume is received at time t. One could also write the rate of return calculation with dividends paid at time $t+1$. While both of these are conceptually very similar, we have much better data coverage for all of these variables at time t than $t+1$, since some houses that are owned at time t are not still owner occupied at $t+1$.

owner-occupied house is never explicitly observed. As a result, much of the literature on inequality in returns has simply used the price appreciation $log(\frac{P_{t+1}}{P_t})$ to measure housing returns. To include all of the relevant cashflows to measure housing returns, we develop and apply a new method to impute the rental value of owner-occupied housing.

2 Data

The bulk of our analysis uses the American Housing Survey (AHS) from 1974-2021. This is a nationally representative survey that is a panel at the housing unit level. From 1974-1981, the survey was done annually, but it switched to every other year in 1983. The AHS interviews the residents of housing units about many aspects of the house itself, many details of their housing costs, and about the housing residents themselves. This survey is ideal for calculating rates of return on housing since it collects raw microdata on housing maintenance, home improvements, property taxes, property values, property rents, and property characteristics. Resident race, which papers using housing transaction data often have to impute from external sources, is also directly reported. The main concern with using survey data is the possibility of measurement error in these variables, as compared to administrative data on property transactions. We investigate these differences in Section 7. Two clear benefits of performing this analysis in the AHS are the very long time series back to 1974 and the fact that the sample is nationally representative. Most datasets of administrative housing transactions begin in 1990s or 2000s [Kermani and Wong, 2022, Goldsmith-Pinkham and Shue, 2023] and exclude regions that do not publicly report housing transaction prices. In addition, the AHS allows us to calculate rates of return for houses that do not sell in sample, while administrative transaction data is restricted to properties that sell in sample. We will investigate all of these differences between the AHS and administrative house transactions data in section 7.

Table 2.1 reports summary statistics on key housing costs in 2019 dollars in the AHS by race. Black homeowners have lower average property values, rents, home improvement spending, and property taxes than White homeowners, but they have similar spending levels on routine maintenance. Hispanic homeowners have similar housing values to Whites, but they have higher spending on routine maintenance and lower spending on home improvement and property taxes.

Table 2.1: Summary Statistics for Owners and Renters in the AHS by Racial Group, 1974-2021

	Num. Obs.	Mean	Median	Std. Dev.
White				
Property value	711,651	239,470	175,780	258,446
Rent (monthly)	285,090	885	757	669
Routine maintenance (annual)	710,935	767	568	1,121
Home improvement (annual)	710,934	2,380	982	8,095
Property tax (annual)	710,933	2,782	1,949	3,435
Black				
Property value	69,374	168,834	119,530	207,478
Rent (monthly)	87,059	729	658	523
Routine maintenance (annual)	69,225	765	556	1,136
Home improvement (annual)	69,225	1,693	521	8,843
Property tax (annual)	69,225	1,985	1,294	2,744
Hispanic				
Property value	54,813	242,293	175,264	258,819
Rent (monthly)	65,810	930	830	628
Routine maintenance (annual)	54,695	880	570	1,296
Home improvement (annual)	54,695	2,117	384	7,428
Property tax (annual)	54,695	2,663	1,857	3,353

Notes: Data from the 1974-2021 AHS. Summary statistics for all data except rent are from owner-occupied homes. Rent is from rented homes.

3 Imputing the Rental Value of Owner-Occupied Housing

One of the most important components of the return on investing in housing is the rental income that a house generates. For a landlord, this is a cash flow paid explicitly by the tenant to the landlord. For a homeowner, the "owner's equivalent rent" measures the market value of the housing services the household consumes by living in that home for a period of time. Because homeowners do not explicitly pay or receive rent, it is necessary to impute this owner's equivalent rent to accurately measure their investment returns.

We develop a new method of imputing owner's equivalent rent that allows for differences in quality across houses that are not captured by their observable characteristics. A house's quality depends on its characteristics X_h, including both observable components X_h^o and unobservable components X_h^u. As shown in [Bajari and Benkard, 2005, Bajari et al., 2005], even when preferences are heterogeneous across households, there exists an equilibrium pricing function that maps house characteristics X_h (this includes physical and location characteristics) to equilibrium prices in the market. For house h with characteristics X_h at time t, the rent r_{ht} and purchase price p_{ht} can be written as $r_{ht} = r_t(X_h)$ and $p_{ht} = p_t(X_h)$, respectively. While we assume households can observe all characteristics X_h of a house, our data only include the subset X_h^o. Without loss of generality, we can write our rent and price functions as $r_{ht} = r_t(X_h^o, X_h^u)$ and $p_{ht} = p_t(X_h^o, X_h^u)$.

The key insight to our method is that the price a house would sell for in the purchase market is a proxy for the unobserved house characteristics, given the observed house characteristics. For example, consider two different 1400 square foot, 2 bed, 2 bath houses built in 1974 located in a given neighborhood. If one would sell for 10% more than the other, despite the observed characteristics being identical, the house with the higher sale price would likely also rent for more, although the gap in rental rates might be more or less than 10%. To formalize this intuition, we make the following key single index assumption.

Assumption 1 *(Single Index)*

There exists at each time t a function f_t that maps the set of unobserved property characteristics X_h^u to a scalar such that the rent and price functions r_t and p_t can written as a function of the observed house characteristics X_h^o and $f_t(X_h^u)$ where p_t is strictly monotone in $f_t(X_h^u)$, given X_h^o. E.g.:

$$
\begin{aligned}
r_{ht} &= r_t(X_h^o, X_h^u) \\
&= r_t(X_h^o, f_t(X_h^u)) \\
p_{ht} &= p_t(X_h^o, X_h^u) \\
&= p_t(X_h^o, f_t(X_h^u)).
\end{aligned}
$$

Under Assumption 1, we can invert the p_t function: $f_t(X_h^u) = p_t^{-1}(X_h^o, p_{ht})$ and plug the inverted price function into the rent function:

$$
\begin{aligned}
r_{ht} &= r_t(X_h^o, X_h^u) && (2) \\
&= r_t(X_h^o, p_t^{-1}(X_h^o, p_{ht})) && (3) \\
&= r_t^*(X_h^o, p_{ht}). && (4)
\end{aligned}
$$

Equation 4 shows that r_{ht} can be written as a flexible function of observed house characteristics, X_h^o and the purchase price of the home, p_{ht}.

To the extent that the purchase price of a house is predictive of its rental value even after controlling for observables X_h^o, a clear bias appears in any procedure that only uses X_h^o to predict rent. For houses with sufficiently low prices, their rent will generally be below $E(r_h|X_h^o)$, so the procedure overestimates the house's rent. Similarly, for houses with high enough prices, $E(r_h|X_h^o)$ is an underestimate of the houses' rent.

Because minority homeowners tend to own lower-price houses than White homeowners, this argument suggests that using only observables to predict a house's rent would induce bias into our calculation of race-specific rates of return on housing. As we show empirically in appendix G, the rental yield earned by minority homeowners would be artificially inflated compared to the rental yield earned by White homeowners if we ignored unobservable differences between homes. As we show empirically in section 4, the procedure we develop next accurately measures average rents for both low- and high- price homes and is therefore ideal for comparing rates of return on housing across racial groups.

3.1 Inferring rental yields using switchers in the AHS

Our baseline rent imputation procedure requires data on a panel of houses, as is provided by the AHS. Each house h at time t is either vacant, rented, or owner-occupied. For rented houses, we observe the rent r_{ht}. For owner-occupied houses, we observed the price $p_{h,t}$. The key difficulty we face is that in each period at most one of r_{ht} and p_{ht} is observed. In addition, we see a vector of house characteristics X_h^o. Our goal is to produce an estimate $\hat{r_t}(X_h^o, p_{ht})$, the rent that could have been obtained for owner-occupied house h at time t.

Equation 4 implies that if we observed the sale price of a house along with observed property characteristics, we should be able to *perfectly* predict its rental price. Since the data are unlikely to provide a perfect rent forecast, we augment equation 4 to include an error term:

$$r_{ht} \quad = \quad r_t^*(X_h^o, p_{ht}) + \epsilon_{ht}. \tag{5}$$

Equation 5 reflects the possibility of measurement error in rents, or that our single index assumption above does not hold perfectly. See appendix A for a precise statement of our necessary assumptions to consistently impute rents, which do not require a perfect single index. For a given set of houses with identical (X_h^o, p_{ht}), we should expect to observe a distribution of rents if ϵ_{ht} has nonzero variance.

Our rent imputation procedure relies on the population of "switcher" houses that switch between owned and rented status between times t and $t+1$. We only can see data for both the price and rent of a single house if that house is a switcher. However, only a small fraction of the houses in our sample at any given time are switchers. In theory, our model suggests that we should estimate the relationship between rents and prices year-by-year, since the hedonic price of a given house changes over time, but we would under-powered due to the low sample size. We therefore develop methods that enable us to pool the sample of switcher houses across many years of our data.

To do so, we transform the relationship our single index provides between rents and prices into one between rent and price ranks. Let ρ_{ht}^r and ρ_{ht}^p respectively be the ranks of house h's rent and price at time t in the nationwide rent and price distributions. If R_t and P_t are respectively the CDFs of the national rent and distributions at time t, we have without loss of generality

$$r_{ht} \quad = \quad r_t^*(X_{ht}^o, p_{ht})) \tag{6}$$

$$\rho_{ht}^r \quad = \quad R_t(r_t^*(X_{ht}^o, P_t^{-1}(\rho_{ht}^p))) \tag{7}$$

$$\rho_{ht}^r \quad = \quad r_t^{*,\rho}(X_{ht}^o, \rho_{ht}^p). \tag{8}$$

This gives us a mapping between the rent and price rank of a given house h at a time t. This mapping is constructed by composing the CDF of the time t rent distribution, our function r_t^* at a given value of observables X_h^o, and the inverse CDF of the time t house price distribution. We then assume that this $r_t^{*,\rho}$ is the same function $r^{*,\rho}$ at all times t. This allows us to estimate a relationship

$$\rho_{ht}^r \quad = \quad r^{*,\rho}(X_h^o, \rho_{ht}^p) \tag{9}$$

between rent and price ranks using data from all times t.

This assumption requires that among houses with the same observables, a given price rank ρ_{ht}^p maps to the same rent rank ρ_{ht}^r, even if the houses are observed in different years. In this relation, these ranks are computed relative to year-specific nationwide distributions. For example, consider all 2-bed 2-bath houses in the years 2000 and 2010. If a house in 2000 has the median price rank and 70th percentile rent rank relative to homes in 2000, it must also be true that the house in 2010 with the median price rank has the 70th percentile rent rank relative to homes in 2010. However, if prices increase and rents fall from 2000 to 2010, a home in 2010 could have a higher price and lower rent than another home in 2000.

We believe this assumption that ranks have a stable relationship over time is more likely to hold than a stable relationship between rents and prices themselves. First, the overall rent-to-price ratio for houses has decreased over time [Davis et al., 2008], similar to the decline in interest rates on bonds since the 1980s. Second, we provide evidence in section 3.3 that the empirical relationship between rank ranks and price ranks is similar across years, while the empirical relationship between dollar rents and dollar prices is not.[2]

3.2 Estimation

Before turning to our rent estimation procedure, we present some descriptive statistics in Table 3.1 on houses that switch between the owner-occupied and rental markets and compare them to units in the overall rental market and owner-occupied market by race. Among White-, Black-, and Hispanic-occupied units, we see the average units that switch between markets are are the 42nd, 33rd, and 43rd percentile of the property value distribution, respectively. These switcher units are negatively selected from the overall owner-occupied distribution, where the average property values for White, Black, and Hispanic homeowners are at the 51st, 36th, and 48th percentile of the property value distribution. However, these switchers are positively selected from the rental market distribution. The average White-, Black-, and Hispanic-occupied rentals are at the 51st, 41st, and 52nd percentiles, respectively, while the switcher units are higher at the 56th, 48th, and 56th percentiles. Taken together, this implies that the average rental until is of lower value than the average owner-occupied unit, highlighting the importance of using property prices to proxy for unobserved quality differences between owned and rented units. These relative differences in quality between owned, rented, and switcher units are also apparent when looking and differences in number of rooms, number of bedrooms, number of bathrooms, and age of the unit.

[2] One additional detail we abstract from here is that for a given house, we observe its rent and its price at different points in time. We state the precise assumption we need to cover this case, together with the possibility of measurement error, in appendix A.

Table 3.1: Property Characteristics by Switcher, Owner-occupied, and Renter-occupied

(1) White	Switchers	Owners	Renters
Number of Rooms	5.53	6.24	4.36
	(1.59)	(1.66)	(1.49)
Number of Bedrooms	2.66	3.00	1.96
	(0.91)	(0.87)	(0.95)
Number of Bathrooms	1.46	1.68	1.22
	(0.60)	(0.67)	(0.46)
Age of the Building	36.06	31.76	34.66
	(20.52)	(20.10)	(20.80)
Number of Units in the Building	1.81	1.23	4.54
	(2.30)	(1.28)	(3.83)
Price Quantile	41.56	50.94	
	(28.78)	(28.53)	
Rent Quantile	55.67		51.48
	(30.69)		(28.98)
Num. Obs.	31,468	751,606	305,684

(2) Black	Switchers	Owners	Renters
Number of Rooms	5.60	6.14	4.38
	(1.50)	(1.53)	(1.36)
Number of Bedrooms	2.79	3.05	2.03
	(0.91)	(0.84)	(0.93)
Number of Bathrooms	1.40	1.55	1.18
	(0.58)	(0.64)	(0.42)
Age of the Building	41.99	36.34	39.66
	(20.74)	(20.62)	(20.26)
Number of Units in the Building	1.82	1.23	4.98
	(2.29)	(1.23)	(3.84)
Price Quantile	33.42	36.19	
	(26.74)	(26.65)	
Rent Quantile	48.27		41.39
	(29.80)		(27.74)
Num. Obs.	5,081	68,305	88,456

(3) Hispanic	Switchers	Owners	Renters
Number of Rooms	5.42	5.97	4.31
	(1.45)	(1.51)	(1.31)
Number of Bedrooms	2.75	3.04	2.00
	(0.93)	(0.86)	(0.92)
Number of Bathrooms	1.51	1.69	1.22
	(0.59)	(0.64)	(0.46)
Age of the Building	41.27	36.59	42.55
	(21.09)	(21.70)	(20.90)
Number of Units in the Building	1.87	1.29	5.02
	(2.38)	(1.43)	(3.86)
Price Quantile	43.20	47.64	
	(29.99)	(29.32)	
Rent Quantile	56.05		52.37
	(29.45)		(27.06)
Num. Obs.	5,128	53,159	67,957

Note: Summary stats of AHS data from 1974-2021. Switchers are housing units that switch between being owner-occupied and renter-occupied between consecutive waves of the AHS. Owners are all owner-occupied units. Renters are all rent-occupied units. We drop observations occupied by other racial groups than Black, White, and Hispanic. Price quantile is the percentile of the house's price within the nationwide distribution of property values observed in the same survey wave. Rent quantile is the percentile the house's rent within the national distribution of property rents observed in the same survey wave.

To implement our rent prediction method, we first need to predict a house's rent rank given its observable characteristics and price rank as in equation 9. To do so, we estimate the regression equation

$$\rho^r_{ht} = \alpha + \beta_1 \rho^p_{ht} + \beta_2 (\rho^p_{ht})^2 + \beta^o X^o_{ht} + \epsilon_i. \tag{10}$$

Because our predicted rent at a time t plugs this rank prediction into the non-linear function R_t, a mean-unbiased rent rank prediction need not result in a mean-unbiased prediction of rent. As a result, we aim instead to construct a median-unbiased rent prediction. To do so, we estimate equation 10 using a median regression.

Our characteristics X^o_{ht} include building age and its square, number of rooms, bedrooms, and bathrooms, units in the structure and its square, dummies for a townhouse and mobile home and the metro-census region level to account for the finding in Demers and Eisfeldt [2001] that rent-to-price ratios vary systematically across regions at a given time.[3] We also include a dummy indicating whether the property was rented in the prior period (and is now owner-occupied), versus previously owner occupied, but now rented. To allow for the possibility that unobserved housing characteristics not captured by price impact rents in a way that is correlated with race, we estimate these regressions separately for each racial group.

Finally, we use a non-parametric estimate of $\hat{R}_t^{-1}$, the CDF R_t of rents at time t. We then simply use the empirical CDF for $p_{h,t}$ as our estimate of the CDF P_t of house prices in a given time t. The imputed rent for each owner-occupied house h is:

$$\hat{r}_{ht} = \hat{R}_t^{-1}(\hat{\rho}^r_{ht}(\hat{P}_t(p_{ht}), X^o_{ht}). \tag{11}$$

In equation 11, we take a house's realized price and first compute its empirical rank $\hat{P}_t(p_{ht})$ using the empirical house price CDF $\hat{P}_t$. Next, we use this price rank and the house's observables to predict its rent rank from our estimate of equation 10. Finally, we plug this estimate of the house's rent rank into the inverse CDF of rents $\hat{R}_t^{-1}$ to obtain an estimate of the house's rent. This rent prediction uses both a house's observables as well as unobservable quality proxied for by its price and is potentially an improvement over hedonic prediction methods that use observables alone.

3.3 Verifying the rank stability assumption

To investigate whether the rank stability assumption behind our rent imputation method is empirically plausible, we estimate the rank-rank regression in equation 10. We run this regression separately year by year on the subset of houses where we see both rent and price. We drop the quadratic term on price ranks and only focus on the linear coefficient as a single summary metric for the relationship between price ranks and rent ranks in each year. For comparison, we also run the same regression, except we regress a house's dollar rent directly on its dollar house price, instead of using ranks. In doing so, we convert all prices into 2019 dollars. To make the units of these regression coefficients comparable between these methods, we replace each rent and price observation by its Z-score across the entire sample to put things in comparable units. We do this both for the rank measures and the dollar measures.

We plot the coefficients on price (both for ranks and dollars) in Figure 3.1. The relationship between dollar rents and dollar prices is unstable over time, with our regression coefficient decreasing over the sample. This coefficient decrease roughly tracks the aggregate decrease in rent-to-price ratios of houses

[3] We also include dummies for each variable if it is equal to the top or bottom code of a given continuous variable (e.g. 10+ bedrooms) and dummies for missing values of each of these variables.

over time. In contrast, our regression coefficient in the rank-rank regressions is quite stable over time, suggesting that our rank stability assumption is plausible.

