Jon Goodman declares that he has done legal work as a consultant for MRI, the company through whom we have accessed the Colorado bid reporting data. This relationship was critical for accessi-ng this novel data set but has not affected how we have conducted the research nor the conclusions that we have drawn. None of the other authors have any other relevant or material financial interested to declare in our submitted manuscript, "Not All Bids Are Equal: Racial Gaps in the Housing Market"

Sincerely,



Christopher Timmins

Corresponding Author

Signed on behalf of Zach Babat, Jon Goodman, and Lu Han.

Not All Bids Are Equal: Racial Gaps in the Housing Market *

Zach Babat
University of Wisconsin–Madison

Jon Goodman
Frascona, Joiner, Goodman and Greenstein P.C.

Lu Han
University of Wisconsin–Madison

Christopher Timmins
University of Wisconsin–Madison

Draft – Do Not Cite or Circulate

March 26, 2026

Abstract

We study racial disparities in housing transactions using novel data that describe the full set of offers submitted for residential properties, including both accepted and rejected bids. Exploiting within-listing comparisons across competing offers, we show that buyers of color bid less on average. However, even when we control for bid amount, rank, buyer agents and detailed offer terms, we find that these buyers are significantly less likely to have their offers accepted. We quantify this disparity in the form of the bid-equivalent penalty faced by buyers of color, which is economically large — these buyers have to bid approximately 2 – 3 percent higher than a white buyer simply to offset their race. Those disparities widen in tighter markets, in listings with more competing bidders and for white sellers – i.e., where sellers have stronger alternative options and where the identity of the marginal buyer is perceived as consequential. Taken together, our results are difficult to reconcile with explanations based solely on buyer preferences, financing constraints, or search costs. Rather, they point to differential acceptance thresholds applied by sellers.

JEL-Codes: R21, R31, J15, J71, D63

*We thank seminar participants at the University of Alabama, University of Pittsburgh, University of North Texas, Universtity of Toronto, University of Wisconsin - Madison, Rice University, Cornell University, Princeton University, Yale University for their helpful comments. We thank Sangyu Qian for excellent research assistance. All errors and omissions are our own.

Babat: *zbabat@wisc.edu*. Goodman: *jon@frascona.com*. Han: *lu.han@wisc.edu*. Timmins: *ctimmins@wisc.edu*.

1 Introduction

Housing markets are central to wealth accumulation and the persistence of racial inequality. A large literature documents that Black and Hispanic households pay more for comparable housing, even after controlling for observable characteristics of buyers, sellers, and neighborhoods. This "tax" eats into the returns from housing ownership, enhancing the racial wealth gap. Yet despite decades of research, the mechanisms underlying these price gaps remain unsettled. Do minority buyers pay more because they value housing differently or face greater search costs? Or is it because sellers treat otherwise identical offers differently?

This question is difficult to answer because most empirical work observes only completed transactions. A literature, beginning with Myers (2004) and continuing through recent work by Ihlanfeldt and Mayock (2009), Bayer et al. (2017), and Box-Couillard and Christensen (2024), documents that Black and Hispanic buyers pay systematically higher prices for similar homes. The dominant interpretation, articulated by Bayer et al. (2017), is that minority buyers face higher search costs. In this view, minority households might search longer or in thinner markets, leading them to pay higher prices for similar homes. When given the chance to bid on a property, these groups will bid more in hopes of not needing re-enter the search process. Importantly, this explanation does not necessarily require direct discriminatory behavior by sellers; rather, it arises from frictions on the buyer side of the market.

In this paper, we bring new evidence from a large-scale dataset of housing bids, observing not only final transactions but *all* submitted offers, including those that were rejected. This allows us to separate two conceptually distinct mechanisms that are confounded in transaction-level data: (1) idiosyncratic taste or search frictions on the buyer side, which affect how much a buyer is willing to bid; and (2) acceptance discrimination on the seller side, which affect whether a given bid is accepted.

This distinction is crucial. If minority buyers pay more because they value certain homes more or face higher search costs, then they should bid more *ex ante*. If instead they pay more because sellers impose higher acceptance thresholds on them, then minority buyers may bid less on average, yet still pay more conditional on winning (keeping in mind that offers are multidimensional and the offer with the highest bid amount is often not the one that wins). These two mechanisms have different implications for theory, welfare, and policy, but they are observationally equivalent in transaction-level data. Our data allow us to disentangle them.

We assemble a novel dataset covering residential real estate transactions in Colorado that links listing information, all submitted offers, bidder characteristics, offer terms, and final transaction outcomes. For each property, we observe the full choice set faced by the seller: who bid, how much they offered, whether the offer included cash or contingencies, which agent was involved, and which bid was ultimately accepted.

This data structure allows us to estimate two complementary objects: (1) reduced-form acceptance regressions on all listings, which relate the probability of winning to bid amount, offer terms, and buyer race; (2) a conditional logit choice model on listings with 2+ bidders, which recover the seller's implicit utility function over competing offers. Crucially, we observe losing bids, which allows us to distinguish seller selection effects from buyers' bidding behavior.

Our identification strategy exploits within-listing comparisons across competing buyers. Because bidders compete for the same property, listed by the same seller at the same time, we difference out property quality, neighborhood characteristics, seller identity, and market conditions by construction. This design directly addresses the unobserved heterogeneity that complicates transaction-based approaches. Moreover, we observe detailed contractual terms — cash share, contingencies, earnest money, escalation clauses — allowing us to condition on dimensions of offer quality typically unobserved in administrative data.

We begin with two distinct stylized facts: First, across otherwise comparable properties, Black and Hispanic buyers submit lower unconditional offers than white buyers. This pattern is robust to controls for property fixed effects, neighborhood characteristics, and financing terms. Second, at the same bid amount, buyers of color are less likely to have their offers accepted. This joint pattern - bid less and face lower acceptance probabilities at the same bid - is difficult to reconcile under a simple interpretation based on buyers' tastes or search-costs. If buyers of color simply had higher idiosyncratic tastes or faced greater search costs, then they should bid more aggressively *ex ante*. Yet the data show the opposite, pointing instead to a mechanism operating at the acceptance stage.

To quantify this mechanism, we estimate a conditional logit model in which sellers choose among competing offers for a given listing. The model recovers the implicit utility weights sellers place on buyer race, bid amount, cash share, earnest money, and bid characteristics including whether buyers waive inspections, have a sale contingency, require an appraisal gap, or have an escalation clause. The results are striking.

We find that the implied bid-equivalent penalty faced by minority buyers is economically large — approximately 2-3 percent of the bid amount — and that it persists even among all-cash offers. Because identification comes from within-listing comparisons and conditions on rich contract terms, these premia cannot be attributed to housing quality, seller heterogeneity, or financing risk. Rather, they reflect systematic differences in the acceptance standards sellers apply across buyers of different races when evaluating otherwise comparable offers.

The magnitude of these penalties varies systematically with market structure. Disparities widen in tighter and faster-turnover markets and in listings with more competing bidders. In search environments where buyer arrival rates are high and the cost of rejecting an offer is low, sellers enjoy stronger outside options and greater discretion. Consistent with Becker (1957), discriminatory behavior intensifies when

its economic cost declines. In addition, we also find that white sellers require a significantly higher racial premium than sellers of color.

Taken together, our findings show that racial disparities in housing operate primarily through endogenous seller acceptance thresholds that vary with market power and neighborhood dynamics. Inequality in housing markets is therefore rooted not only in differences in bidding behavior, but in how market competition governs the exercise of seller discretion.