Figure 3.1: Stability of rank-rank versus dollar-dollar relationship between rents and prices

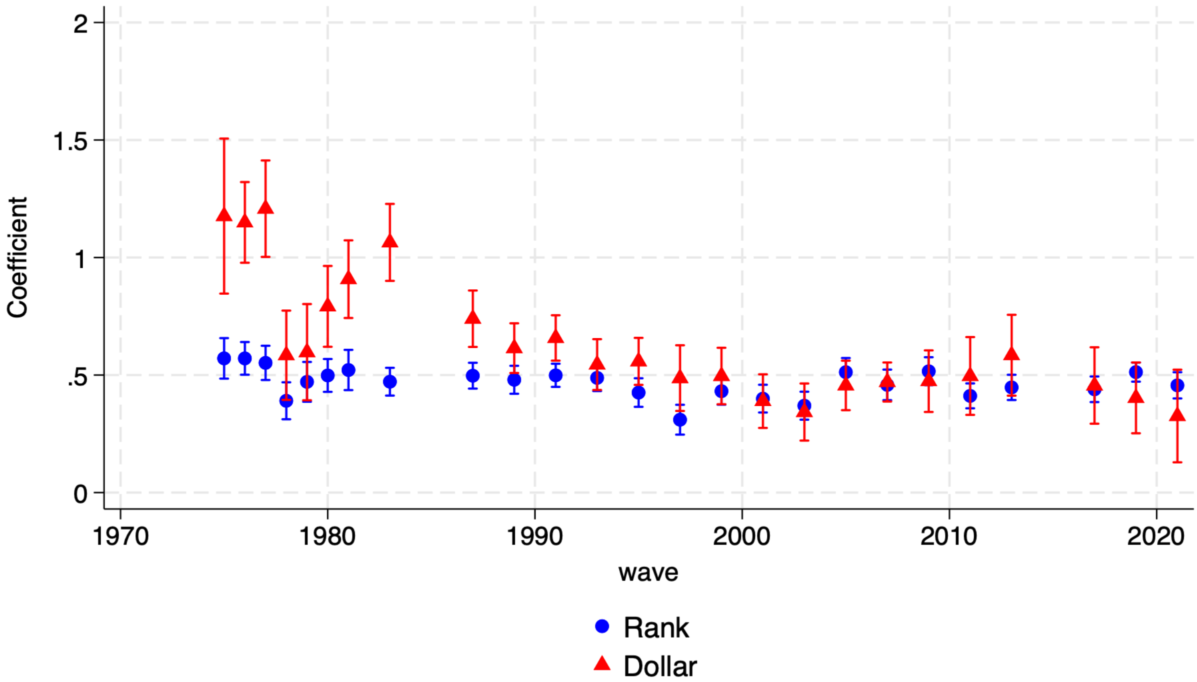

Notes: The figure plots the coefficient and 95% confidence interval of price in equation 10. The blue dots come from a median regression of rent ranks on price ranks, along with controls. The red triangles come from a median regression of rent dollars on price dollars, in units of 2019 dollars, along with controls. Both measures of rents and prices have been z-scored to make units comparable.

3.4 Implied Rental Yields by Racial Group

Table 3.2 reports summary statistics on the annual rent-to-price ratio by racial group from 1974 to 2021. We find Black homeowners' rents are equal to 8.86% of their property value on average. Hispanic homeowners earn rents equal to 8.36% of their property value, and White homeowners earn 7.83% of their property values. Both racial minority group homeowners earn higher rents as a share of their property values than White homeowners. This is also apparent in the raw data on property that switch between the owner and renter markets, as shown in Table 3.2 Black-owned properties that switch into/out of the rental market earn rents equal to 9.91% of their property value, Hispanic-owned switcher properties earn 9.42% of their value, and White-owned switchers earn 9.04% of their value. All of these rent-to-price ratios are higher for the switcher properties than the overall owner occupied housing stock. This reflects the fact that the properties that switch into/out of the rental market are on average lower value than the typical owner-occupied property.

Table 3.2: Rent-to-Price Ratio in the AHS by Owner Race, 1974-2021

Panel A: Switchers	Mean	Median	SD	25th Pct.	75th Pct.	Num Obs.
White	9.04	7.58	6.85	4.90	10.60	21,259
Black	9.91	8.23	7.37	5.10	12.11	4,100
Hispanic	9.42	7.44	7.61	4.56	11.32	4,586
Panel B: Owners	Mean	Median	SD	25th Pct.	75th Pct.	Num Obs.
White	7.83	7.06	4.11	5.52	8.82	697,025
Black	8.86	7.79	4.74	5.82	10.33	66,548
Hispanic	8.36	7.27	4.84	5.26	9.72	52,079

Note: Summary stats on rent-to-price ratios by race. Panel A reports raw data from our switcher sample. Panel B reports our imputed rent-to-price ratios for the owner-occupied sample using our switcher method in the AHS. Data are winsorized at the 2.5th and 97.5th percentiles to limit the effects of outliers.

Figure 3.2 plots the median rent-to-price ratio for homeowners by race over time. We see rent-to-price ratios falling in the 1970s, rising in the 1980s, staying flat in the 1990s, falling in the 2000s, rising in the early 2010s and then falling again until 2021. These patterns are broadly similar for all racial groups. Consistently over time, the Black rent-to-price ratios are about a percentage point higher than the White levels. Hispanic rent-to-price ratios are about 0.5 percentage points higher than White levels on average but are more cyclical and sometimes fall below White rent-to-price ratios.

Figure 3.2: Plot of Median Rent-to-Price Ratio by Owner Race and Year, Imputed Rent for all Homeowners

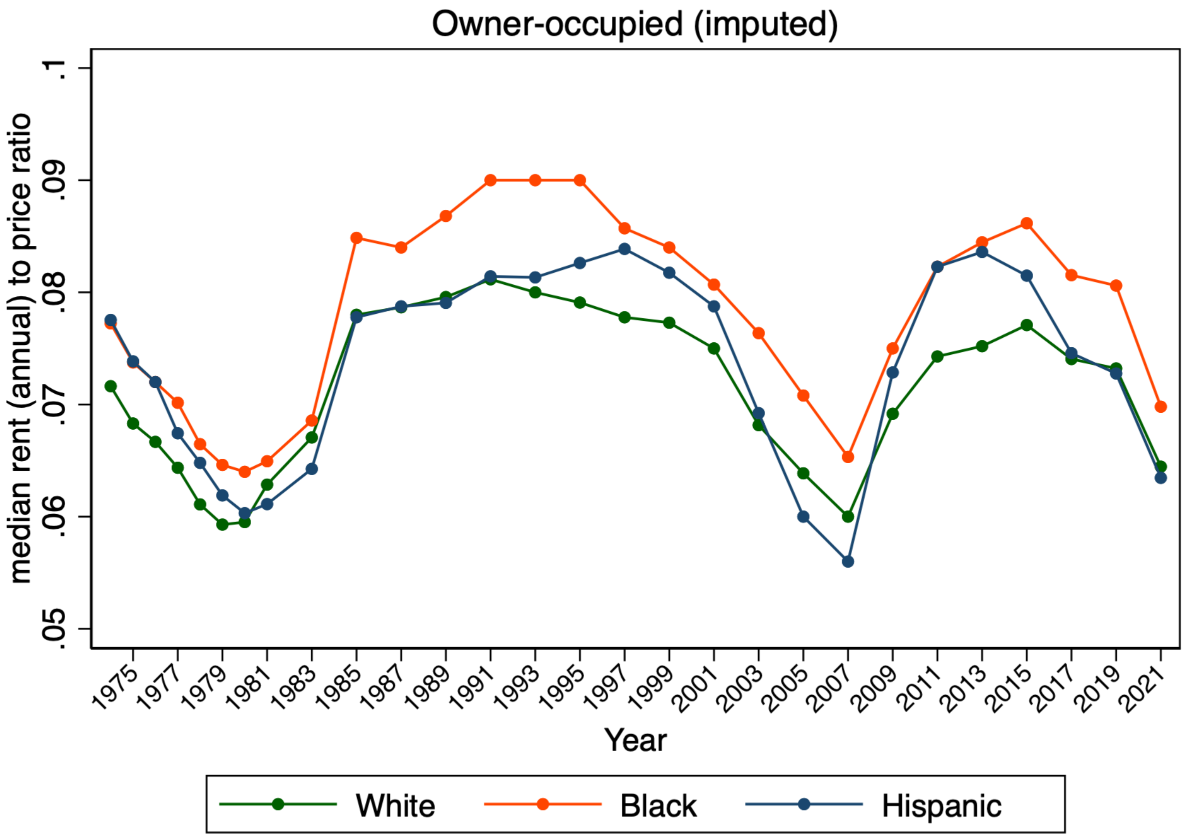

Note: Median annual rent-to-price ratios by race over time for the owner-occupied housing stock, as measured from our switcher imputation method in the AHS.

4 Comparing Rent imputation methods

We now compare our rent imputation method to simpler methods that only use the observable characteristics of a property. If a property has a rent r, observable characteristics X and a market price p, then we expect $E(r|X, p)$ to be increasing in the price p. As a result, $E(r|X)$ should overestimate the rent of low-price properties and underestimate the rent of high-price properties.

To impute a property's rent using only observable characteristics, we run the same median regression as above that we used to study rent ranks, but we switch the dependent variable to rents measured in dollars. We also drop the house sale price as a covariate. This standard hedonic price regression is then estimated on the entire sample of rental properties in the AHS, not just the subset of switchers.

Figure 4.1: Performance of Quantile-based and Hedonic Rent Imputation Methods on AHS Data

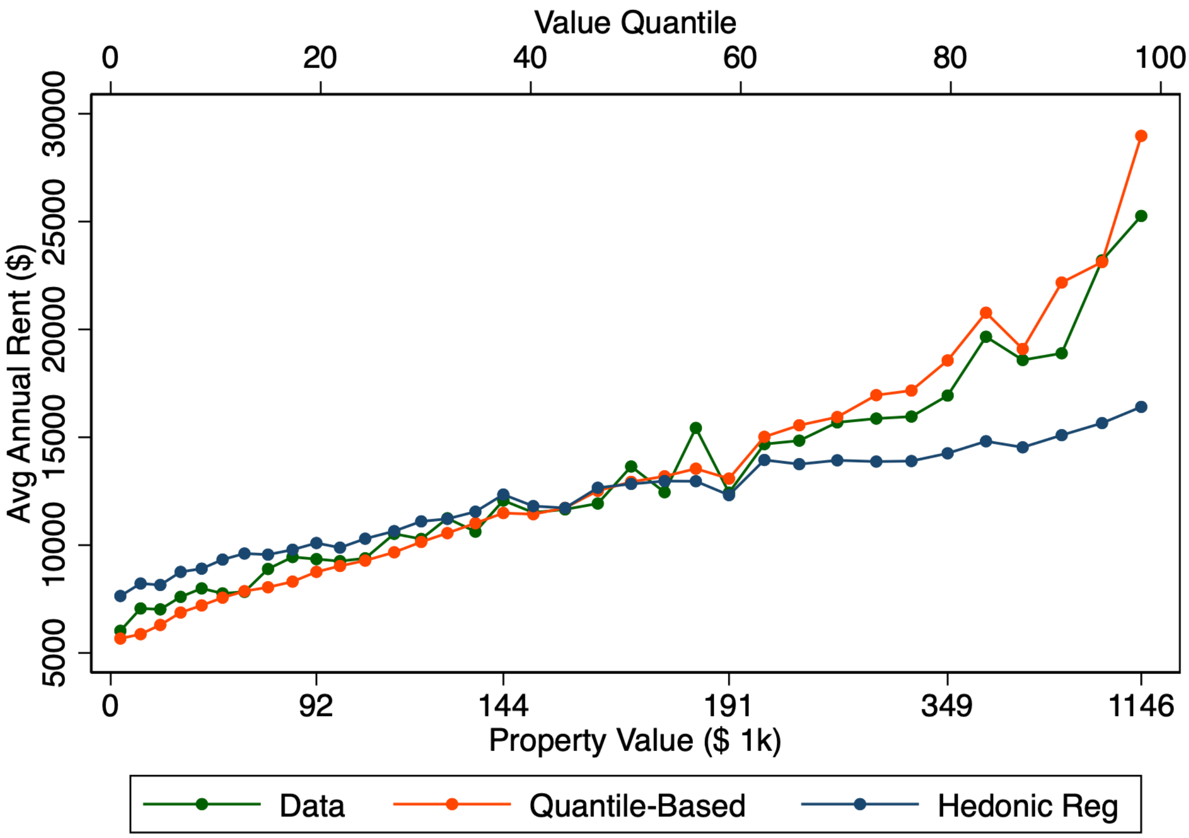

Note: This figure tests the results of our imputation methods against a 20% holdout sample of "switcher" houses with both rent and price data in the AHS from 1980 to 2019. All values are in 2019 dollars. The scale on the x-axis is equally spaced by quantiles of property value.

In each year, we use all rented properties in the AHS as our sample except for a random holdout sample of 20% of those properties that switch between being owned and rented. To test our paper's quantile-based method which uses a property's price to help predict its rent, we run our specification on the 80% of switcher properties that are not included in our holdout sample. We then test its predictions on our 20% holdout sample and present the results in figure 4.1. This graph shows the mean level of rent for homes of a given price in the raw switcher holdout data, based on the predictions of our quantile-based procedure that uses prices, and based on a hedonic regression that ignores prices.

Unlike our quantile-based procedure, the hedonic rent imputation over-predicts the rents of low-price properties and under-predicts the rent of high-price properties. The orange line representing the average predicted rent from our quantile-based procedure stays close to the green line that presents the average rent in the raw data of our holdout sample. For the bottom quartile of property prices the hedonic method predicts rents that are 25% too high, while our switcher method predicts rents that are within 5% of the true rents in the holdout sample. Similarly, in the top quartile of property values, the hedonic method predicts rents that are 25% too low, while our swither method predicts rents within 5% of the true rents.

Using a method that accurately estimates the expected rent $E(r|p)$ of a property with a given price p is crucial for our application to understanding racial differences in housing returns. Because minority homeowners tend to live in lower-price houses on average than White homeowners, the hedonic method would artificially increase the rental yield of minority homeowners relative to White homeowners. As a result, the return on housing investment for minority homeowners would be artificially increased relative to that of White homeowners, which we show empirically in appendix G.

5 Evidence on rent-to-price ratios by race from other sources

This section compares the rent-to-price ratios implied by our new rent imputation method to evidence from other data sources. We first examine raw data from the Consumer Expenditure Survey(CEX), which asks homeowners both about what they believe they could sell their homes for as well as what they think they could rent their homes for. As far as we are aware, this is the only dataset that directly measures rents and sale prices for owner-occupied properties in the US, though it relies heavily on homeowners' subjective beliefs. Second, we use large administrative datasets on housing transactions from Corelogic and on rents from Multiple Listing Services to estimate rent-to-price ratios using a method similar to our approach in the AHS.

5.1 Consumer Expenditure Survey

Figure 5.1 presents a time series plot of the median rent-to-price ratio in the CEX from 1994-2022 by racial group. This figure looks qualitatively quite similar to our imputed AHS rents in Figure 3.2 with a flat period in the 1990s, declines in the 2000s, growth in the early 2010s and then declines to 2022. The CEX data also shows Black rent-to-price ratios above Hispanic and Hispanic rent-to-price ratios above White. However, the magnitudes are a bit higher than what we find in the AHS. Table 5.1 reports our AHS imputed rent-to-price ratios by race for the same sample period as our CEX data. In the CEX, the average rent-to-price ratio for White homeowners is 9.40%, while we find 7.96% in the AHS. For Black homeowners the CEX shows 11.51%, while the AHS shows 8.93%. For Hispanic homeowners, the rent-to-price ratio is 11.21% in the CEX and 8.47% in the AHS. The racial gaps in rent-to-price ratios are somewhat wider in the CEX than what we find in the AHS. While these results broadly confirm the rental yield disparities we found in the AHS, the quantitative difference between the findings in the two data sources raises the question of which is most correct.

One possibility is that most homeowners who do not rent their homes out are overly optimistic about the rent they could charge. This bias is likely worse for rent than for price data, biasing rental yields upwards. While there is always some concern with bias and measurement error in survey-reported variables, owner occupants did purchase their house, making them at least somewhat informed on property values. However, most owners will never rent out their house, possibly making them very uninformed about the rental value of their property. Alternatively, it could be that the raw data in the CEX is more accurate since our rent imputation approach in the AHS requires us to make statistical assumptions which could be violated. While our estimation procedure allows for unobservable differences in quality between the switchers we use for estimation and other owner-occupied homes, we have to assume that these differences are captured by a home's price. To adjudicate this discrepancy with a third data source, we now turn to administrative data on property transaction prices, both in the rental market and the purchase market, where we believe that rents and prices are most accurately measured.

Figure 5.1: Median Rent-to-Price Ratios for Homeowners by Racial Group over Time, 1994-2022 in the Consumer Expenditure Survey

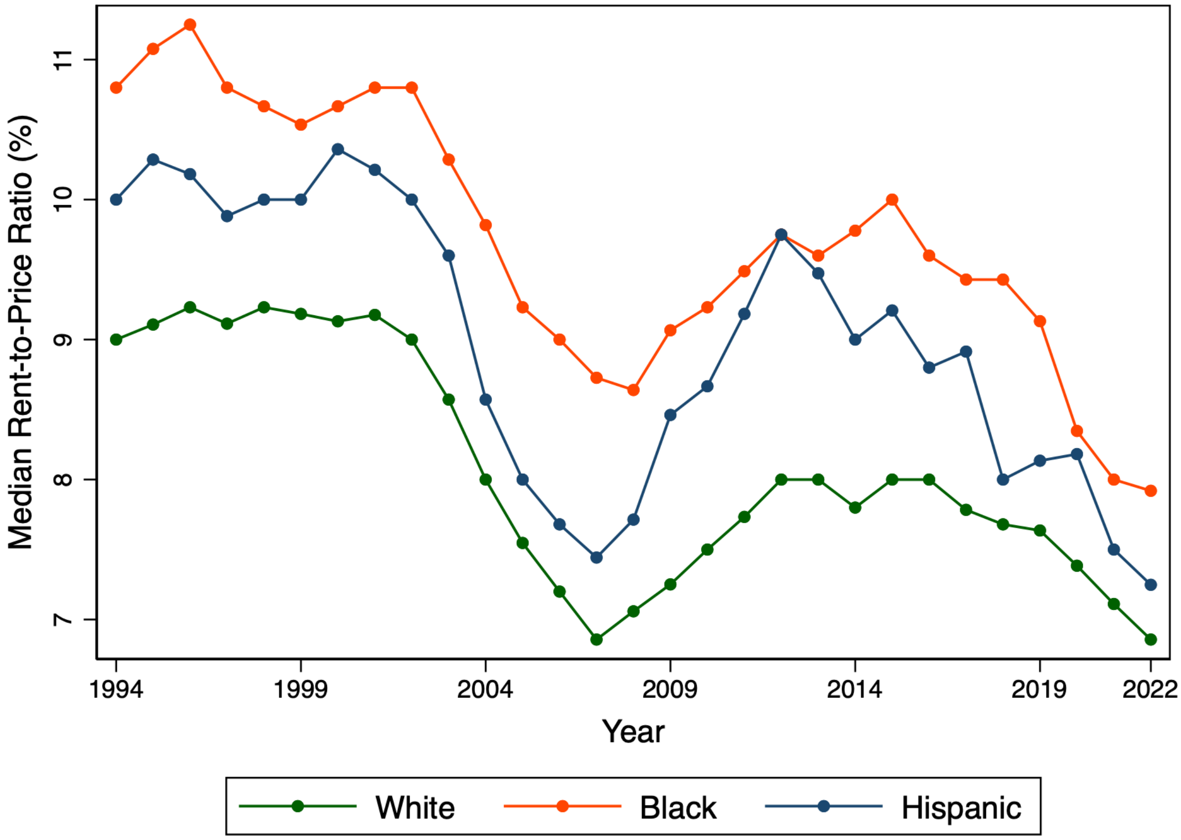

Note: Data from the Consumer Expenditure Survey. The CEX asks homeowners the property value and what their property could rent for. We use these to calculate rent-to-price ratios.

Table 5.1: Rent-to-Price Ratio for Owner-occupied units in the CEX and AHS by Owner Race, 1994-2021

CEX	Mean	Median	SD	25th Pct.	75th Pct.	Num Obs.
White	9.40	7.60	7.58	5.76	10.06	127,222
Black	11.51	9.29	8.47	6.96	12.51	12,962
Hispanic	11.21	8.47	9.65	6.17	12.00	14,757
AHS	Mean	Median	SD	25th Pct.	75th Pct.	Num Obs.
White	7.96	7.20	4.49	5.43	9.00	335,161
Black	8.93	7.91	5.00	5.77	10.44	37,204
Hispanic	8.47	7.34	5.22	5.14	9.85	39,126

Note: The sample includes all homeowners and is winsorized at the 2.5th and 97.5th percentiles.

5.2 Rental yields from administrative data

To complement our evidence on racial disparities in rental yields from the AHS and the CEX, this section reports evidence from rental yields inferred from data on actual transaction prices. We turn to administrative deeds records on property transactions, assessor records for property characteristics, and MLS rental listings from Corelogic. We also use data on property resident characteristics from Infutor. The deeds records on property transactions prices are linked to assessor records and MLS rental listing through unique parcel id numbers (APNs).[4] These data are then linked to the property resident listed in the Infutor data through a link using address strings and spatial linking through GPS coordinates. See Appendix B for more details.

A key upside of using these administrative data sources is the very large sample size. This linking procedure provides us with 2,845,590 observations that have a reported rent and a property transaction which we use in estimation. However, our data only start in 2007 and are not geographically representative, as we show in appendix B. Over 50 percent of our sample comes from 5 MSAs: Las Vegas, NV; Los Angeles, CA; Phoenix, AZ; Washington DC, and West Palm Beach, FL.

Since our sample is so much larger than the AHS switcher sample, we run our rent estimation regression separately by metropolitan statistical area and by year. For our baseline specification, we run the following regression of housing rents on housing prices and other covariates

$$rent_{ht} = \alpha_{jt} + \beta_{jt} X_{ht} + \epsilon_{ht}. \tag{12}$$

The regression includes the house's transaction price and a rich set of controls and fixed effects described in appendix B. We run this regression separately by MSA-year if we have at least a sample size of 1,000 observations. If this threshold fails, we run the same regression pooled across years with year fixed effects if there is a minimum of 2,000 observations. Regions that do not have enough data are excluded. One issue we must account for that was absent in the AHS is that the year at which the housing unit is rented need not be the same year in which the housing unit transacts in the purchase market. To deal with this, we compute the quantile of each housing unit's price in a county-year-specific distribution of house values from the American Community Survey. We then assign the house an imputed purchase price with the same quantile as the house value distribution which matches the county and year of the rental transaction. We use this imputed price as a regressor to help predict rents.

We use our estimates of equation 12 to construct fitted values for the rental value of owner-occupied housing units. We then compute the rental yield of an owner-occupied home by dividing our predicted rent value by the house's resale value $price_{ht}$. In years where the house does not transact, we use the approach described above using price quantiles to impute a price. The final sample we analyze covers 192 million observations of property values and imputed rents for owner-occupied properties between 2007 and 2016.

Table 5.2 presents regressions that analyze racial disparities in the log rent-to-price ratios we esti-

[4] We only study MLS rental listings that indicate that the rental listing was actually rented and lists the "closing" rent, which may differ from the original post rent in the listing.

mate.[5] Column 1 controls for year fixed effects and finds Black homeowners earn a 1.6 percentage point higher rental yield than White homeowners. Hispanic homeowners earn 0.39 percentage points more than White homeowners. To handle our non-representative geographic sample, column 2 adds MSA fixed effects and column 3 adds MSA-year fixed effects, which look similar. When looking within MSAs, Hispanics now earn 1.02 percentage points higher rent-to-price ratios than Whites. This substantial widening comes from the fact that our sample over-weights California, where rent-to-price ratios are low and many Hispanics live.

To validate our AHS rent imputation method against this administrative data, columns 5 and 6 present regressions in our AHS sample restricted to the same time period and MSAs as covered in our administrative data. Just including year fixed effects, we find Black homeowners earn 1.49 percentage points higher rent-to-price ratios than White homeowners and Hispanic homeowners earn 0.19 percentage points more than White homeowners. This is extremely economically close to the Black-White gap we find of 1.61pp and Hispanic-White gap of 0.39pp in the administrative and easily within our 95% confidence intervals. Column 6 adds MSA-year fixed effects to our AHS sample and we find a within-MSA Black-White gap of 1.47pp and Hispanic-White gap of 0.92pp. These are again very economically and statistically close to our administrative data findings of 1.67pp and 1.02pp, respectively. This suggests that the racial disparities we found in our AHS data are the right size, while those in the CEX survey data are potentially too large.

Finally, we add zip-year fixed effects into our administrative data regression in column 4 of Table 5.2. The within zip-year Black-White gap falls to 0.48pp, and the Hispanic-White gap falls to 0.35pp. This implies that 29% (0.48/1.67) of the within-MSA Black-White and 34% (0.35/1.02) of the Hispanic-White rental yield gap occurs *within* zipcodes. This suggests that imputing rent-to-price ratios at the zipcode level using standard hedonic methods, as done in Kermani and Wong [2022] and Demers and Eisfeldt [2001], misses a substantial portion of racial inequality in rent-to-price ratios that happens within zipcodes.

Table 5.2: Racial Gaps in Log(1 + Rent-to-Price Ratio)

	CoreLogic				AHS	
	(1)	(2)	(3)	(4)	(5)	(6)
Black	0.01607***	0.01664***	0.01671***	0.00480***	0.0149***	0.0147***
	(0.00037)	(0.00031)	(0.00029)	(0.00015)	(0.00121)	(0.00107)
Hispanic	0.00392***	0.01025***	0.01026***	0.00346***	0.00190	0.00925***
	(0.00030)	(0.00025)	(0.00024)	(0.00007)	(0.00138)	(0.00129)
Observations	192,329,864	192,329,864	192,329,864	192,329,864	26,857	26,857
Year FEs	Yes	Yes	-	-	Yes	-
MSA FEs	No	Yes	-	-	No	-
MSA × Year FEs	No	No	Yes	-	No	Yes
Zip Code FEs	No	No	No	-	No	No
Zip Code × Year FEs	No	No	No	Yes	No	No

Note: This table documents racial gaps in log rent-to-price ratios in MSAs where we have observations both in Corelogic and the AHS from 2007-2016. For consistency with the AHS, Corelogic data are weighted per household rather than per person. Standard errors (clustered at the Zip Code level for Corelogic data and the unit level for AHS data) are in parentheses.

[5] We use log(1+rent/price) as our rent-to-price metric as it will be identical to the metric used in the next section where we decompose total rates of return into their individual sources. This metric will measure the contribution of rents to total rates of return.

6 Racial Disparities in the Return on Owner-Occupied Housing in the AHS

We now apply our quantile-based rent imputation to compare the total rates of return on owner-occupied housing earned by different racial groups. Our analysis so far has shown that our rent imputation method provides accurate rent predictions both for low- and high- price homes and that rental yields vary significantly between racial groups. This section now shows that the higher rent-to-price ratios for minority homeowners are reflected in a higher average return on housing investment as well.

For each homeowner, we calculate the log-return on its housing to account for their home's rental value, property taxes, and maintenance. First, we compute the total "dividend" of house i at time t as

$$dividend_{it} = rent_{it} - proptax_{it} - maint_{it}, \tag{13}$$

where the rent is imputed using our quantile-based methodology and property taxes and maintenance costs are directly reported in the AHS. Our maintenance measure includes both routine maintenance and occasional home improvement projects. For years when the AHS is reported annually, we measure log-returns as

$$log(P_{it+1}) - log(P_{it} - dividend_{it}). \tag{14}$$

After 1981, when the AHS becomes only bi-annual, we annualize our rate our rate of return measure by dividing the expression in equation 14 by 2. This yields a house-time level panel of rates of returns. We CPI-adjust all dollars, so rates of return are real.

Table 6.1 presents summary statistics on our return measure by race. The average return on housing is 6.5% for White homeowners, 7.2% for Black homeowners, and 7.9% for Hispanic homeowners. We find that the standard deviation of returns is highest for Black and then for Hispanic homeowners. Figure 6.1 plots these rates of return by race over time. We find the rates of returns of all three racial groups track macroeconomic housing conditions, with lower return rates during recessions. This figure suggests that the higher returns earned by Black and Hispanic homeowners come in non-recession times, while it appears these groups frequently earn lower returns during recessions than White homeowners. This figure highlights the importance of studying these returns over a long time period, since shorter run business cycle dynamics impact the overall level of returns, as well as the racial disparities.

We decompose racial rate of return differences into property appreciation rates and dividend yields in tables 6.2 and 6.3. To measure the contribution of dividends to the rate of return, we set the property appreciation rate to zero and recalculate the rate of return. Similarly, we set the dividend to zero to measure the contribution of appreciation rates. We find slightly higher appreciation rates for White (0.9%) than Black (0.7%) homeowners but a substantially higher appreciation rate of 1.6% for Hispanic homeowners. In addition, we find that Hispanic and particularly Black homeowners have higher average dividend yields (5.9% and 6.0%, respectively) than White homeowners (5.3%). The high 29% standard deviation of total returns for Black homeowners is explained primarily by their property appreciation rate standard deviation being 25.9%, the highest of all racial groups. This is broadly consistent with the Kermani and Wong [2022] finding that a disproportionate share of distressed sales for minority homeowners raises their probability of a large negative investment return.

Table 6.1: Log rate of return

	White	Black	Hispanic	All groups
Mean	0.065	0.072	0.079	0.066
Median	0.052	0.058	0.068	0.053
StdDev	0.245	0.291	0.270	0.251
Observations	626,636	59,142	44,810	758,828

Note: Data from 1974-2021 AHS home owners. Summary statistics of annual log rates of return by race.

Table 6.2: Log property value rate (excl. dividends)

	White	Black	Hispanic	All groups
Mean	0.009	0.007	0.016	0.009
Median	0.001	0.001	0.014	0.002
StdDev	0.220	0.259	0.237	0.225
Observations	626,636	59,142	44,810	758,828

Note: Data from 1974-2021 AHS home owners. Summary statistics of annual log property value appreciation by race.

Table 6.3: Log dividend rate

	White	Black	Hispanic	All groups
Mean	0.053	0.060	0.059	0.053
Median	0.045	0.052	0.049	0.045
StdDev	0.065	0.075	0.076	0.067
Observations	626,636	59,142	44,810	758,828

Note: Data from 1974-2021 AHS home owners. Summary statistics of annual log dividend rate by race.

Figure 6.1: Average Annual Rate of Return by Racial Group, 1974-2021

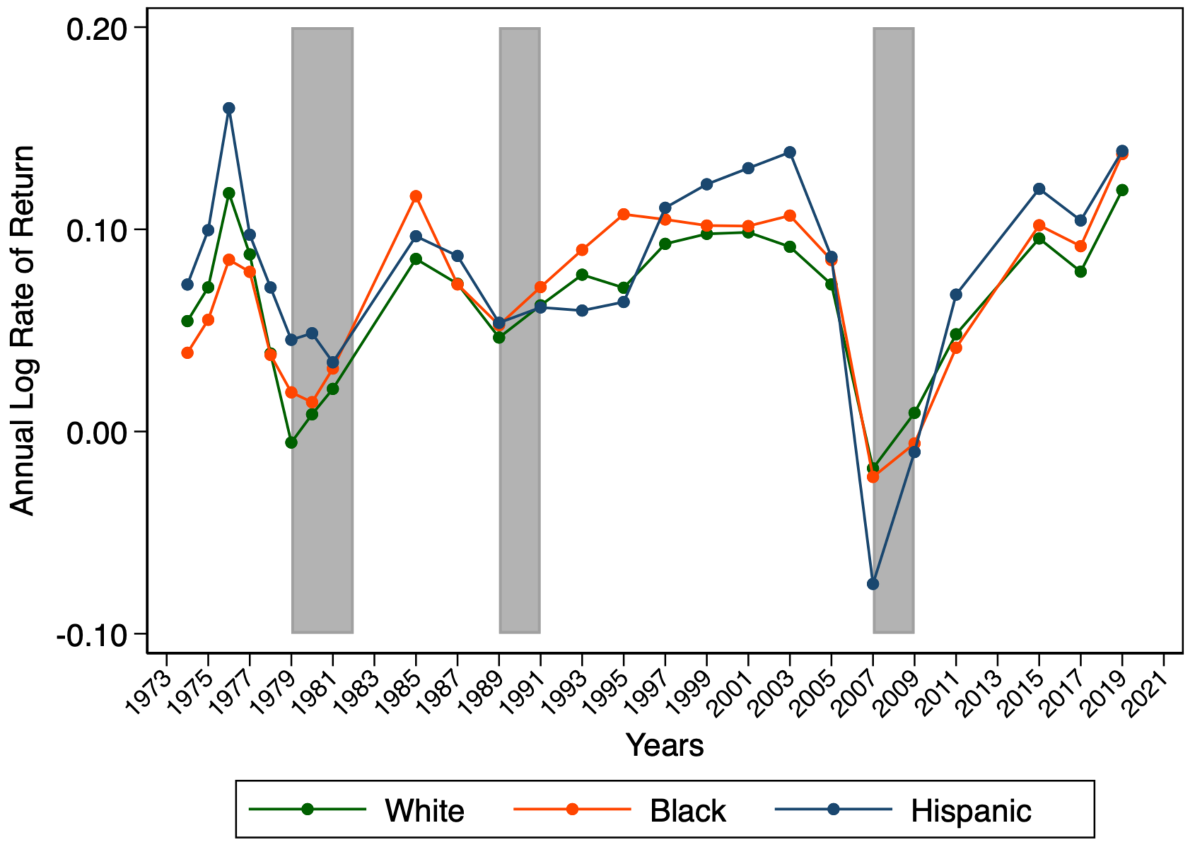

Note: Data from 1974-2021 AHS home owners. Plots average annual log rates of return by race. Shaded gray areas indicate periods of recession.

Table 6.4 investigates the extent to which we can account for disparities in housing rates of returns between racial groups with observable characteristics. To do so, we regress our individual-level log rates of return on a dummy variable for each racial group together with a growing list of control variables. Controlling only for year fixed effects indicates a 0.62pp higher rate of return for Black homeowners and 1.19pp higher rate of return for Hispanic homeowners relative to White homeowners. Column 2 adds MSA-year fixed effects. This does little to the Black-White gap, but it narrows the Hispanic-White gap a bit to 0.94 pp. Column 3 controls for household demographics: household head age, number of children in the household. These controls slightly increase the Black-White gap to 0.68pp and the Hispanic-White gap to 0.99pp. Controlling for household economic characteristics (household income, household head education, and years since household head moved into the housing unit) more than entirely accounts for the entire Black-White and Hispanic-White difference, leading to -0.36pp and -0.33pp respectively. This suggests that houses bought by lower income/education homeowners tend to have higher average returns and that these houses are disproportionately bought by minority homeowners. Adding additional controls for housing characteristics fully closes the Black-White and Hispanic-White gaps, leading to statistically and economically insignificant gaps of -0.16pp (Black-White) and -0.009pp (Hispanic-White).

We use a Gelbach [2016] decomposition in figure 6.2 to quantify our ability to explain differences in

Figure 6.2: Gelbach [2016] Decomposition of Explainable Share of Racial Disparities in Returns

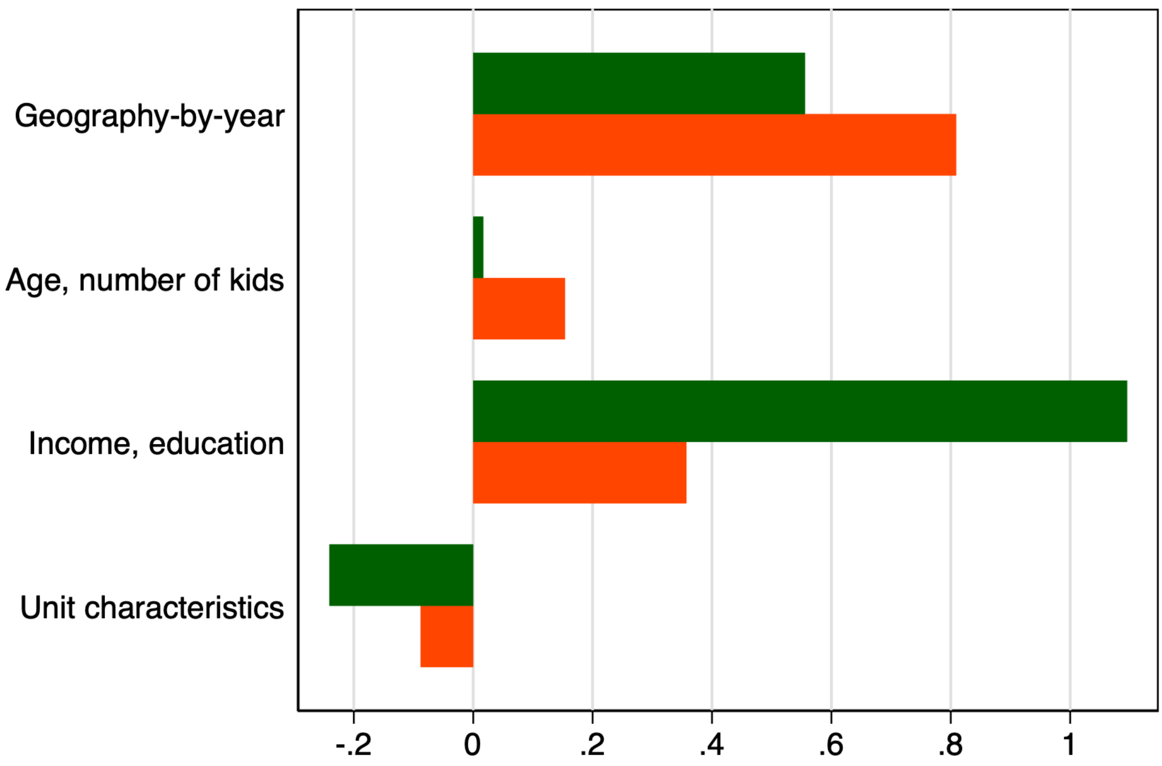

Note: Data from 1974-2021 AHS home owners. Each bar reports the fraction of the total return disparity explained by each set of control variables using the Gelbach [2016] method. The green bar represents the Black-White total rate of return gap, while the orange bar represents the Hispanic-White total rate of return gap.

average returns across homeowners of different races. This decomposition compares a regression with all covariates in column 5 of table 6.4 to the one in the first column. The decomposition uses a multivariate version of the formula for omitted variable bias to account for how much each variable helps to explain the change in our regression coefficient when all are included jointly. Unlike the standard approach of adding regressors one-by-one, this decomposition does not require us to arbitrarily choose the order in which regressors are added. We find consistent with table 6.4 that the income/education controls are by far the most important for explaining Black-White differences in rates of return and play an important role for the Hispanic-White gap. However, geography also plays a large role, especially for the Hispanic-White gap.