There is a large literature on discrimination in housing markets to which this paper contributes.[1] First, we provide direct evidence that racial disparities in housing operate primarily through seller-side acceptance behavior rather than solely through buyer-side bidding differences. By observing losing bids, we disentangle mechanisms that are observationally equivalent in transaction data. Second, we quantify discrimination in economically interpretable terms. The conditional logit framework expresses racial disparities as price-equivalent trade-offs in seller utility, allowing us to measure the implicit premium required to offset buyer race. Third, we show that racial disparities are not spatially uniform. They concentrate in transitional neighborhoods - areas with improving amenities, rising demand, and demographic flux — where the identity of the marginal buyer is perceived as consequential. Discrimination thus shapes not only who transacts, but also where, influencing access to neighborhoods experiencing upward transition.

Our findings have important implications for housing policy. Interventions aimed solely at reducing search costs, such as improving information or expanding listings, will not eliminate price disparities if sellers continue to apply differential acceptance standards. Policies must instead address discrimination at the point of offer evaluation, where market power and discretion are greatest.

2 Data

2.1 MRI Data

This paper uses a novel dataset comprised of submitted bids for real estate transactions. These data are collected by MRI Software, a multinational company offering a wide array of software products related to real estate investment and management. MRI owns a system called CTM eContracts ("CTMe"), a platform for sending and storing real estate contracts. CTMe has grown into a near monopoly across the Denver

[1] Work has documented discriminatory behavior across multiple stages of the housing process, including search, access, finance, and transaction outcomes (Yinger (1995); Galster and Godfrey (2005); Dymski (2006); Hanson and Hawley (2011); Ambrose et al. (2021); Christensen et al. (2022); Bartlett et al., 2022; Frame et al., 2023.) Much of the empirical evidence on housing discrimination derives from the U.S. Department of Housing and Urban Development's Housing Discrimination Studies (HDS), which have been used to measure the presence of and mechanisms underlying discriminatory behavior (e.g., animus-based discrimination versus statistical discrimination). (Ondrich et al. (1998); Ondrich et al. (2003); Zhao et al. (2006); Dymski (2006); Guryan and Charles (2013)). Christensen and Timmins (2022) use the 2012 HDS to measure the welfare implications of discriminatory steering in owner-occupied housing search, finding that properties shown to testers of color differed in many important neighborhood characteristics including pollution exposure, crime rates, poverty levels, and local human capital–relative to those shown to white testers.

metro area (and much of the remainder of Colorado) as the platform through which real estate brokers send and sign real estate contracts: CTMe is the official contract software for the Colorado Association of Realtors; in three separate conversations with industry leaders at large brokerage firms in the Denver metro area, estimates of the share of deals done through CTMe ranged from 95% to 100%.

Some types of real estate transactions tend not to use real estate brokers, including but not limited to non-arm's-length transactions (e.g. conveyance of property to family members, foreclosure sales, and conveyances for estate planning purposes). In these cases, it is more common for lawyers, rather than real estate brokers, to handle the contracting and transaction process. However, for the broad segment of the Denver metro real estate market where homeowners wish to sell their homes and are seeking a buyer on the open market, these transactions tend to use real estate brokers, who overwhelmingly use the CTMe platform.

The Colorado Division of Real Estate requires that real estate brokers use certain template forms for submitting offers, meaning that contracts submitted through the CTMe platform all follow a standardized format. MRI maintains a dataset where observations correspond to submitted offers (sometimes referred to as "bids" or "contracts" below), and columns in this dataset correspond to the contents of the fields filled out on each contract. The number of filled fields varies with the contract structure, but can range into the hundreds. This dataset serves as our source for submitted home purchase offers.

2.2 Data Assembly

We construct a panel of listings, comprised of CoreLogic data on residential properties in the Denver MSA, CoreLogic data on the history of transactions on these properties, and MRI data on bids submitted through the CTMe platform as homeowners sold these properties. In our ideal dataset, we observe every offer that owner-occupied homeowners received between purchasing and selling their homes, observe which offer was chosen and resulted in a closed deal, and link together every cycle of receiving offers/completing transactions on the same property.