We examine the observable contribution of rental yield differences to total rates of return across racial groups in table 6.5. To do this, we examine what returns homeowners would earn in equation 14 if all terms other than rental yields were set to zero. Similar to our overall summary statistics, we find log rates of return due only to rental yields that are 1.2pp higher for Black homeowners and 0.79pp higher for Hispanic homeowners relative to White homeowners, controlling for year fixed effects. Few of our controls help to explain this disparity except the household economic characteristics added in column 4. This control explains 70% of the Black-White rental yield gap and essentially all of the Hispanic-White gap. This suggests that lower income/education homeowners live in cheaper homes, which tend to have higher rental yields. Adding in property characteristic controls widens the gap, leaving a Black-White gap of 0.53pp and 0.30pp Hispanic-White gap, even with all controls.

Potential differences in property taxes also contribute to dividend rates. Table 6.6 compares racial differences in log rates of return driven only by property taxes. Note that since taxes are a cost, all rates of return here will be negative and that a negative racial minority gap implies higher racial minority property tax rates. Controlling only for year fixed effects, column 1 shows property taxes contribute to a small, positive Black-White return gap of 0.029%, while the Hispanic-White gap is even closer to zero at -0.004%. This implies Black homeowners face slightly lower property tax rates than White homeowners, while Hispanics face a tiny bit higher. However, both are economically very small contributors to rates of return. Adding MSA-year fixed effects now shows slightly negative (i.e. higher tax rates) rates of return for both Black-White and Hispanic-White gaps. Adding in additional controls for demographics, household economic conditions, and property characteristics does little to mitigate the Hispanic-White gap, but fully explains the Black-White gap in returns driven by property taxes. Overall, since property taxes are small in magnitude relative to other forces driving rates of return, racial gaps in property taxes have negligible effects on racial differences in total rates of return.

This may seem at odds with recent work on racial inequality in property taxes [Avenancio-León and Howard, 2022] that shows Black and Hispanic homeowners paying higher property tax rates than White homeowners. Column 2 of Table 6.6, which includes only MSAs-year fixed effects, is the most comparable specification to analysis done in [Avenancio-León and Howard, 2022]. This specification shows both Black and Hispanic homeowners paying higher tax rates within MSA-years compared to Whites, although our magnitudes are smaller than those in Avenancio-León and Howard [2022]. This smaller magnitude appears to be driven by differences in time period and geographic span of the data in Avenancio-León and Howard [2022] and AHS. In Appendix C, we replicate the main findings of Avenancio-León and Howard [2022] in administrative Corelogic property tax data from 2007 to 2016. We then show how changes in sample definitions to align the time period and geography examined in the AHS and in Corelogic yield similar results across both datasets. Using only the years 2007-2016 and MSAs that are observed in the administrative property tax data, we find in the AHS that property taxes lower Black homeowner returns by 0.17 percent relative to White homeowner returns. This is substantially more negative than what we find for the overall nation from 1974-2021. This quantitatively small impact on housing returns is consistent with the magnitudes in Avenancio-León and Howard [2022]. A back of the envelope calculation using numbers from Avenancio-León and Howard [2022] indicates the median home worth $207,000 is on average taxed at 1.4 percent. Assuming their headline number that Black homeowners pay 12% higher tax rates within an MSA, this tax rate would be $1.4*(1.12) = 1.568$ percent. This would imply a Black-White rate-of-return gap of 0.19%, very close to what we find.

Looking at differences in rates of return driven by spending on housing maintenance and home improvement, Table 6.7 shows that Black homeowners spend a greater fraction of their home's value on maintenance and home improvement than White homeowners, leading to a 0.38% lower rate of return. Hispanic homeowners spend more than White homeowners as well, lowering their rate of return by 0.07% relative to White homeowners. This offsets some but not all of the higher rental yields Black and Hispanic homeowners earn. The Black-White and Hispanic-White gaps become more similar to each other, -0.28% and -0.24%, respectively, when adding MSA-year controls. Controlling for demographics has little impact on these numbers. Adding in household economic characteristics shrinks these gaps to -0.21% for Black homeowners and -0.11% for Hispanic homeowners. Adding in housing characteristics further narrows these gaps to -0.19% and -0.10% for Black and Hispanic homeowners, respectively, relative to White homeowners.

Next, we consider in Table 6.8 the determinants of racial disparities in property appreciation rates. Including only year fixed effects, we find a Black-White appreciation gap of -0.26% and a Hispanic-White gap of 0.48%. The fact that Hispanic homeowners earn higher appreciation rates is driven by their differences in geographic exposure to local appreciation rates. Column 2 adds MSA-year fixed effects, which leads to a Black-White gap of -0.18% and a Hispanic-White gap of 0.17%. Adding demographic controls does little to these gaps. Adding controls for household economic characteristics makes these gaps more negative. The Black-White gap is now -0.46% and Hispanic-White is -0.06%. Adding house characteristics makes these a bit more negative at -0.47% and -0.10%, respectively.

While we are able to replicate the previous literature's finding that Black and Hispanic homeowners earn lower appreciation rates than White homeowners, our magnitudes are less negative than those found in Kermani and Wong [2022] using administrative data. We return to this point and reconcile the returns we find in the AHS and in Corelogic administrative data with Kermani and Wong [2022] in Section 7.

Table 6.4: Log rate of return

	(1)	(2)	(3)	(4)	(5)
Black	0.00619***	0.00604***	0.00682***	-0.00355**	-0.00165
	(0.00116)	(0.00118)	(0.00118)	(0.00118)	(0.00118)
Hispanic	0.0119***	0.00940***	0.00994***	-0.00329*	-0.000942
	(0.00137)	(0.00137)	(0.00138)	(0.00135)	(0.00132)
Observations	742,745	742,745	742,745	742,745	742,745
Dummy for other racial group	Yes	Yes	Yes	Yes	Yes
Year FE (not interacted)	Yes	-	-	-	-
Geography-by-year FE	No	Yes	Yes	Yes	Yes
Demographics (age, #kids)	No	No	Yes	Yes	Yes
Income, education, tenure	No	No	No	Yes	Yes
Housing characteristics	No	No	No	No	Yes

Note: Data from 1974-2021 owner-occupied properties in the AHS. Standard errors are clustered at the unit level. The excluded race group is White owner-occupants. Geography FEs include the most granular-defined geography variables in the AHS for each year, which change some over time.

Table 6.5: Log rate of owner-occupied rent

	(1)	(2)	(3)	(4)	(5)
Black	0.0124***	0.0109***	0.0114***	0.00372***	0.00527***
	(0.000747)	(0.000699)	(0.000697)	(0.000684)	(0.000672)
Hispanic	0.00794***	0.0118***	0.0118***	0.000391	0.00297***
	(0.000962)	(0.000897)	(0.000899)	(0.000852)	(0.000760)
Observations	742,745	742,745	742,745	742,745	742,745
Dummy for other racial group	Yes	Yes	Yes	Yes	Yes
Year FE (not interacted)	Yes	-	-	-	-
Geography-by-year FE	No	Yes	Yes	Yes	Yes
Demographics (age, #kids)	No	No	Yes	Yes	Yes
Income, education, tenure	No	No	No	Yes	Yes
Housing characteristics	No	No	No	No	Yes

Note: Data from 1974-2021 owner-occupied properties in the AHS. Standard errors are clustered at the unit level. The excluded race group is White owner-occupants. Geography FEs include the most granular-defined geography variables in the AHS for each year, which change some over time.

Table 6.6: Log rate of annual property taxes

	(1)	(2)	(3)	(4)	(5)
Black	0.000287*	-0.000251*	-0.000260*	0.000121	-0.0000404
	(0.000143)	(0.000126)	(0.000126)	(0.000127)	(0.000127)
Hispanic	-0.0000432	-0.00153***	-0.00149***	-0.000963***	-0.00119***
	(0.000177)	(0.000170)	(0.000170)	(0.000172)	(0.000168)
Observations	742,745	742,745	742,745	742,745	742,745
Dummy for other racial group	Yes	Yes	Yes	Yes	Yes
Year FE (not interacted)	Yes	-	-	-	-
Geography-by-year FE	No	Yes	Yes	Yes	Yes
Demographics (age, #kids)	No	No	Yes	Yes	Yes
Income, education, tenure	No	No	No	Yes	Yes
Housing characteristics	No	No	No	No	Yes

Note: Data from 1974-2021 owner-occupied properties in the AHS. The dependent variable measures the rate-of-return due only to property tax payments. A negative racial gap indicates higher property tax rates for the racial minority group. Standard errors are clustered at the unit level. The excluded race group is White owner-occupants. Geography FEs include the most granular-defined geography variables in the AHS for each year, which change some over time.

Table 6.7: Log rate of maintenance cost

	(1)	(2)	(3)	(4)	(5)
Black	-0.00375***	-0.00278***	-0.00263***	-0.00214***	-0.00189***
	(0.000290)	(0.000259)	(0.000260)	(0.000260)	(0.000252)
Hispanic	-0.000731**	-0.00244***	-0.00198***	-0.00110***	-0.000994***
	(0.000261)	(0.000273)	(0.000275)	(0.000279)	(0.000267)
Observations	742,745	742,745	742,745	742,745	742,745
Dummy for other racial group	Yes	Yes	Yes	Yes	Yes
Year FE (not interacted)	Yes	-	-	-	-
Geography-by-year FE	No	Yes	Yes	Yes	Yes
Demographics (age, #kids)	No	No	Yes	Yes	Yes
Income, education, tenure	No	No	No	Yes	Yes
Housing characteristics	No	No	No	No	Yes

Note: Data from 1974-2021 owner-occupied properties in the AHS. The dependent variable is the rate of return driven only by maintenance spending. A negative racial gap indicates higher spending on maintenance and home improvement as a share of the property value by the racial minority group. Standard errors are clustered at the unit level. The excluded race group is White owner-occupants. Geography FEs include the most granular-defined geography variables in the AHS for each year, which change some over time.

Table 6.8: Log rate of property values (excluding dividends)

	(1)	(2)	(3)	(4)	(5)
Black	-0.00261**	-0.00176*	-0.00169*	-0.00461***	-0.00467***
	(0.000813)	(0.000841)	(0.000844)	(0.000853)	(0.000854)
Hispanic	0.00478***	0.00168	0.00178	-0.000560	-0.00110
	(0.000944)	(0.000978)	(0.000983)	(0.000996)	(0.000996)
Observations	742,745	742,745	742,745	742,745	742,745
Dummy for other racial group	Yes	Yes	Yes	Yes	Yes
Year FE (not interacted)	Yes	-	-	-	-
Geography-by-year FE	No	Yes	Yes	Yes	Yes
Demographics (age, #kids)	No	No	Yes	Yes	Yes
Income, education, tenure	No	No	No	Yes	Yes
Housing characteristics	No	No	No	No	Yes

Note: Data from 1974-2021 owner-occupied properties in the AHS. The Dependent variable is the rate of return driven only by price appreciation. Standard errors are clustered at the unit level. The excluded race group is White owner-occupants. Geography FEs include the most granular-defined geography variables in the AHS for each year, which change some over time.

6.1 Exposure to Business Cycles

We explore racial disparities in the total rate of return on housing over the business cycle in Table 6.9. In this table, we regress the realized log total return homeowners' earn on dummy variables for race as well as interactions between race and an indicator for recession years with a growing list of control variables. In all of our specifications, the interaction effect between a recession indicator and an indicator for either a Black or Hispanic homeowner is negative. This implies that Black and Hispanic homeowners have returns on housing that are more sensitive to the business cycle than White homeowners. In column 1, the Black-White return gap is 0.82% in non-recession times, but it falls by 0.67 percentage points during recessions. The Hispanic-White gap is very sensitive to the business cycle. In non-recession times the gap is 2.05%, but it falls by 3.32 percentage points during recessions and becomes negative. A substantial share of Hispanic homeowners' excess exposure to boom-bust cycles appears to be due to the MSAs that they live in. Controlling for MSA-year effects in column 2 reduces the Hispanic-White gap to 1.14% in non-recessions, which now falls by 0.80 percentage points in recessions. MSA-year fixed effects have little effect on the Black-White gap, and if any widen the gap a bit, both during booms and busts. Adding controls for household demographics does little to these gaps for either racial group.

However, column 4 shows that controlling for household economic characteristics largely accounts for the greater return cyclicality of Hispanic homeowners and explains the excess returns of the Black-White gap during boom times. However, we find that these observables have almost no ability to account for the excessively negative returns of Black homeowners during recessions. Column 5 adds house characteristics controls. The Hispanic-White gap was already statistically indistinguishable from zero without these controls in both boom and bust times, and this specification pushes it even closer to zero. However, the Black-White return gap remains negative and significant in recession times at -0.68%. One possibility is that segregation between Black and White homeowners results in Black homeowners living in more cyclical areas. Indeed, geographic racial segregation indices show that Blacks are twice as geographically segregated from Whites as Hispanics are [Elbers, 2021]. Another possibility is Black homeowners are more likely to sell their homes as part of a foreclosure [Kermani and Wong, 2022], particularly in recessions.

We decompose this excess cyclicality into price appreciation and rental yields in tables 6.10 and 6.11. For both Black and Hispanic homeowners, we find that their rental yields only fall a bit in recessions relative to White homeowners. Even with all of our observable controls the Black-White and Hispanic-White return gap from rents remains positive and significant in boom and bust times at 0.56% and 0.31%, respectively during boom times and 0.44% and 0.25% during bust times. Turning to price appreciation rates, table 6.11 shows Blacks have slightly lower price appreciation in boom times at -0.14% and this falls to -0.55% in bust times. This excess cyclicality is even larger for the Hispanic-White gap, with a positive gap of 1.25% in boom times and -1.74% during bust times. Controlling for geography-year fixed effects narrows these gaps for Hispanics to 0.29% during boom times and -0.18% during bust times. However, this does little to explain the excess exposure of Black homeowners to the business cycle.

Table 6.9: Log rate of return

	(1)	(2)	(3)	(4)	(5)
Black	0.00815***	0.00856***	0.00924***	-0.00143	0.000310
	(0.00147)	(0.00150)	(0.00150)	(0.00149)	(0.00147)
Black=1 × Recession=1	-0.00667*	-0.00884**	-0.00853**	-0.00731*	-0.00677*
	(0.00284)	(0.00289)	(0.00290)	(0.00291)	(0.00292)
Hispanic	0.0205***	0.0114***	0.0115***	-0.00263	-0.0000337
	(0.00159)	(0.00166)	(0.00166)	(0.00163)	(0.00161)
Hispanic=1 × Recession=1	-0.0332***	-0.00799*	-0.00643	-0.00318	-0.00396
	(0.00342)	(0.00346)	(0.00346)	(0.00348)	(0.00343)
Observations	742,745	742,745	742,745	742,745	742,745
Recession-by-race FE (fully interacted)					
Dummy for other racial group	Yes	Yes	Yes	Yes	Yes
Year FE (not interacted)	Yes	-	-	-	-
Geography-by-year FE	No	Yes	Yes	Yes	Yes
Demographics (age, #kids)	No	No	Yes	Yes	Yes
Income, education, tenure	No	No	No	Yes	Yes
Housing characteristics	No	No	No	No	Yes

Note: Data from 1974-2021 owner-occupied properties in the AHS. The dependent variable is the total log rate of return. Standard errors are clustered at the unit level. The excluded race group is White owner-occupants. Geography FEs include the most granular-defined geography variables in the AHS for each year, which change some over time. Recession year dummies (1975,1979,1980,1981,1989,1991,2007,2009) come from FRED.

Table 6.10: Log rate of owner-occupied rent

	(1)	(2)	(3)	(4)	(5)
Black	0.0129***	0.0114***	0.0118***	0.00416***	0.00564***
	(0.000814)	(0.000775)	(0.000772)	(0.000759)	(0.000752)
Black=1 × Recession=1	-0.00153	-0.00164*	-0.00158	-0.00144	-0.00121
	(0.000814)	(0.000810)	(0.000810)	(0.000819)	(0.000817)
Hispanic	0.00903***	0.0124***	0.0121***	0.000223	0.00310***
	(0.00103)	(0.000977)	(0.000979)	(0.000929)	(0.000831)
Hispanic=1 × Recession=1	-0.00424***	-0.00202	-0.00152	0.000320	-0.000629
	(0.00112)	(0.00118)	(0.00118)	(0.00116)	(0.00104)
Observations	742,745	742,745	742,745	742,745	742,745
Recession-by-race FE (fully interacted)					
Dummy for other racial group	Yes	Yes	Yes	Yes	Yes
Year FE (not interacted)	Yes	-	-	-	-
Geography-by-year FE	No	Yes	Yes	Yes	Yes
Demographics (age, #kids)	No	No	Yes	Yes	Yes
Income, education, tenure	No	No	No	Yes	Yes
Housing characteristics	No	No	No	No	Yes

Note: Data from 1974-2021 owner-occupied properties in the AHS. Standard errors are clustered at the unit level. The excluded race group is White owner-occupants. Geography FEs include the most granular-defined geography variables in the AHS for each year, which change some over time. Recession year dummies (1975,1979,1980,1981,1989,1991,2007,2009) come from FRED.

Table 6.11: Log rate of property values (excluding dividends)

	(1)	(2)	(3)	(4)	(5)
Black	-0.00137	-0.0000327	-0.0000783	-0.00318**	-0.00323**
	(0.00112)	(0.00115)	(0.00115)	(0.00116)	(0.00116)
Black=1 × Recession=1	-0.00417	-0.00607*	-0.00568*	-0.00494	-0.00502
	(0.00252)	(0.00257)	(0.00258)	(0.00259)	(0.00259)
Hispanic	0.0125***	0.00327*	0.00292*	0.000332	-0.000209
	(0.00123)	(0.00129)	(0.00130)	(0.00131)	(0.00131)
Hispanic=1 × Recession=1	-0.0299***	-0.00622*	-0.00468	-0.00369	-0.00369
	(0.00295)	(0.00294)	(0.00295)	(0.00299)	(0.00298)
Observations	742,745	742,745	742,745	742,745	742,745
Recession-by-race FE (fully interacted)					
Dummy for other racial group	Yes	Yes	Yes	Yes	Yes
Year FE (not interacted)	Yes	-	-	-	-
Geography-by-year FE	No	Yes	Yes	Yes	Yes
Demographics (age, #kids)	No	No	Yes	Yes	Yes
Income, education, tenure	No	No	No	Yes	Yes
Housing characteristics	No	No	No	No	Yes

Note: Data from 1974-2021 owner-occupied properties in the AHS. The dependent variable is the rate of return driven only by price appreciation. Standard errors are clustered at the unit level. The excluded race group is White owner-occupants. Geography FEs include the most granular-defined geography variables in the AHS for each year, which change some over time. Recession year dummies (1975,1979,1980,1981,1989,1991,2007,2009) come from FRED.

7 Reconciling Estimates of House Price Appreciation Rates

This section attempts to reconcile the measured rate of house price appreciation observed in the AHS to estimates from administrative datasets such as Corelogic. As documented by Kermani and Wong [2022], if we measure the return on housing by looking only at housing spells with a measured buy and sell price, White homeowners earn higher realized returns than Black or Hispanic homeowners. They demonstrate that racial disparities in foreclosure rates drive much of this overall return difference, and we find similar results to them when examining housing spells with measured buy and sell prices. However, we also show that in our Corelogic data, a significant share of housing spells either begin at the start of the sample or end at the conclusion of the sample and are therefore missing a buy price and/or a sell price. Our analysis is restricted to four counties (Cook County, IL, Philadelphia County, PA, King County, WA, and Santa Clara County, CA), where we study foreclosure discounts in detail in appendix D. Racial disparities in house price appreciation look much closer to our AHS results once we account for the entire housing stock including houses that do not transact. In addition, under the very conservative assumption that our AHS returns entirely miss foreclosure discounts, we use a measure of foreclosure rates and discounts from Kermani and Wong [2022] and show that minorities still earn a higher rate of return after accounting for racial disparities in foreclosures.

We first show in table 7.1 that a significant share of housing spells either begin at the start of our sample or are still continuing at the end, particularly for minorities. If we weigh these spells proportional to their length, we find that 74.9% of the years of overall homeownership and 86.1% of the years of Black homeownership are part of a spell that begins and/or ends outside of our sample. To make the rates of return measured in AHS and Corelogic comparable, we need to account for the fact that the AHS tries to measure returns on the entire housing stock while Corelogic only reports data on housing units that transact.

Table 7.1: Summary Statistics: Completeness of Housing Spells by Race

	White		Black		Hispanic		All	
	Num.	Pct.	Num.	Pct.	Num.	Pct.	Num.	Pct.
Panel A: At the Spell Level								
Purchased and Sold Only	2,370,788	40.78	231,359	24.93	294,445	38.99	2,896,592	38.63
Purchased but Unsold Only	1,851,049	31.84	372,765	40.16	365,627	48.41	2,589,441	34.54
Not Purchased Only	1,592,150	27.38	324,090	34.92	95,177	12.60	2,011,417	26.83
Not Purchased or Unsold Only	3,443,199	59.22	696,855	75.07	460,804	61.01	4,600,858	61.37
Overall	5,813,987	-	928,214	-	755,249	-	7,497,450	-
Panel B: At the Spell-Year Level								
Purchased and Sold Only	29,684,541	25.08	3,301,150	13.90	3,522,036	23.88	36,507,727	23.27
Purchased but Unsold Only	38,725,323	32.72	8,390,092	35.32	8,004,645	54.28	55,120,060	35.14
Not Purchased Only	49,956,200	42.20	12,066,557	50.79	3,220,805	21.84	65,243,562	41.59
Not Purchased or Unsold Only	88,681,523	74.92	20,456,649	86.10	11,225,450	76.12	120,363,622	76.73
Overall	118,366,064	-	23,757,799	-	14,747,486	-	156,871,349	-

Note: This table presents summary statistics on the fraction of housing spells that begin with a purchase and/or terminate with a sale in sample, by race. Data is from 1975-2016 across four counties: Cook County, IL, Philadelphia County, PA, King County, WA, and Santa Clara County, CA.

Table 7.2 analyzes how racial disparities in house price appreciation vary with the sample used to compute average returns. In each column, we regress the log of house price appreciation on dummy variables for each White/Black/Hispanic racial group and change our included sample and/or fixed effect controls across columns. We use data from four counties: Cook County, IL, Philadelphia County, PA, King County, WA, and Santa Clara County, CA. In appendix D, we discuss how we have carefully measured the returns on foreclosed homes specifically in these counties. The two panels A and B of the table measure the return on foreclosed properties in different ways. Panel A uses the price at which houses sold in a foreclosure auction, while Panel B uses a measure of the value of debt which a homeowner expunges when it "sells" its home for the face value of its outstanding mortgage, as described in detail in appendix D. In the columns on the left of this table, we see that racial return disparities are sensitive to foreclosure measurement, while they become much less so in the columns on the right hand side of the table that use more and more expansive sets of homes.

In the table's first specification in panel A, we find a 5.3% lower rate of house price appreciation for Hispanic and 9.9% lower rate of house price appreciation for Black homeowners relative to White homeowners. This column looks only at complete housing spells with a buy and a sell, restricts to a 2000-2016 sample consistent with Kermani and Wong [2022], and only examines homes bought with a mortgage (since Kermani and Wong [2022] use HMDA mortgage data to infer the race of a homeonwer). We then include in column 2 the set of homeowners who buy houses in sample but do not sell. To measure the rate of appreciation on houses that do not have both a buy and a sell, we follow Kermani and Wong [2022] in using a year-county specific return on an FHFA local house price index as a proxy for the return on the individual house. We find a 3.0% Hispanic-White and 4.3% Black-White return gap in column 2, somewhat larger than the comparable 2.26% and 2.32% gaps Kermani and Wong report. This difference is likely due to the fact that our sample of 4 counties contains Chicago and Philadelphia, poor/segregated cities with larger than usual racial disparities. The disparities we measure shrink as we expand our sample to 1975-2016 matching the AHS, include houses purchased without a mortgage, and re-weight our average over property ownership spells proportional to the length of the spell. This re-weighting is consistent with how we measure returns in the AHS, since each house-year observation gets equal weight and long ownership spells would have several observations.

In the final column, after we have controlled for time/county fixed effects and included the entire housing stock, even those that never transact, we find much smaller racial disparities in property appreciation rates than in column 1. The Black-White gap is now only 50 basis points, while the Hispanic-White gap is only 18 basis points. These disparities are very close to the 47 basis point and 11 basis point return disparities we see in the last column of table 6.8 with the most exhaustive set of controls. This suggests that the racial disparities in house price appreciation rates we observe in the AHS are consistent with data in Corelogic and with the findings in Kermani and Wong [2022] once differences in sample selection are accounted for.

Table 7.2: Racial Gaps in Log House Price Appreciation Rates: Reconciling with Kermani and Wong [2022]

	Purchased during 2000-2016	Never Sold	Purchase Year × County FEs	Purchased during 1975-2016	Cash Purchased	Tenure Length Weighted	Purchase Year × Sale Year × County FEs	Not Purchased
Panel A: Auction Price-based Return for Foreclosed Spells								
Black	-0.09913***	-0.04313***	-0.03234***	-0.02845***	-0.02207***	-0.01348***	-0.00930***	-0.00505***
	(0.00154)	(0.00051)	(0.00050)	(0.00033)	(0.00028)	(0.00009)	(0.00008)	(0.00004)
Hispanic	-0.05285***	-0.03037***	-0.02019***	-0.01479***	-0.00965***	-0.00655***	-0.00291***	-0.00180***
	(0.00104)	(0.00042)	(0.00040)	(0.00031)	(0.00026)	(0.00008)	(0.00007)	(0.00005)
Panel B: Outstanding Mortgage-based Return for Foreclosed Spells								
Black	-0.06741***	-0.03322***	-0.02501***	-0.02382***	-0.01922***	-0.01303***	-0.00812***	-0.00441***
	(0.00123)	(0.00041)	(0.00041)	(0.00027)	(0.00025)	(0.00008)	(0.00007)	(0.00004)
Hispanic	-0.04349***	-0.02675***	-0.01689***	-0.01290***	-0.00839***	-0.00633***	-0.00255***	-0.00157***
	(0.00093)	(0.00038)	(0.00036)	(0.00029)	(0.00024)	(0.00008)	(0.00006)	(0.00005)
Observations	476,115	1,129,740	1,129,740	3,011,164	5,486,006	5,486,006	5,486,006	7,497,430

Note: This table documents racial gaps in log house price appreciation for our Corelogic housing spell sample from 1975-2016 in four counties: Cook County, IL, Philadelphia County, PA, King County, WA, and Santa Clara County, CA. Data are at the individual-spell level, with observations weighted by household each year. All numbers are adjusted for inflation. Standard errors (clustered at the household-spell level) are in parentheses.

As additional robustness, we explore the sensitivity of our average rate of return estimates in the AHS to potential mis-measurement of the impact of foreclosure sales. Kermani and Wong [2022] show that Black and Hispanic homeowners have a higher probability of foreclosure than White homeowners, especially during downturns. They provided us with a time series of estimates of the probability of foreclosure and size of the foreclosure discount for Black, White, and Hispanic homeowners, by year. Since our house price data is a survey-based measure of owners' perceived value of their house, it is possible that these values don't fully include the effect of foreclosure on home values. For example, if a house flipper buys a foreclosed home and then sells it to a new homeowner at a higher price between waves of our survey, we may miss the impact of this foreclosure. If, however, the new homeowner also buys at a distressed price, our data should likely reflect the impact of this foreclosure. Under the very conservative assumption that foreclosure discounts are entirely ignored in our estimates from the AHS, we show that our measure of the rates of return is still higher for Black and Hispanic homeowners than White homeowners with this discount added.

Table 7.3: AHS Log Rate of Return: 2001-2021

	Num. Obs.	Mean	Median	Std. Dev.
White	218,688	0.065	0.055	0.256
Black	24,792	0.071	0.061	0.287
Hispanic	27,490	0.078	0.071	0.276

Note: Data from 2001-2021 owner-occupied properties in the AHS. The sample period is restricted to 2001-2021 since we only have data from Kermani and Wong [2022] on foreclosure discounts in this period. The table reports our baseline log rate of return estimates in this restricted sample period, without the foreclosure adjustment from Kermani and Wong [2022].

Table 7.4: AHS Log Rate of Return with added foreclosure discount from Kermani and Wong [2022]: 2001-2021

	Num. Obs.	Mean	Median	Std. Dev.
White	218,688	0.064	0.053	0.256
Black	24,792	0.065	0.056	0.288
Hispanic	27,490	0.072	0.066	0.277

Note: Data from 2001-2021 owner-occupied properties in the AHS and augmented with data shared by Kermani and Wong [2022] on foreclosure rates and foreclosure discounts by race. These rates of returns make the conservative assumption that the AHS completely misses the effects of foreclosures on property values. See text for details of foreclosure adjustment.

We present our results in tables 7.3-7.4. The average return for White homeowners decreases by 0.1 pp, while Black and Hispanic homeowners respectively have a decrease of 0.6 pp and 0.8 pp in their average rates of return. These differences in returns are far smaller than the 1.37 pp, 3.30 pp, and 5.14 pp decreases in returns Kermani and Wong [2022] respectively report for White, Black and Hispanic homeowners when including foreclosure sales relative to a non-distressed only subsample. One key reason why our average rate of return estimates are less sensitive to foreclosure discounts than those in Kermani and Wong [2022] is that we average over homeowners at a given time, while they average over realized transactions. The average of the one-period returns in equation 14 is over the I_t houses at time t+1, indexed by i, is:

$$\frac{1}{I_t}\sum_i [log(P_{it+1}) - log(P_{it} - dividend_{it})]. \tag{15}$$

Suppose each house has a probability $pr_{for,t}$ of foreclosure and is then sold for a fraction f_t of its non-foreclosure value P^*_{it+1}. This average return can be written as

$$\frac{1}{I_t}\sum_i [log(P^*_{it+1}) - log(P_{it} - dividend_{it})] + pr_{for,t}log(1 - f_t), \tag{16}$$

so the reduction in the annual rate of return is $pr_{for,t}log(1 - f_t)$. We use the annual series of foreclosure rates $pr_{for,t}$[6] and discounts f_t provided by Kermani and Wong [2022] for each racial group. To get a sense of magnitudes, for an annual foreclosure rate of 2%, close to the peak after the 2008 financial crisis, and a large 30% foreclosure discount, this would equal $.02 * log(.7)$, a relatively modest 31 basis point decrease in the average rate of return.

In contrast, Kermani and Wong [2022] compute average returns by considering pairs of transactions where a house i is bought at a time $t_{1,i}$ for a price $P_{b,i,t_{1,i}}$ and then later sold at time t_2 for a price

[6] We sum the authors' foreclosure probabilities across the quarters in a year, which slightly overestimates the odds of foreclosure due to double counting of houses foreclosed twice.

$P_{s,i,t_{2,i}}$, with I total transactions. If we ignore dividend payments for this calculation, their average rate of return can be computed as

$$\frac{1}{I}\sum_{i}(\frac{P_{s,i,t_{2,i}}}{P_{b,i,t_{1,i}}})^{\frac{1}{t_{2,i}-t_{1,i}}}. \tag{17}$$

Unlike our estimate, which uses the entire housing stock at a given time including homes that do not transact, this summary statistic considers only realized transactions. Suppose that a randomly selected fraction of transactions with probability $pr(foreclose|transact)$ are foreclosures. The probability $pr(foreclose|transact)$ is higher than the foreclosure probability we use, since most houses do not transact in any given year and all foreclosures occur among houses that are sold. If the non-distressed sale price of a house would be $P^*_{s,i,t_{2,i}}$ and all foreclosures have the same discount f, then this sales-based average return equals

$$pr(foreclose|transact)\frac{1}{I}\sum_{i} f^{\frac{1}{t_{2,i}-t_{1,i}}}(\frac{P^*_{s,i,t_{2,i}}}{P_{b,i,t_{1,i}}})^{\frac{1}{t_{2,i}-t_{1,i}}}. \tag{18}$$

From this expression, we see that a foreclosure on a short tenure transaction has a large impact on this average rate of return statistic, since the discount f is taken to the power $\frac{1}{t_{2,i}-t_{1,i}}$, which can be large. Because only examining transactions increases the weight on foreclosures in this statistic, and because foreclosures on short holding periods can dramatically impact the implied rate of return, the return summary statistic in Kermani and Wong [2022] is more sensitive to foreclosure discounts than ours. While we believe that our notion of average return, which averages equally over all homeowners of a racial group in a given year without requiring them to transact, is an intuitive measure of expected return, there is no contradiction between our results and the greater impact of a foreclosure discount in Kermani and Wong [2022].

8 Conclusion

This paper quantifies racial differences in the total rate of return on owner-occupied housing. The total rate of return of buying a house equals the appreciation in its price plus the rental value of its housing services minus taxes and costs. Our results rely crucially on a new method we develop to impute the rental value of owner-occupied housing that performs well for both low-value and high-value homes. Unlike existing hedonic methods, our estimator accurately predicts rents of both low- and high-value homes by using price information to account for house quality. This accuracy across the value distribution is crucial for our application to racial inequality, since White, Black, and Hispanic homeowners vary in the average value of their homes. We document similar racial disparities in the rental yield of owner-occupied housing across multiple data sources, from raw survey data in the CEX to a more sophisticated imputation approach using actual transactions from Corelogic and MLS.

With our total return measure, we document a new stylized fact that Black and Hispanic homeowners earn higher average total returns on housing than White homeowners. This racial disparity in returns is primarily due to differences in rental yields across racial groups. Our new method of rent imputation that accurately matches rental yields across the house price distribution is crucial for measuring this precisely. This rental yield disparity is largely driven by a higher rent-to-price ratio for lower value homes, and controlling for homeowner income/education largely explains the return disparity it creates. In our nationally representative sample from 1974-2021, including the many homes that do not transact, differences in house price appreciation rates across racial groups are somewhat smaller. We also find that Black and Hispanic homeowners face riskier and more cyclical returns than White homeowners, and that the excess negative returns Black homeowners face in recessions cannot be explained by our observables. However, our results imply that at least on average, racial disparities in the total rate of return on owner-occupied housing reduce the racial wealth gap over time, once we account for rental yields and all other relevant cashflows.

References

John Y. Campbell. Household Finance. *Journal of Finance*, 61(4):1553–1604, 2006.

Ellora Derenoncourt, Chi Hyun Kim, Moritz Kuhn, and Moritz Schularick. Wealth of two nations: The U.S. racial wealth gap, 1860-2020. *Quarterly Journal of Economics, forthcoming*, 2023.

Robert Barsky, John Bound, Kerwin Ko' Charles, and Joseph P Lupton. Accounting for the black–white wealth gap: a nonparametric approach. *Journal of the American statistical Association*, 97 (459):663–673, 2002.

Moritz Kuhn, Moritz Schularick, and Ulrike I Steins. Income and wealth inequality in america, 1949–2016. *Journal of Political Economy*, 128(9):3469–3519, 2020.

Arpit Gupta, Christopher Hansman, and Pierre Mabille. Financial Constraints and the Racial Housing Gap. *working paper*, 2023.

Kerwin Kofi Charles and Erik Hurst. The Transition to Home Ownership and the Black/White Wealth Gap . *Review of Economics and Statistics*, 84(2):281–297, 2002.

Amir Kermani and Francis Wong. Racial Disparities in Housing Returns. *working paper*, 2022.

Prottoy A. Akbar, Sijie Li Hickly, Allison Shertzer, and Randall P. Walsh. Racial Segregation in Housing Markets and the Erosion of Black Wealth. *The Review of Economics and Statistics*, pages 1–45, 11 2022. ISSN 0034-6535. doi: 10.1162/rest_a_01276. URL https://doi.org/10.1162/rest_a_01276.

Edward N. Wolff. Heterogenous rates of return on homes and other real estate: Do the rich do better? Do Black households do worse? *working paper*, 2022.

Matthew E. Kahn. Racial and ethnic differences in the financial returns to home purchases. *Real Estate Economics*, 52(3):908–927, 2024. doi: https://doi.org/10.1111/1540-6229.12475. URL https://onlinelibrary.wiley.com/doi/abs/10.1111/1540-6229.12475.

David Chambers, Christophe Spaenjers, and Eva Steiner. The rate of return on real estate: Long-run micro-level evidence. *The Review of Financial Studies*, 34(8):3572–3607, 2021.

Piet Eichholtz, Matthijs Korevaar, Thies Lindenthal, and Ronan Tallec. The total return and risk to residential real estate. *The Review of Financial Studies*, 34(8):3608–3646, 2021.

Andrew Demers and Andrea Eisfeldt. Total Returns to Single Family Rentals. *working paper*, 2001.

Caitlin Gorback and Gregor Schubert. The Impact of Cultural Preferences on Homeownership. *working paper*, 2024.

Jonathan Halket, Lara Loewenstein, and Paul S. Willen. The cross-section of housing returns. 2024.