To construct this dataset, we begin by linking together CoreLogic Property Basic data and CoreLogic Owner Transfer data for the state of Colorado. These datasets contain CoreLogic's proprietary Composite Property Linkage Key that enables us to identify individual properties, and link a given property to its entire history of sales.

We then link the offers we observe in the CTMe data to each property, using the address information reported on each bid. First, we match offers to properties by executing a simple merge, linking together an offer from CTMe with a property in CoreLogic if the following seven fields match exactly: county, five-digit

zip code, street number, street direction, street name, street type, and unit number (if applicable). (Example: if an offer in CTMe and a property in CoreLogic both contain the address "Denver County, 80203, 200 E Colfax Ave.," we form a match.)

We notice that a significant number of offers contain apparent typos written in the address fields. In order to match these offers to CoreLogic, we execute a geocoding procedure. We take the bids that do not find a match in the first round, and use their address data to query the Google Geocoding API. The API returns latitude and longitude coordinates for the address, typo-corrected address information, and several fields describing the confidence of the geocoder (e.g. a variable stating whether the query is a partial match for multiple addresses, a variable stating whether the geocoder is returning a precise "rooftop" address or an approximate location). We then perform a second merge round, in which we link bids to properties if their latitude and longitude coordinates match to three decimals of precision, as well as their street numbers and unit number.

Out of 3,885,440 observations in the original bid data, we drop 106,166 bids that likely refer to multiple properties or contain invalid address information (e.g. terms like "TBD" or "N/A" or "Multiple"). Of the 3,779,274 bids that have valid address information, 2,986,908 match to properties in CoreLogic during the first round comparing seven address fields. 289,645 properties merge during the second round, using latitude and longitude coordinates.[2] We ultimately match 86% of the bids that contain usable address information to property records in CoreLogic.

The procedure above ensures that offers are properly matched to the correct address, but we must also take care to group together offers along the timeline and identify which offers were in competition with one another. To do this, we use the CoreLogic transactions data to identify the dates on which properties are sold, and construct windows between sale dates. We say that all offers are in competition with one another if they fall within the same window. When we use the term "listing" below, it refers to this entire window between when someone buys the property and the subsequent sale.

2.3 Sample Construction

As described above, MRI has a virtual monopoly in the Denver metro area as a platform for real estate brokers to create and deliver offer documents. However, certain types of transactions likely do not involve real estate brokers (see Section 2.1). To construct a universe of home sales that excludes transactions MRI is unlikely to observe, we restrict the set of transactions we consider according to the following parameters:

- Properties must be located in the five most populous counties in the Denver metro area: Denver,

[2] 40% of properties have the same coordinates in CoreLogic and from the geocoder, to four decimal places. We flag 6.34% of matches where $\sqrt{(lat_1 - lat_2)^2 + (lng_1 - lng_2)^2} > .0003$, and manually audit a subsample to confirm that these are valid matches.

Arapahoe, Jefferson, Adams, Douglas.

- CoreLogic indicators for the following sale attributes must be switched off: sale is non-arm's-length or between potentially related buyer and seller, sale is of a foreclosed/REO property, sale is multi-parcel, property is new construction, property is flagged as non-residential.

- If a property sells twice within a 365-day period, we drop both sales. These sales likely correspond to either home-flipping, or are "paper sales" that merely capture the transfer of the home between (e.g.) a trust and a related interest.

In addition, we restrict to sales of single-family homes (i.e., we exclude properties where CoreLogic reports the land is used for condos, duplexes, townhomes, or mobile homes). We also restrict to transactions where the sale price was between $50,000 and $2,500,000.

We identify 293,132 home sales between 2014-2024 that meet all of the parameters described above. We successfully link bids to 240,753 (82.13%) of these home sales. For 131,116 (44.73%) of home sales, we observe at least one bid that the CTMe platform indicates was accepted by the home sellers. We perform a back-of-the-envelope exercise to compute a maximum theoretical match rate of 59.5%. Our merge therefore links at least one accepted bid to a home sale at a rate of 75% of the possible maximum.[3] This suggests that buyers have the ability to negotiate sale terms away from the formal offer submission platform. Despite this, we will show that among the home sales where we observe at least one accepted bid, the bid that we identify as the winner tends to have characteristics that match the publicly known sale details from CoreLogic.

In the future, we intend to improve upon this 45% coverage rate in at least two ways. First, we have identified avenues for continuing improve the merge that links bids to property sales. Second, we intend to make better use of MLS data to identify the ways in which different listings were advertised – we believe that home sales that were not listed on an MLS are less likely to be home sales where bids were sent through the CTMe platform, and we can exclude these listings from the universe of home sales we consider.

[3] We compute the maximum possible match rate as follows: *i)* There are 336,403 accepted bids in the Denver MSA, excluding bids with invalid address information or which refer to multiple properties; *ii)* there are 565,114 home sales that could plausibly be mediated by brokers (arm's length sales that not foreclosures, REO, new construction, multi- or split-parcel, without restrictions to single-family, by price, or by property-type). If we were to perfectly match every accepted bid in the Denver metro to a sale, and assume only one accepted bid per sale, the maximum possible match rate is thus 59.5%.

2.4 Key Variables from the CTMe Platform

2.4.1 Basic offer and bidder information

Identifiers: Observations contain unique numeric identifiers to distinguish contracts. Observations also contain numeric identifiers for buyer's and seller's agents, and a series of fields identifying the address of the property related to the bid. It is exceedingly rare for buyer's agents to represent multiple prospective buyers competing to purchase the same home, as this represents a conflict of interest. As such, within a given listing, we assume that a buyer's agent represents a single prospective buyer, enabling us to identify repeat offers by the same person. This method does not allow us to identify offers submitted to multiple different properties by the same person.

Offer amount variables: When homebuyers describe an offer's purchase amount, they complete a section where they write the overall total and fill out several boxes describing how they will finance this total. Simplifying this contract section for the sake of exposition, we identify four key variables: the total purchase amount, the amount of earnest money (a deposit enclosed with the offer), the amount of money secured from a loan, and cash to be delivered at closing.

When we refer below to the share of an offer that is financed by "cash", we are referring to the percentage of the purchase amount that is covered by the earnest money and the cash to be delivered at closing. We identify "all-cash offers" as offers that are 100% cash.

When we refer below to measures of loan-to-value (LTV), we simply compute the ratio of the loan amount written on a bid to the total purchase amount.

2.4.2 Bidder race

We do not observe the bidder's name directly, and bidders do not directly report their race. Upon request, MRI processed their dataset through an algorithm that uses a bidder's name and home address to impute their race. Although self-reported race is often useful for validating the accuracy of imputed race, in our context, it is reasonable to assume that a prospective buyer's race is likely being communicated via the buyer's name and their home address – thus, such an algorithm captures the most relevant information that sellers observe and possibly act upon.

With each contract, we observe five variables, which report the algorithm's suggested probability that the bidder is Asian, Black, Hispanic, White, or some other race. In Appendix A.1, we perform an exercise with HMDA data to validate the imputed race measures from MRI. We estimate a multinomial logit model, in which the vector of race probabilities from MRI predict self-reported race in HMDA data. We estimate that if MRI reports a bidder has a 100% probability of being Asian, HMDA classifies that person as Asian

82% of the time. Corresponding figures for other race groups are: Black, 79%; Hispanic, 85%; White, 96%.

When including race as an explanatory variable in both reduced-form and structural estimations, we do not use the single-most likely race for a bidder, but rather their vector of predicted race probabilities (see Equation 1). This smooths out the noise in race classifications: while some bidders are assigned to a single race category with high confidence, others have moderate probabilities of being multiple races. Using the full vector of probabilities preserves the underlying uncertainty of the demographic assignment, protecting our estimates from the measurement error inherently introduced by arbitrary, discrete categorization. (Appendix A.1 also shows how using the single-most likely category performs significantly worse in validation exercises.)