Sonia Gilbukh, Andrew Haughwout, Rebecca J Landau, and Joseph Tracy. The price-to-rent ratio: A macroprudential application. *Real Estate Economics*, 51(2):503–532, 2023.

Stefano Colonnello, Roberto Marfè, and Qizhou Xiong. Housing yields. In *Proceedings of Paris December 2021 Finance Meeting EUROFIDAI-ESSEC*, 2024.

Theodore M. Crone, Leonard I. Nakamura, and Richard P. Voith. Hedonic estimates of the cost of housing services: Rental and owner-occupied units. *Price and Productivity Measurement*, 4:51–68, 2009.

Thesia I Garner and Kathleen Short. Accounting for owner-occupied dwelling services: Aggregates and distributions. *Journal of Housing Economics*, 18(3):233–248, 2009.

Carlos Avenancio-Leon and Troup Howard. The Assessment Gap: Racial Inequalities in Property Taxation . *Quarterly Journal of Economics*, 137(3):1383–1434, 2022.

Oscar Jorda, Katharina Knoll, Dmitry Kuvshinov, Moritz Schularick, and Alan Taylor. The Rate of Return on Everything, 1870–2015. *Quarterly Journal of Economics*, 2019.

Neil Bhutta and Aurel Hizmo. Do Minorities Pay More for Mortgages? . *Review of Financial Studies*, 34(2):763–789, 2021.

Peter Ganong, Damon Jones, Pascal J Noel, Fiona E Greig, Diana Farrell, and Chris Wheat. Wealth, race, and consumption smoothing of typical labor income shocks. *working paper*, 2024.

Sasha Indarte, Bronsyn Argyle, Ben Iverson, and Christopher Palmer. Explaining Racial Disparities in Personal Bankruptcy Outcomes . *working paper*, 2023.

Paul Goldsmith-Pinkham and Kelly Shue. The gender gap in housing returns. *The Journal of Finance*, 78(2):1097–1145, 2023.

Patrick Bajari and C. Lanier Benkard. Demand Estimation with Heterogeneous Consumers and Unobserved Product Characteristics: A Hedonic Approach. *Journal of Political Economy*, 113(6), 2005.

Patrick Bajari, C. Lanier Benkard, and John Krainer. House Prices and Consumer Welfare. *Journal of Urban Economics*, 58(3), 2005.

Morris A. Davis, Andreas Lehnert, and Robert F. Martin. The rent-price ratio for the aggregate stock of owner-occupied housing. *The Review of Income and Wealth*, 54(2):279–284, 2008.

Jonah B. Gelbach. When Do Covariates Matter? And Which Ones, and How Much? *Journal of Labor Economics*, 34(2), 2016.

Carlos F Avenancio-León and Troup Howard. The assessment gap: Racial inequalities in property taxation. *The Quarterly Journal of Economics*, 137(3):1383–1434, 2022.

Benjamin Elbers. Trends in u.s. residential racial segregation, 1990 to 2020. *Socius*, 7:23780231211053982, 2021. doi: 10.1177/23780231211053982. URL https://doi.org/10.1177/23780231211053982.

Junting Ye, Shuchu Han, Yifan Hu, Baris Coskun, Meizhu Liu, Hong Qin, and Steven Skiena. Nationality Classification Using Name Embeddings. *Proceedings of the 2017 ACM on Conference on Information and Knowledge Management*, page 1897–1906, 2017.

Rebecca Diamond, Tim McQuade, and Franklin Qian. The effects of rent control expansion on tenants, landlords, and inequality: Evidence from san francisco. *American Economic Review*, 109(9):3365–3394, 2019.

Appendix

A Formal Statement of Rank Stability Assumption

We state a version of the rank stability assumption behind our rent imputation method that accounts for the fact that rent and price observations on a given home cannot be seen at the same time. Our rent imputation procedure uses "switcher" homes that either have a rent observation at time t and a house price observation at time t+1, or a house price observation at time t and a rent observation at time t+1. In addition, it is possible that the observed rent and price ranks we see do not perfectly satisfy our single index assumption. Allowing for possible violations of the single index assumption, we assume that the joint distribution of rent and price ranks is stable over time.

Assumption 2 *(Rank stability). Let $(\rho_{rh,t}, \rho_{rp,t})$ be the rank of house h's rent r and price p within the distribution of all observed rents and prices at time t. These ranks are normalized to be uniformly distributed on* $[0, 1]$

Let $F_{switcher,t,t+1}(\rho_{rh}, \rho_{ph}, X_h^o)$ be the joint CDF of rent ranks and price ranks observed for homes that switch between being rented and being owner-occupied between times t and t+1, conditional on observables X_h^o. We assume that this joint CDF is the same function $F_{switcher}(\rho_{rh}, \rho_{ph}, X_h^o)$ at all times t.

Under the idealized version of our single index assumption, we showed in equation 9 that one can exactly solve for the price rank of a house by knowing its observable characteristics and its rent rank. Assumption 2 allows for there to be a non-degenerate distribution of price ranks given an observed rent rank but requires this distribution not to vary over time. If, for example, renters care more about the (unobserved) quality of restaurants and coffee shops near a home, while owners care more about the (unobserved) quality of local schools, this relative difference in preferences could lead to small violations of our single index assumption in the main text.

Even if our single index assumption does not hold, the joint distribution $F_{switcher}(\rho_{rh}, \rho_{ph}, X_h^o)$ can be observed directly from data. As a result, the median regression we run in equation 10 still provides a median-unbiased prediction of a house's rent rank given its price rank and observables without any additional assumptions. However, the purpose of the regressions we run on the switcher population is to impute rent estimates for the population of owner-occupied houses. We therefore have to assume that the rents we observe for switchers are representative of the rent of owner-occupied homes, conditional on home observables and prices.

Assumption 3 *(Representativeness) Let $F_{owner}(\rho_{rh}, \rho_{ph}, X_h^o)$ be the joint distribution of rent ranks between what an owner-occupied house would hypothetically rent for and the price rank of its observed price, conditional on observables X_h^o. We assume that this is the same joint distribution as the object $F_{switcher}(\rho_{rh}, \rho_{ph}, X_h^o)$ we can estimate from the switcher population*

This representativeness assumption trivially holds under our single index assumption, since one can perfectly solve for the rent value of a home by knowing its observables and purchase price. However, it is strictly more general. The assumption allows for some degree of selection between owned and rented properties, if for example unobservably higher quality homes tend to be owned rather than rented. However, all of the selection on unobservables must be due to unobservables which are reflected in the house's price. Under this assumption, a median-unbiased prediction of the rent rank of a switcher also yields a median-unbiased prediction of the rent rank of an owner-occupied home with the same observables and price rank.

B Imputing Rents Using Corelogic and MLS

This appendix describes our methodology for estimating the rental yields of owner-occupied housing in the Corelogic database. Unlike the AHS, Corelogic is a large database of realized property transactions rather than a year-by-year survey of houses regardless of their transaction status. Corelogic has the benefit of being a considerably larger database than the AHS, but it has the downsides of not being geographically representative and only reporting property values at the time of a transaction. We therefore develop a modified rent imputation procedure tailored to the strengths and weaknesses of this database.

Our analysis begins with a database of rents from Multiple Listing Services (MLS), where real estate agents list properties for sale and for rent. We use all rental observations from the MLS where the rent paid by an actual tenant who occupies the property is reported and merge these rent observations to our property transaction data using each property's APN. This gives us a database of "switcher" properties for which we jointly observe rents and sales prices, allowing us to use the sales price of a property to proxy for its unobserved quality. Because this database is considerably larger than the AHS, we are able to run regressions predicting the rent of switcher properties separately by MSA-Year.

With each MSA-Year, we collect all MLS rental observations and merge them to all housing purchases on the same housing unit that occur within at most 10 years. We restrict this data to only include transactions where either the buyer or seller is an owner occupant, based on matching the name of the resident in Infutor to the Corelogic deed transaction. We then use our name-location-based measure of the owner's race, housing unit covariates from Corelogic's tax history file, and the price of the property to predict our rental yields. In doing so, we must account for the fact that the year at which the housing unit is rented need not be the same year in which the housing unit transacts in the purchase market. To do so, we compute the quantile of each housing unit's price in the distribution of house values in the American Community Survey in the same year and county.[7] We then assign the house an imputed purchase price with the same quantile during the year of the rental transaction as the house had during the year of the purchase transaction.

For our baseline specification, we run the following regression of housing rents on housing prices and

[7]For some observations in the ACS, we observe the PUMA but not the county FIPS of a home for data privacy reasons. For such homes, we compute the fraction of residents in the PUMA that reside in each neighboring county. We then treat the observation as if a fraction of it was seen in each neighboring county, weighted by the population weights of the PUMA. Our marginal distributions of house prices are computed after these fractional assignments of units which are missing a county FIPS.

other covariates in a given county and a given year

$$rent_{ht} = \alpha_{jt} + \beta_{jt} X_{ht} + \epsilon_{ht}. \tag{19}$$

The regression includes the following covariates X_{ht}. First, it includes the full set of interactions between the price and its square, the absolute value of the difference in years between the year t at which the unit is rented and the year of the housing transaction from which the price $price_{ht}$ is derived, and an indicator function for whether this difference in years is positive. Any term composed of multiplying together one, two, three, or all four of these four variables is included in the regression. In addition, we include linear and quadratic terms in the square footage of the housing unit, a linear term in the age of the housing unit, a dummy variable for whether the housing unit is a single family residence rather than an apartment, a dummy for each racial group of the property owner (Black/Hispanic/White), and zipcode fixed effects. On top of this, we include when possible the number of bedrooms, bathrooms, and total rooms in a housing unit. For each of these three remaining variables, we only include them in the specification when they are missing in no more than 10% of observations. In addition, we trim the top and bottom 1% of our rent and price variables each year to account for errors in data entry and outliers. We require a minimum sample size of 1,000 observations to run regression 19 in a given MSA-year. If this threshold fails, we run the same regression pooled across years with year fixed effect added when there is a minimum of 2,000 observations, in which case we replace our zipcode fixed effects with zipcode-year fixed effects.

After estimating the coefficients in equation 19, we use them to construct fitted values for the rental value of owner-occupied housing units. We do this for any housing unit for which we see a transaction in 1990 or later and for which the covariates needed to impute a rent are not missing. Finally, we compute the rental yield of an owner-occupied housing using by dividing our predicted rent value by the house's resale value $price_{ht}$. In years where the house does not transact, we construct this resale value in the same manner as we did to construct our house price covariates for equation 19. We first compute the quantile of the house's value within the county-year specific house price distribution in the ACS when it transacted. We then find the price level at the year for which we are imputing a rent which would have the same quantile in that county-year price distribution. The final data objects we analyze are the ratio of our rent fitted values divided by this time-matched value of the resale value of a home.

Geographical Representativeness of Corelogic-MLS Sample

To construct our rental yield estimates in Corelogic, we need accurate data on a house's characteristics, the race of its owner, and both a resale value of the house and data on its rent. This restricts us to a sample of 2007-2016, since our data on a house's characteristics from Corelogic tax history begin in 2007. In addition, MLS data exists only in certain subsets of the country, meaning that the unconditional means of our estimated rental yields cannot directly be compared to the nationally representative results in the AHS.

Table B.1 documents how our sample differs from the overall country. Our data comes from places that are richer, less White, and more concentrated in the West than the country as a whole. The fact that the West tends to have both a large Hispanic population as well as high house prices and low rent-to-price ratios makes aggregate statistics from this data difficult to compare to the AHS. We therefore compare our Corelogic/MLS data on rental yields to the AHS only in a sample that is matched both on time and on geography.

Table B.1: Summary Statistics Based on County-Level ACS Data

	Whole US (Population Weighted)	Counties in Sample (Sample Weighted)
2012-2016 White Share (%)	61.96 [22.07]	56.31 [17.79]
2012-2016 Black Share (%)	12.28 [12.65]	13.39 [11.03]
2012-2016 Hispanic Share (%)	17.3 [16.90]	20.00 [14.70]
2012-2016 Median Housing Value ($)	231,282.74 [146,813.42]	273,331.84 [146,434.20]
2012-2016 Mean Household Income ($)	78,182.99 [20,647.29]	87,492.76 [21,114.44]
2012-2016 Northeast Share (%)	17.60	12.29
2012-2016 Midwest Share (%)	21.24	21.31
2012-2016 South Share (%)	37.59	33.28
2012-2016 West Share (%)	23.56	33.13

Note: This table presents summary statistics for the counties included in our Corelogic-MLS sample and compares them to summary statistics on the overall country. We compute these statistics by aggregating county-level data from the American Community Survey (ACS) in 2012-2016. 38.41% of the country's population lives within counties where we can construct rental yields using the Corelogic-MLS data. Standard deviations are in brackets.

Summary Statistics in our Corelogic-MLS Sample

Table B.2 shows the summary statistics of unit attributes and other covariates in the rental regression for switchers, owned, and rented homes in the Corelogic-MLS sample. We find that switcher homes are somewhat less expensive than the overall population of owner-occupied homes, despite having similar numbers of rooms, bedrooms, and bathrooms. Owner-occupied homes are also somewhat older and more likely to be single family homes. We find that the average monthly rent we imputed for owner-occupied homes is $2045.98, higher than the average $1626.57 we see for switchers. MLS renters have an average monthly rent of $1646.33, very close to the average for switchers. Rentals are less likely to be single family homes than either switchers or owner-occupied homes and have slightly lower square footage. Overall, MLS switcher homes seem to be of similar quality to other rental homes, but owner-occupied homes seem to be of higher quality. The average owner-occupied home sells for $487,928, well above the average $318,759 price of a switcher. Because our rent imputation procedure uses a home's price to help measure quality, it is ideally suited to our sample, where the price differences between switcher and owner-occupied homes seem larger than the differences in their hedonic characteristics. These patterns in relative quality across owners/renters/switchers are similar to those presented in table 3.1 for the AHS.

Table B.2: Summary Statistics of Corelogic-MLS sample

	MLS Switchers	Owner-Occupied Sample	MLS Rentals
Monthly Rent ($)	1,626.57	2,045.98	1,646.33
	[908.38]	[1,270.25]	[965.16]
Monthly Rent ($, in 2019)	1,826.69	2,298.12	1,837.75
	[1,012.53]	[1,423.88]	[1,077.82]
Year of Rental Contract	2,012	2,012	2,012
	[2.57]	[2.75]	[2.59]
Building Square Footage	1,919.88	2,210.11	1,672.86
	[2,008.39]	[1,165.03]	[746.43]
Total Number of Rooms	6.28	6.63	6.00
	[1.84]	[1.69]	[1.70]
Number of Bathrooms	2.51	2.43	2.36
	[1.66]	[1.67]	[0.88]
Number of Bedrooms	3.12	3.20	2.93
	[1.16]	[0.98]	[1.05]
Building Age	24.30	36.27	27.62
	[24.10]	[25.09]	[23.44]
Share of Single Family Residence (%)	80.29	91.20	65.88
Transaction Price ($)	262,606.96	366,449.89	-
	[213,587.63]	[363,719.55]	-
Transformed Price ($)	261,200.44	343,419.64	-
	[227,919.37]	[366,145.06]	-
Transaction Price ($, in 2019)	318,759.17	487,927.87	-
	[252,927.95]	[499,419.48]	-
Transformed Price ($, in 2019)	294,605.59	386,524.52	-
	[256,667.19]	[410,566.70]	-
Price Quantile (%, County Level)	48.53	52.23	-
	[24.95]	[27.14]	-
Year of Transaction	2,009	2,006	-
	[5.09]	[7.24]	-
Diff.: Transaction Year - Rental Year	-3.47	-6.35	-
	[5.01]	[7.50]	-
Absolute Diff.: Transaction Year - Rental Year	5.29	7.94	-
	[3.03]	[6.25]	-
Transaction Before Rental	0.72	0.70	-
	[0.45]	[0.41]	-
White Share (%)	87.85	82.88	-
Black Share (%)	4.70	6.29	-
Hispanic Share (%)	7.45	10.83	-
Observations	2,845,590	192,329,864	2,516,997

Note: This table presents summary statistics on rents and the main covariates we use to impute rents using equation 12. We have three populations: 1. "MLS Switchers" that have both rent observations from MLS and transaction prices from Corelogic deed. 2. "Owner-Occupied Sample" of all homes for which we are able to successfully impute a rent using equation 12. 3. "MLS Rentals" of all closed rental contracts in MLS, excluding advertised rents that do not result in a tenancy. The sample includes properties (single-family residences, condos, and townhouses) transacted or rented from 2006-2017 in the counties where the Corelogic-MLS data is available. The switchers' sample is at the transaction-rental pair level. The owner-occupied sample is at the housing unit-individual-year level. The rental sample is at the transaction level. We weight our data to ensure comparability. For switchers, data are weighted by the inverse of the number of transaction-rental pairs per housing unit each year. For the owner-occupied sample, data are weighted by the inverse of the number of residents in a housing unit each year. For the rental sample, data are weighted by the inverse of the number of rentals per housing unit each year. We topcode the total number of rooms at 10 for both the owner-occupied and rental samples. Additionally, for the MLS rental sample, we winsorize rent and unit square footage at the 1st and 99th percentiles and topcode the number of bathrooms at 5 and the number of bedrooms at 10. Standard deviations are in brackets.

C Property Tax Rates in Corelogic

This section shows that our results on racial disparities in property taxes are broadly consistent with the analysis of property taxes in Avenancio-León and Howard [2022]. Their headline result considers racial disparities in the "assessment gap," the ratio between the assessed value of a home which determines property tax payments and the market value of a home in transactions. They find that Black homeowners have an assessment gap that is 12% higher than that of White homeowners. On the surface, this result seems to differ from our finding in table 6.6, where we find nearly the same property tax rate paid by each racial group. However, we show in table C.1 that our results can be reconciled and are internally consistent.

Their headline number of a 12% Black-White gap in assessment ratios comes from comparisons within taxation jurisdictions. We do not have data directly on the universe of tax jurisdictions, however they then add census tract-year fixed effects to this regression, which shrinks their Black-White assessment gap to 6.4%. We can run this specification in our data. The first column of table C.1 runs a regression that is nearly identical to the second regression in column 2 of table 2 of Avenancio-León and Howard [2022]. They find that in their sample from 2003-2016, Black homeowners have a 6.4% higher assessment ratio than White homeowners, after controlling for census-tract-year fixed effects. With these same controls, we find that Black homeowners have a 5.6% higher assessment ratio in our sample from 2007-2016.[8] Our 95% confidence interval of [.046, 0.067] contains their estimate.

In our next 3 columns, we first add Hispanic homeowners to our regression. Next, we reduce the geographic granularity of our fixed effects to MSA by year to match our analysis in the AHS. Finally, we add California back to our sample, which was dropped previously, and drop outlier property tax rates above 10% and assessment ratios that are not between 0.5 and 2. Column 4 shows that leaves us with a similar Black-White assessment gap of 6.5%.

In column 5 of table C.1, we switch the dependent variable of our regression from a homeowner's assessment ratio to the rate of return the household would get due only to property tax payments (setting appreciation, rents, and maintenance to zero), holding the sample fixed. Here we find a 38 basis point lower rate of return due to taxes for Black homeowners and 17.5 basis point lower return for Hispanic homeowners than for White homeowners. In the next two columns, we then add back in home transactions which were not preceded by a purchase mortgage and then finally include the entire housing stock inclusive of non-transacting homes. For non-transacting homes, we observe the nearest transaction in time and assign a property value at the time of the property tax observation that has the same value quantile, following the same procedure described in section B. In this final specification, we observe a 16.5 basis point lower rate of return from taxes for Black homeowners and 8.1 basis point lower rate of return for Hispanic homeowners relative to White homeowners.

In the final column of the table, we run an identical regression in the AHS, restricting ourselves to the MSAs that exist in our Corelogic data and matching our 2007-2016 time sample. Here we find a 16.6 basis point lower rate of return for Black homeowners and 21 basis point lower rate of return for Hispanic howeowners relative to White homeowners. This almost exactly matches the equivalent Corelogic specification for Black homeowners and if anything has a slightly lower rate of return for Hispanic

[8] Both our regression and that of Avenancio-León and Howard [2022] exclude California due to the unique impact of proposition 13 on its property tax assessments. We also restirct our sample to properties that transact in the sample year as the assessment measure and had a prior transaction in sample for the sample property that also originated a mortgage. We follow these sample restrictions from Avenancio-León and Howard [2022].

homeowners. This shows that within a geography- and time-matched sample, the racial disparities in property tax rates in the AHS are consistent with our data in Corelogic and with Avenancio-León and Howard [2022]. However, in the overall, nationally representative AHS sample from 1974 to 2021, racial disparities in property tax rates are close to zero on average.

Table C.1: Racial Gaps in Property Assessments and Rate of Return due to Taxes

	CoreLogic: Log(Assessment Ratio)				CoreLogic: Log(1 + Property Tax Rate)			AHS: Log(1 + Property Tax Rate)	
	Census Tract × Year FEs	Black and Hispanic	MSA × Year FEs	Add California/Drop Outliers	MSA × Year FEs	All Transacted	Non-transacted	Year FEs	MSA × Year FEs
Black	0.05646***	0.05795***	0.10589***	0.06473***	0.00381***	0.00336***	0.00165***	0.00248***	0.00166***
	(0.00536)	(0.00526)	(0.00795)	(0.00303)	(0.00023)	(0.00017)	(0.00014)	(0.000511)	(0.000485)
Hispanic		0.01921***	0.01811***	0.05013***	0.00175***	0.00159***	0.00081***	-0.000566	0.00211***
		(0.00261)	(0.00456)	(0.00233)	(0.00008)	(0.00006)	(0.00006)	(0.000458)	(0.000429)
Observations	5,942,075	6,088,651	6,088,651	5,410,397	5,410,397	17,921,862	200,452,594	26,857	26,857

Note: This table documents racial gaps in property assessment ratios and in effective property tax rates for the counties in our Corelogic sample from 2007-2016. The sample includes any owner-occupied housing units (single-family residences, condos and townhouses) with transactions recorded from 1990 onward and information on the owner's race. Standard errors (clustered at the Zip Code level for CoreLogic sample and the unit level for AHS sample) are presented in parentheses.

D Foreclosure Discount Measurement

This appendix considers the robustness of our measure of the rate of return on owner-occupied housing to different ways of measuring foreclosure discounts. In a foreclosure, borrowers who do not pay an existing mortgage have their homes auctioned off in order to recoup losses for the lender. We consider two ways of measuring the price a borrower receives in a foreclosure. First, we examine their unlevered return on housing (ignoring dividends) as if they sold their home at the foreclosure auction price. Second, we compare this to a rate of return (ignoring dividends) where they are considered to sell their home to their lender at the balance of their outstanding mortgage, with any remaining losses on the foreclosure being born by the lender rather than the borrower.

We estimate the outstanding balance on a property's mortgage as follows. First, we assume that borrowers pay a mortgage interest rate equal to the average 30-year fixed mortgage interest rate at the month of mortgage origination. Then, we use the following expression to calculate the current outstanding mortgage balance:

$$\text{Outstanding Mortgage Balance} = P \times \frac{(1+\frac{r}{12})^n - (1+\frac{r}{12})^m}{(1+\frac{r}{12})^n - 1}. \tag{20}$$

Here P is the principal amount of the mortgage loan, r is the annual mortgage interest rate, n is the total number of payments (e.g., n equals 360 for a 30-year fixed-rate mortgage), and m is the number of payments the homeowner has made. For foreclosure cases, we assume the default occurs 6 months before an observed foreclosure auction, resulting in 6 months of unpaid interest. Our final expression for the amount the borrower owes at the time of foreclosure is then

$$\text{Outstanding Mortgage Balance Amount} \times [1 + (\frac{r}{12} \times 6)].$$

For each foreclosed-upon homeowner, we calculate the annual house price appreciation rate as

$$\text{Annualized Log Appreciation Rate} = \frac{log(P_t) - log(P_0)}{\text{Num. Years of Occupation}},$$

where P_t is the sale price and P_0 is the purchase price. We examine how our average rate of house price appreciation varies with how we compute the sale prices in foreclosures.

We compare our data in Corelogic to a more detailed dataset, RIS, with information on all foreclosures in Cook County, IL. We match these two datasets based on a property's APN and the year and month of a foreclosure transaction from 2000 to 2016.[9] In each foreclosure, we observe a "complaint amount" which represents the amount of remaining money the lender believes is owed by the borrower.[10] We show that this complaint amount is approximated well by our method using outstanding mortgage balances, at least for the set of households who never refinance their mortgages.

Table D.1 presents regressions that quantify the relationship between different measures of a home's value in foreclosure in Cook County, IL. For households which never refinance their mortgages, our outstanding mortgage balance imputation predicts the complaint amount with a slope of 1.037 and an R-squared of .896. This shows that our mortgage balance imputation procedure is an accurate proxy for the true complaint amount in a foreclosure for homeowners who do not refinance. This R-squared falls to .599 when we separately examine homeowners who refinance, since we are not able to tell which refinanced mortgages did or did not involve prepaying old ones. The relationships we find between complaint amount and foreclosure auction prices, as well as auction prices and imputed mortgage balances, are weaker. This shows that a foreclosure auction price is not a particularly accurate measure of the amount of debt a homeowner erases in a foreclosure. These results suggest that we accurately estimate the value of debt a non-refinancing homeowner erases in a foreclosure with our imputed outstanding mortgage balance. However, auction prices seem only moderately related to the complaint amounts reported in RIS.

Table D.1: Regressions showing the relationship between foreclosure price measures in Cook County, IL using RIS and Corelogic Data

Sample	Y-variable	X-variable	Coefficient	R-squared
All Foreclosures	Complaint/Judgement Amount	Foreclosure Auction Price	0.839 (0.009)	0.567
Non-Refinancer	Complaint/Judgement Amount	Outstanding Mortgage Balance	1.037 (0.002)	0.896
Non-Refinancer	Foreclosure Auction Price	Outstanding Mortgage Balance	0.719 (0.016)	0.485
Refinancer	Complaint/Judgement Amount	Outstanding Mortgage Balance	1.255 (0.005)	0.599
Refinancer	Foreclosure Auction Price	Outstanding Mortgage Balance	0.993 (0.014)	0.543

Note: This table presents regression results of each price measure listed as the Y-variable on the alternative price measure listed as the X-variable. We control for interacted fixed effects both for the year of foreclosure and the year of mortgage origination. The data are at the foreclosure case level, with multiple foreclosures on a home in the same month weighted as a single observation. Complaint/Judgement Amount comes from RIS, while Outstanding Mortgage Balance and Foreclosure Auction Price are from Corelogic. Standard errors are in parentheses.

Next, we consider in table D.2 how our choice of foreclosure price measurement impacts the rate of

[9] When we see foreclosures that match on APN and year but not month, we include these as well but weigh any multiple matches for an APN-year as a single observation.

[10] We drop matched cases with missing complaint amounts. We also limit our analysis to foreclosures that report a purchase price and at least one of a purchase mortgage amount or foreclosure auction price in Corelogic. We restrict our analysis to complaint amounts, purchase mortgages, and auction prices between \$10,000 and \$2,000,000. Additionally, for cases with non-missing purchase mortgages, the initial LTV must be less than 1.05 at purchase.

house price appreciation earned by homeowners who experience foreclosure. We find that using auction prices results in considerably lower rates of house price appreciation than either our mortgage balance imputation or using RIS complaint amounts. For homeowners who never refinance, the mortgage balance yields a rate of house price appreciation that is only around 1-2 percentage points lower than the rate implied by using the complaint amount, while the gap is larger for homeowners who refinance. In all specifications, we find that the RIS complaint amount results in foreclosed homeowners earning a higher rate of house price appreciation than either using a mortgage balance imputation or an auction price. If we believe that the return earned by a homeowner should be measured by them "selling" their house for the amount of outstanding debt they erase in the foreclosure, then our mortgage imputation method should yield a somewhat more negative rate of return than the baseline truth provided in RIS.

Finally, in table D.3, we analyze homeowners' rates of house price appreciation in the 4 counties where we have carefully measured foreclosure sales. In Cook County and especially Philadelphia County, we find that foreclosure auctions imply considerably lower rates of house price appreciation earned by foreclosed homeowners than our method that imputes outstanding mortgage balances. For King County and Santa Clara County, we find if anything that foreclosure auction prices suggest a slightly higher rate of house price appreciation. In panel B of this table, we find that overall rates of house price appreciation are only minimally sensitive to our method of computing foreclosure discounts, except in Philadelphia where overall rates are 1.3% higher using the mortgage-based method than the auction price in foreclosure cases.

Table D.2: Log Housing Appreciation Rates For Foreclosed Properties Based on Different Measures of Foreclosure Discounts in Cook County, IL

	Complaint Amount-based (RIS)				Outstanding Mortgage-based (Corelogic)				Foreclosure Auction Price-based (Corelogic)			
	Mean	Median	Std. Dev.	Num. Obs.	Mean	Median	Std. Dev.	Num. Obs.	Mean	Median	Std. Dev.	Num. Obs.
Panel A: All Foreclosures in Corelogic/RIS Match												
Non-Refinancer	-0.050	-0.038	0.094	43,914	-0.069	-0.053	0.086	43,914	-	-	-	-
Refinancer	-0.010	-0.010	0.062	45,286	-0.054	-0.046	0.049	45,286	-	-	-	-
All Foreclosures	-0.029	-0.024	0.082	89,200	-0.061	-0.048	0.070	89,200	-	-	-	-
Panel B: Only Foreclosures with Non-Missing Auction Price												
Non-Refinancer	-0.038	-0.040	0.192	2,383	-0.051	-0.053	0.170	2,383	-0.168	-0.142	0.250	2,383
Refinancer	-0.011	-0.017	0.118	3,985	-0.046	-0.048	0.110	3,985	-0.077	-0.069	0.143	3,985
All Foreclosures	-0.021	-0.026	0.151	6,368	-0.048	-0.050	0.135	6,368	-0.111	-0.092	0.195	6,368

Note: This table presents summary statistics on log housing appreciation rates based on different measures of foreclosure discounts in Cook County, IL. The data includes cases with observed foreclosure auctions from 2000 to 2016 in Corelogic deeds which can be matched to RIS. All numbers are adjusted for inflation.

Table D.3: Overall Log Housing Appreciation Rates Based on Different Measures of Foreclosure Discounts

	Outstanding Mortgage-based				Foreclosure Auction Price-based			
	Mean	Median	Std. Dev.	Num. Obs.	Mean	Median	Std. Dev.	Num. Obs.
Panel A: Foreclosure Sales Only								
Cook	-0.026	-0.019	0.042	36,249	-0.041	-0.016	0.134	36,249
King	-0.035	-0.029	0.057	24,619	0.009	0.029	0.105	24,619
Santa Clara	-0.051	-0.042	0.078	10,862	-0.038	-0.006	0.200	10,862
Philadelphia	-0.017	-0.013	0.031	31,101	-0.174	-0.086	0.294	31,101
Panel B: All Sales								
Cook	0.072	0.073	0.140	1,408,185	0.072	0.073	0.142	1,408,185
King	0.088	0.086	0.111	806,746	0.090	0.086	0.111	806,746
Santa Clara	0.093	0.096	0.182	391,749	0.093	0.096	0.185	391,749
Philadelphia	0.080	0.067	0.156	310,625	0.063	0.067	0.197	310,625

Note: This table presents summary statistics on log housing appreciation rates for all complete housing spells observed in Corelogic in four counties: Cook County, IL, Philadelphia County, PA, King County, WA, and Santa Clara County, CA. Panel A presents only spells that end with foreclosure sales, while Panel B presents all completed spells, including regular sales. All numbers are adjusted for inflation.

E Data Section: AHS

The data come from the American Housing Survey (AHS). The AHS is a nationally representative panel survey of housing units that has been conducted by the U.S. Census Bureau since 1973. The latest survey we use comes from 2021.

Housing costs

All monetary values are CPI-adjusted to 2019.

Property values

Property values in the AHS are directly reported by respondents, but the way they are recorded changes over time.

- Between 1973 and 1983, property values are categorical with bracket sizes that increase from \$2,500 for low-value properties to \$50,000 for high-value properties. We assign the mean value in each bracket to all properties in the bracket. To determine which value to assign to the top-coded bracket we utilize information on the distribution of property values in 1985 – when property values are first reported on a continuous scale. Specifically, we determine the percentile of properties in the top-bracket in each year before 1985 and compare this percentile to the distribution of property values in 1985. We then adjust the value for the top-coded bracket in each year before 1985 to match the 1985 distribution. Since property values in 1985 are top-coded, in some cases, i.e. when the percentile of properties prior to 1985 surpasses the top-coded percentile in 1985, there is no adjustment and the value assigned to units in the top-bracket remains unchanged.
- After 1985, property values are reported on a continuous scale and top-coded at levels that vary by metropolitan area.

There are some property values that are most likely misreported. First, we flag any property value below $10,000 in 2019 dollars are bad. This flags 1.5 percent of property value reports. Next, we flag the top and bottom one percent of annualized appreciation rates between survey waves as unreliable. This corresponds to annual returns less than -65 percent or greater than 195 percent. This flags extreme appreciation rates, but does not indicate whether the house value at time t or $t+1$ is the outlier value. To identify the outlier property value, we first look at adjacent appreciation rates for these houses with previous/subsequent survey ways. If the appreciation rate from $t-1$ to t is not an outlier, then we flag the property value at time $t+1$ as the bad value. If the appreciation rate from time $t+1$ to $t+2$ is not an outlier, then we identify the value at time t to be driving the outlier appreciation. For the remaining property values that cannot be flagged as causing the outlier appreciation rates from these methods, we use a hedonic regression of property values on property characteristics. Among the property values contributing to the outlier appreciation rates, we flag the property value as bad as the one which is further from the hedonic predicted property value.

We then impute missing property values. Property values are essentially never missing in the raw data for owned units since 1985. Prior to 1985, properties that on 10+ acres of land or were mobile homes were not asked their property values. We will exclude these properties from our data, as imputing them would be very hard since there are no observed property values for similar housing in the AHS during that time period.

We also need to impute property values for units that are owned in year t, but are not observed in the subsequent survey wave as still owner-occupied. If they are vacant or rented, the property value is not asked of the survey respondent, so we need to impute it. Alternatively, the unit could have exited the sample.

To impute missing property values, we run a hedonic regression of observed appreciation rates on property and household characteristics. We use the same covariates in our appreciation imputation as those included in our regressions decomposing our racial differences in total rates of return. We then impute the missing property value by taking the prior observed property value of the unit and multiplying it by the predicted appreciation rate from the hedonic regression.

Property taxes

Property taxes are reported on a continuous scale prior to 1985 and reported in steps of $50-$100 thereafter. For reports after 1985, we assign the mean value of each bracket to all observations in the bracket. Property taxes are also top-coded. We predict the property tax in the top-coded bracket using a regression of property taxes on property values and region-by-metro fixed effects. To impute missing reports and impute property taxes for units that are owner-occupied in t but not t+1 we regress property taxes on property value, race, year, building year cohort, and region-by-metro fixed effects.

Maintenance cost

Our variable for maintenance cost includes routine maintenance cost as well as home improvement projects. The AHS contains information on maintenance costs after 1985. Routine maintenance is asked for directly and we determine home improvement costs by aggregating information from a number of variables on cost for home improvement jobs (e.g. for kitchen, bath, etc.), alterations, and additions, over the past two years and halving the aggregate to obtain annual maintenance cost.

To impute missing maintenance cost for units that are owned in t, we regress maintenance cost on property value, property tax, building age (current year - built year), building age squared, unit size (in square feet), geographic region, and race dummies. The regression equation is also fully interacted with the race dummies. The maintenance cost variable is only available after 1985 so we also impute maintenance cost prior to 1985. To do so, we predict maintenance cost in the last period the unit is owned, using the same regression as described above, and linearly extrapolate backwards using the average change in maintenance cost across the reporting horizon (1985 to 2021). Last, we predict maintenance costs for units that are owner-occupied in t but no longer in t+1 using the same regression of maintenance cost on observables as above.

F Data Section: Corelogic

This section presents details on how we assembled our large database of housing transactions in Corelogic. Corelogic data comes from 4 separate sources: tax data, deed data, tax history data, and MLS data. Tax data provides information about the names of residents and the characteristics of individual homes at the end of our sample in 2017. Tax history provides analogous data going back each year until 2007. Deed data provides a systematic list of all housing purchases and mortgage issuances, providing information about the names of people active in each transaction, the price of any purchase or sale, the quantity lent in any mortgage, and a unique house identifier. Finally, the subset of MLS data we use provides information about the monthly rent advertised to live in a property at a given point in time. All of these datasets are easily merged to each other based on the Assessor's Parcel Number (APN) assigned to each property.

Corelogic-Infutor Merge

We merge all of these Corelogic databases to another database, Infutor, which provides information about the residents of each property. Infutor provides individual-level information about the name of residents, the sequence of addresses where they have lived over time, and a reported time at which they moved in and out. Merging Infutor to Corelogic allows us to compare the names of a home's owner listed in deed or tax data to the name of the resident, so that we are able to identify owner-occupied housing.

Our merge of Infutor to Corelogic proceeds in two steps. First, we attempt to merge the data based on the listed zipcode and address string of each property. The string merge takes place in several steps. First, we try to match addresses on entire strings, which yields 74% of our entire set of matched observations. We then look at additional matches with either have different zip codes, different city/town names but identical zip codes, or missing apartment numbers or matches only on the numeric portion of apartment numbers (e.g. apartment 3 instead of 3A).