For some descriptive statistics, we use the single-most likely race to describe bidders. If no single race is most likely, we assign bidders to the "other" category.

2.4.3 Important legal clauses

We make use of the contents of a free-text box, titled "additional provisions," to identify when contracts contain certain terms that significantly affect the liquidity of an offer. We describe what these clauses are, how we identify them, and their frequency in our cleaned panel of listings.

Escalation clauses: Alicia has submitted an offer to purchase a home from Bob at a certain amount. However, if Bob were to receive another offer for a higher amount, Alicia is willing to increase her offer by a stipulated increment, and is willing to continue increasing her offer up to a specified maximum possible bid amount. Usually, Alicia would include a requirement that for Bob to trigger the escalation clause, he must provide proof of the legitimacy of a competing offer, as well as proof that Bob did not tamper with the homebuying process by soliciting a non-arm's-length bid. As such, there is some ambiguity about whether sellers view escalation clauses favorably: they can ultimately fetch a higher sales price, but they can also invite liquidity risk in the form of pre-closing arguments about whether the escalation clause should actually go into effect. We say an offer has an escalation clause if the string "escalat" appears in the additional provisions sections of offers. 15,005 bids have these clauses (3.83%).

Waive-inspection clauses: Alicia has submitted an offer to purchase a home from Bob, and is agreeing to forego certain forms of due diligence that Bob might usually expect a homebuyer to require before closing a deal. Normally, inspections can be grounds for canceling an offer and requiring Bob to release any earnest money a prospective buyer has put forth. Thus, if Alicia waives her right to inspect the home, her offer presents less liquidity risk for Bob. We say an offer has a waive-inspection clause if any of the following terms appear in the additional provisions sections of offers: "waive inspection," "waives inspection,"

"waiving inspection," "inspection is waived," "waive all inspections," and various spellings of the phrase "as-is" ("Alicia agrees to purchase the home as-is"). 57,690 bids have these clauses (14.74%).

Cover-appraisal-gap clauses: Alicia submits an offer to purchase a home from Bob, and proposes to secure a mortgage to finance her purchase. The bank's appraisal of the home will determine how much financing they are willing to offer Alicia. In the event that the bank appraises Bob's home at a lower value than anticipated, Alicia is guaranteeing that she can provide additional cash (up to a specified limit) to ensure that she can still purchase the house at the originally specified price. Thus, Alicia's offer to cover such an appraisal gap presents less liquidity risk for Bob. We say an offer has an appraisal gap clause if the phrases "appraisal gap" or "appraised shortfall" appear in the additional provisions sections of offers. 15,059 bids have these clauses (3.85%).

Home sale contingencies: Alicia is in the process of selling her own home. In the meantime, she makes an offer to Bob, but stipulates that her offer is contingent upon her successfully selling her own home. Thus, Alicia's offer presents significant liquidity risk for Bob: if he accepts her offer, he is foregoing the opportunity to work with other potential buyers in the expectation that Alicia is able to successfully complete her own home sale. We say an offer has a home sale contingency by searching the additional provisions sections of offers using the following regex search keys: "contingent.*sale", "sale of.*property," and "subject to.*sale". 9,926 bids have these clauses (2.54%).