After our string merge, we attempt to use a spatial join to match additional addresses. To do so, we take the list of all our addresses from Infutor and feed them to the Geocodio geocoding service, who provided us with detailed latitudes and longitudes for each address. Next, we use parcel Polygons from a data provider Landgrid to map each APN in Corelogic to an associated geographic area. Finally, we observe all of the matches of latitudes and longitudes within each APN-associated area to construct matches from Infutor addresses to Corelogic APNs.

Table F.1 describes, of those rows of Infutor who are (A) not PO-Boxes, (B) whose cleaned addresses are not singletons (so that we see a nontrivial address history for the person), and (C) which successfully map to CoreLogic addresses, what proportion of mappings derive from each stage. There are of the 812 million rows that meet conditions (A) and (B), 84.0% can be successfully merged.

Table F.1: Success Rate of Merge between Infutor Person-Address Observations and Corelogic Addresses

Overall Match Rate	
Number of Total Infutor Observations	812 million
Fraction of Infutor Observations Merged	84.0 %
Fraction of Match Types Among Merged	
Complete String Matches	74.3 %
Incomplete String Matches	16.4 %
Spatial Join Matches	9.3 %

Determining if a Home is Owner-Occupied

Our Infutor-Corelogic match allows us to observe both the names of a housing unit's residents in Infutor and the names of its owners using deed and tax data from Corelogic. We determine whether a home is owner-occupied based on comparing these resident and owner names. For each housing unit, we obtain the most recent owner's name (or possibly multiple owner names) from Corelogic tax data and the names of all buyers and sellers of the property listed in Corelogic deed data in our sample. Since CoreLogic owner names sometimes include e.g. trusts (which sometimes contain owner last names i.e. "JOHN SMITH TRUST"), we remove a set of irrelevant strings that are unlikely to be named. We then employ two methods for identifying congruence of CoreLogic and Infutor names.

First, we calculate the Jaro-Winkler distance between the Infutor last name and each CoreLogic last name, which is a distance metric on the similarity of two strings. We say that a match is obtained if the string similarity is above 0.9. Second, we run a regular expression check for Infutor names being contained in each CoreLogic name (including searching for space- and hyphen-delimited tokens). If either of these matches, we say that the resident of a unit is an owner-occupant. Because we believe the Corelogic's buy and sell dates are more accurate than Infutor's dates when a resident moves in or out of a property, we use Corelogic's purchase and sale data to infer when an owner-occupant moves in or out of a property.

Identifying Homeowner Race

We impute homeowner race following [Ye et al., 2017, Diamond et al., 2019]. The first step uses the NamePrism machine learning algorithm to predict race given a homeowner's name. Next, the racial demographics of each census block are used to reweight the probabilistic race predictions of NamePrism using Bayes' rule. We then classify any homeowner with greater than a .8 chance of being White, Black or Hispanic as part of that racial group. This leads to a modest overestimate of the White share of the population, since the larger White population in the US will be reflected in the Bayesian weights coming from population shares. Moreover, Black and Hispanic homeowners are more likely to be identified in minority-dominant areas, leading to a greater degree of geographic segregation than in the true data. That said, name and geography are jointly a highly accurate predictor of race, and our large sample size allows for precise statistical inference even if we underestimate the size of the minority population.

For the set of owner-occupied homes we have identified, we use our name-based method of identifying race to infer the race of the homeowner. When there is a single homeowner whose name we can match between Infutor and Corelogic, this is straightforward. When there are several matching names in the same ownership spell, we assign the home the most common race of all of the house's matched owners. When we have the same number of homeowner name matches from two different racial groups in the same ownership spell, we choose of the two racial groups randomly and i.i.d. across such homes.

G Robustness: Rate of Return Using Alternative Rent Imputation Methods

This appendix considers the impact of three robustness checks on our measured racial disparities in the total rate of return on housing in the AHS. First, we replace our new method for rent imputation with a standard hedonic method and find that the disparities we measure would be even larger. Second, we run our quantile regressions for rent imputation separately across three distinct periods when the AHS refreshes its panel of homes and find similar results to the main text. Third, we separately estimate our rental yields first using only homes that switch from being rented to being owned and then only using homes that switch from being owned to being rented, whereas we use both directions of switchers together in our baseline approach.

Our hedonic rent imputation procedure is the same as used in section 4, using data on all AHS rentals to estimate the regression.

We present the results using hedonic rent imputation in tables G.1-G.2, which are analogous to tables 6.4 and 6.9 in the main text. In table G.1, the left column begins with a 2.85% higher rate of return for Black and 2.28% higher rate of return for Hispanic homeowners relative to White homeowners. In the main text, these disparities are only 0.62% and 1.19%. After accounting for our full set of controls in the fifth column of table G.1, we still have a remaining 1.8% Black-White return gap and 1.1% Hispanic-White return gap, while there is no remaining gap in the fifth column of table 6.4. This shows that if we had used hedonic methods to impute rental yields, we would have found an even higher rate of return for Black and Hispanic homeowners relative to White homeowners.

Table G.1: Log rate of return (Hedonic rent)

	(1)	(2)	(3)	(4)	(5)
Black	0.0285***	0.0280***	0.0290***	0.0158***	0.0184***
	(0.00142)	(0.00140)	(0.00139)	(0.00137)	(0.00134)
Hispanic	0.0228***	0.0240***	0.0252***	0.00826***	0.0108***
	(0.00163)	(0.00163)	(0.00163)	(0.00156)	(0.00155)
Observations	768,356	768,356	768,356	768,356	768,356
Dummy for other racial group	Yes	Yes	Yes	Yes	Yes
Year FE (not interacted)	Yes	-	-	-	-
Geography-by-year FE	No	Yes	Yes	Yes	Yes
Demographics (age, #kids)	No	No	Yes	Yes	Yes
Income, education, tenure	No	No	No	Yes	Yes
Housing characteristics	No	No	No	No	Yes

Note: Data from 1974-2021 owner-occupied properties in the AHS. Rent is imputed using a hedonic method. Standard errors are clustered at the unit level. The excluded race group is White owner-occupants. Geography FEs include the most granular-defined geography variables in the AHS for each year, which change some over time.

Table G.2: Log rate of return (Hedonic rent)

	(1)	(2)	(3)	(4)	(5)
Black	0.0294***	0.0294***	0.0304***	0.0168***	0.0193***
	(0.00170)	(0.00171)	(0.00171)	(0.00169)	(0.00163)
Black × Recession=1	-0.00294	-0.00488	-0.00473	-0.00350	-0.00318
	(0.00310)	(0.00314)	(0.00315)	(0.00316)	(0.00316)
Hispanic	0.0318***	0.0263***	0.0272***	0.00937***	0.0121***
	(0.00183)	(0.00190)	(0.00191)	(0.00184)	(0.00183)
Hispanic × Recession=1	-0.0350***	-0.00933*	-0.00812*	-0.00502	-0.00575
	(0.00359)	(0.00364)	(0.00367)	(0.00367)	(0.00361)
Observations	768,356	768,356	768,356	768,356	768,356
Recession-by-race FE (fully interacted)					
Dummy for other racial group	Yes	Yes	Yes	Yes	Yes
Year FE (not interacted)	Yes	-	-	-	-
Geography-by-year FE	No	Yes	Yes	Yes	Yes
Demographics (age, #kids)	No	No	Yes	Yes	Yes
Income, education, tenure	No	No	No	Yes	Yes
Housing characteristics	No	No	No	No	Yes

Note: Data from 1974-2021 owner-occupied properties in the AHS. Rent is imputed using a hedonic method. Standard errors are clustered at the unit level. The excluded race group is White owner-occupants. Geography FEs include the most granular-defined geography variables in the AHS for each year, which change some over time.

In table G.2, we show that using a hedonic rent measure, we would have found that Hispanic but not Black homeowners have a greater return exposure to the business cycle than White homeowners. The Black/ recession interaction coefficient is not significant in any specification, while in table 6.9 this coefficient is significant at either the 90% or 95% confidence level in all specifications. In both tables 6.9 and G.2, the excess exposure to the business cycle of Hispanics in table can entirely be accounted for with observable controls used in the 4th and 5th columns, where the income/education of the homeowner is included.

Table G.3: Log rate of return (Separating Panels)

	(1)	(2)	(3)	(4)	(5)
Black	0.00490***	0.00478***	0.00548***	-0.00463***	-0.00211
	(0.00118)	(0.00119)	(0.00119)	(0.00119)	(0.00119)
Hispanic	0.0146***	0.0122***	0.0124***	-0.000131	0.00225
	(0.00135)	(0.00138)	(0.00138)	(0.00136)	(0.00133)
Observations	742,629	742,629	742,629	742,629	742,629
Dummy for other racial group	Yes	Yes	Yes	Yes	Yes
Year FE (not interacted)	Yes	-	-	-	-
Geography-by-year FE	No	Yes	Yes	Yes	Yes
Demographics (age, #kids)	No	No	Yes	Yes	Yes
Income, education, tenure	No	No	No	Yes	Yes
Housing characteristics	No	No	No	No	Yes

Note: Data from 1974-2021 owner-occupied properties in the AHS. Rent is imputed using our switcher method and allows coefficients in switcher regression to change across the three distinct panels in the AHS: 1974-1983, 1985-2013, and 2015-2021. Standard errors are clustered at the unit level. The excluded race group is White owner-occupants. Geography FEs include the most granular-defined geography variables in the AHS for each year, which change some over time.

Tables G.3 and G.4 examine how our results change when we impute rental yields with separate regressions across 3 sub-periods when the AHS follows distinct panels of residents: 1974-1983, 1985-2013, and 2015-2021. Here, we find a 0.49% basis point rate of return for Black and 1.46% for Hispanic homeowners in the first column, close to the 0.62% and 1.19% we found in the first column of table 6.4. By column 5, we find that observable controls can account for effectively all of these racial disparities and slightly over-explain the Black-White return gap. In table G.4, we find like in table 6.9 that both Black and Hispanic homeowners have more cyclical returns than White homeowners. Like in the main text, observable differences by column 5 can account for most of the Hispanic-White gap but not the Black-White gap in exposure of housing returns to the business cycle.

Finally, tables G.5-G.8 use data constructed by running our switcher regressions for imputing rental yields by breaking the switcher population in two. Tables G.5-G.6 uses rental yields estimated only from properties that switch from owned to rented, while tables G.7-G.8 use rental yields estimated only from properties that switch from being rented to owned. We find that the estimated racial disparities in returns are somewhat higher when we use rent-to-own switchers than when we use own-to-rent switchers. However, even using the own-to-rent switchers, we find a 0.54% higher rate of return for Black homeowners and 1.04% for Hispanic homeowners relative to White homeowners in the first column of table G.5. These disparities increase to a Black-White return gap of 0.78% and a Hispanic-White return gap of 1.49% in the first column of table G.7. In both cases, we find by column 5 that observables can

Table G.4: Log rate of return (Separating Panels)

	(1)	(2)	(3)	(4)	(5)
Black	0.00793***	0.00799***	0.00862***	-0.00213	0.000189
	(0.00148)	(0.00149)	(0.00149)	(0.00149)	(0.00146)
Black × Recession=1	-0.0104***	-0.0112***	-0.0110***	-0.00859**	-0.00797**
	(0.00286)	(0.00291)	(0.00292)	(0.00293)	(0.00293)
Hispanic	0.0237***	0.0151***	0.0150***	0.00103	0.00359*
	(0.00159)	(0.00167)	(0.00168)	(0.00166)	(0.00163)
Hispanic × Recession=1	-0.0352***	-0.0111**	-0.0101**	-0.00537	-0.00590
	(0.00339)	(0.00343)	(0.00344)	(0.00346)	(0.00341)
Observations	742,629	742,629	742,629	742,629	742,629
Recession-by-race FE (fully interacted)					
Dummy for other racial group	Yes	Yes	Yes	Yes	Yes
Year FE (not interacted)	Yes	-	-	-	-
Geography-by-year FE	No	Yes	Yes	Yes	Yes
Demographics (age, #kids)	No	No	Yes	Yes	Yes
Income, education, tenure	No	No	No	Yes	Yes
Housing characteristics	No	No	No	No	Yes

Note: Data from 1974-2021 owner-occupied properties in the AHS. Rent is imputed using our switcher method and allows coefficients in switcher regression to change across the three distinct panels in the AHS: 1974-1983, 1985-2013, and 2015-2021. Standard errors are clustered at the unit level. The excluded race group is White owner-occupants. Geography FEs include the most granular-defined geography variables in the AHS for each year, which change some over time.

Table G.5: Log rate of return (Using only own-to-rent switchers)

	(1)	(2)	(3)	(4)	(5)
Black	0.00540***	0.00515***	0.00598***	-0.00465***	-0.00282*
	(0.00115)	(0.00116)	(0.00116)	(0.00116)	(0.00116)
Hispanic	0.0104***	0.00811***	0.00888***	-0.00462***	-0.00252
	(0.00135)	(0.00135)	(0.00135)	(0.00132)	(0.00129)
Observations	742,742	742,742	742,742	742,742	742,742
Dummy for other racial group	Yes	Yes	Yes	Yes	Yes
Year FE (not interacted)	Yes	-	-	-	-
Geography-by-year FE	No	Yes	Yes	Yes	Yes
Demographics (age, #kids)	No	No	Yes	Yes	Yes
Income, education, tenure	No	No	No	Yes	Yes
Housing characteristics	No	No	No	No	Yes

Note: Data from 1974-2021 owner-occupied properties in the AHS. Rent is imputed using our switcher method only using switchers that switch from owner-occupied status in the prior period to rented in the current period. Standard errors are clustered at the unit level. The excluded race group is White owner-occupants. Geography FEs include the most granular-defined geography variables in the AHS for each year, which change some over time.

Table G.6: Log rate of return (Using only own-to-rent switchers)

	(1)	(2)	(3)	(4)	(5)
Black	0.00722***	0.00752***	0.00826***	-0.00267	-0.000965
	(0.00145)	(0.00148)	(0.00148)	(0.00147)	(0.00144)
Black × Recession=1	-0.00618*	-0.00832**	-0.00804**	-0.00681*	-0.00641*
	(0.00283)	(0.00288)	(0.00289)	(0.00291)	(0.00291)
Hispanic	0.0190***	0.0102***	0.0106***	-0.00391*	-0.00158
	(0.00159)	(0.00165)	(0.00165)	(0.00162)	(0.00160)
Hispanic × Recession=1	-0.0335***	-0.00827*	-0.00674*	-0.00340	-0.00413
	(0.00335)	(0.00338)	(0.00338)	(0.00340)	(0.00334)
Observations	742,742	742,742	742,742	742,742	742,742
Recession-by-race FE (fully interacted)					
Dummy for other racial group	Yes	Yes	Yes	Yes	Yes
Year FE (not interacted)	Yes	-	-	-	-
Geography-by-year FE	No	Yes	Yes	Yes	Yes
Demographics (age, #kids)	No	No	Yes	Yes	Yes
Income, education, tenure	No	No	No	Yes	Yes
Housing characteristics	No	No	No	No	Yes

Note: Data from 1974-2021 owner-occupied properties in the AHS. Rent is imputed using our switcher method only using switchers that switch from owner-occupied status in the prior period to rented in the current period. Standard errors are clustered at the unit level. The excluded race group is White owner-occupants. Geography FEs include the most granular-defined geography variables in the AHS for each year, which change some over time.

either entirely account for or more than account for this racial disparity in returns.

We find in tables G.6 and G.8 nearly identical results for how much more cyclical the returns of Black and Hispanic homeowners are compared to White homeowners. In both tables, we find a 3.3-3.4% greater drop in returns for Hispanic homeowners relative to White homeowners during recessions, which can be almost entirely explained by observable controls in column 5. In addition, we find a 0.62-0.68% greater drop in Black homeowner returns relative to White homeowners during recessions in column 1 of both tables, and this disparity cannot be explained by controlling for additional observables.

The fact that we find smaller racial disparities when using own-to-rent switchers mitigates a potential worry that racial disparities in foreclosure rates could be biasing our results. Because only own-to-rent switchers can follow a homeowner's foreclosure, we would expect a greater degree of foreclosures among minorities to depress their house prices relative to the rents earned on such houses after a foreclosure. This would result in higher measured rental yields for minorities, increasing our measured racial disparities in returns. We find the opposite result that estimating rental yields using rent-to-own switchers instead results in greater racial disparities in measured returns on housing.

Table G.7: Log rate of return (Using only rent-to-own switchers)

	(1)	(2)	(3)	(4)	(5)
Black	0.00782***	0.00755***	0.00829***	-0.00160	0.000246
	(0.00114)	(0.00115)	(0.00115)	(0.00115)	(0.00116)
Hispanic	0.0149***	0.0123***	0.0126***	0.0000725	0.00239
	(0.00138)	(0.00140)	(0.00140)	(0.00137)	(0.00133)
Observations	742,817	742,817	742,817	742,817	742,817
Dummy for other racial group	Yes	Yes	Yes	Yes	Yes
Year FE (not interacted)	Yes	-	-	-	-
Geography-by-year FE	No	Yes	Yes	Yes	Yes
Demographics (age, #kids)	No	No	Yes	Yes	Yes
Income, education, tenure	No	No	No	Yes	Yes
Housing characteristics	No	No	No	No	Yes

Note: Data from 1974-2021 owner-occupied properties in the AHS. Rent is imputed using our switcher method only using switchers that switch from rented in the prior period to owner-occupied in the current period. Standard errors are clustered at the unit level. The excluded race group is White owner-occupants. Geography FEs include the most granular-defined geography variables in the AHS for each year, which change some over time.

Table G.8: Log rate of return (Using only rent-to-own switchers)

	(1)	(2)	(3)	(4)	(5)
Black	0.00981***	0.0101***	0.0108***	0.000541	0.00220
	(0.00144)	(0.00146)	(0.00146)	(0.00146)	(0.00145)
Black × Recession=1	-0.00680*	-0.00900**	-0.00868**	-0.00738*	-0.00677*
	(0.00283)	(0.00286)	(0.00288)	(0.00289)	(0.00290)
Hispanic	0.0234***	0.0145***	0.0144***	0.000866	0.00343*
	(0.00161)	(0.00168)	(0.00168)	(0.00166)	(0.00162)
Hispanic × Recession=1	-0.0331***	-0.00849*	-0.00699*	-0.00370	-0.00445
	(0.00344)	(0.00348)	(0.00349)	(0.00350)	(0.00345)
Observations	742,817	742,817	742,817	742,817	742,817
Recession-by-race FE (fully interacted)					
Dummy for other racial group	Yes	Yes	Yes	Yes	Yes
Year FE (not interacted)	Yes	-	-	-	-
Geography-by-year FE	No	Yes	Yes	Yes	Yes
Demographics (age, #kids)	No	No	Yes	Yes	Yes
Income, education, tenure	No	No	No	Yes	Yes
Housing characteristics	No	No	No	No	Yes

Note: Data from 1974-2021 owner-occupied properties in the AHS. Rent is imputed using our switcher method only using switchers that switch from rented in the prior period to owner occupied in the current period. Standard errors are clustered at the unit level. The excluded race group is White owner-occupants. Geography FEs include the most granular-defined geography variables in the AHS for each year, which change some over time.