2.5 Identifying Winning Bidders and Collapsing Repeat-Bidders

The CTMe dataset comes with a set of binary variables indicating the number of people who need to sign a contract for it to be valid, and how many people signed the contract. MRI considers an offer to have been accepted if the two numbers are equal. Note that an accepted offer is different from a closed transaction: a crucial aspect of this study context is that even after an offer is accepted, the buyer and seller must often navigate a gauntlet of additional difficulties (e.g., completing inspections, securing mortgage financing, additional bargaining) before a deal closes and a transaction occurs. As such, we occasionally see multiple accepted offers within a single listing.[4]

However, only one bid can actually be the bid that secured the purchase of a home. When the CTMe data indicates only one bid was accepted before a home sold, we identify that bid as the winner. If there are multiple accepted bids, we use the following tiebreakers (in order, until only one bid remains): smallest difference between the submitted offer price and sale price reported in CoreLogic; bid submitted on a later date; bid with the higher numeric contract identifier (these identifiers appear to be assigned in the order

[4] 83.63% of listings have a single accepted bid; 13.46% of listings have 2 accepted bids; 2.32% of listings have 3 accepted bids; the remaining 0.59% have 4 or more accepted bids up to a maximum of 11.

that bids are received).

Finally, when a bidder submits multiple offers on the same listing, we collapse these repeat offers into a single observation representing their best offer. If the bidder submitted the winning bid, we select that bid as the bidder's best offer. If the bidder did not submit the winning bid, we use the following hierarchy of tiebreakers to select a bid: accepted offer, highest offer, latest offer, highest numeric contract identifier.

2.6 Cleaning Listing Sample

We are concerned with estimating how various characteristics of a bid affect the probability that it wins the listing. We perform the following cleaning steps to ensure that every offer competing for a given listing in our estimation sample has reasonably comparable characteristics and high data quality (i.e., was not corrupted by typos, information written in the wrong part of a contract form, etc.). If any offer submitted for a listing meets any of the following criteria, drop that listing: non-numeric bid amount and accepted bid; total purchase price outside of $[\$50\,000, \$2\,750\,000]$; amount of earnest money or cash at closing was non-missing and non-numeric; sum of earnest money and cash at closing, as a share of total bid amount, falls outside the range [-.05, 1.1]; bid contains negative earnest money; earnest money as a share of total purchase price is above the 99.9th percentile. These steps drop 18,153 offers submitted across 1,741 listings from our sample our 654,180 offers submitted across 131,116 listings.

2.7 Additional sale data

CTMe is a unique dataset for information about bidders who failed to win a listing. However, we also make use of several other sources for information pertaining to winning bidders and the home sale process.

CoreLogic MLS: We merge CoreLogic MLS data to transactions using CoreLogic's proprietary property linkage keys and information about listing timing. We use this data to compute a measure of time-to-close, defined as the difference between contract date (when buyer and seller agree to a contract in the MLS data), and sale date (from the publicly available deeds data). In the descriptive statistics section below, we also make use of the MLS data to compute measures of market tightness.

CoreLogic Mortgage and HMDA: We merge transactions to CoreLogic mortgage data, then use property and loan characteristics to merge on loan-level HMDA data for the years 2000-2021. This provides us with an external source for homebuyer race, and we use the buyer race for a given sale on a property to construct seller race for the subsequent sale on a property (as in Bayer et al. (2017)). HMDA also provides us with a measure of buyer income, which we use in supplemental analyses.

2.8 Neighborhood Covariates

In many of the results presented below, we split our dataset of listings into subsamples depending on the characteristics of the neighborhoods in which a given home resides. We are interested in how bid-equivalent penalties vary based on the following neighborhood characteristics: population shares of non-Hispanic white people, educational attainment at or above a Bachelor's degree, poverty rate, and air quality.

From the US Census, we use the ACS5 for time-varying ZCTA-level data on racial demography, educational attainment, and poverty rate. We match homes to ZCTAs using zip code. A listing is matched to the ACS5 data published two years after a home's sale date (so a home sold in 2020 is matched to the 2022 ACS5 estimates, which pool data from 2018-2022). Homes sold in 2023 and 2024 are all matched to the 2024 ACS5, which is the latest available data.

To measure air quality, we use estimates of the concentration of fine particulate matter ($PM_{2.5}$) from van Donkelaar et al. (2021). We use this for yearly measures of air quality within a zip code; homes sold in 2018-2024 are all matched to the 2018 values of air pollution within the zip code.

3 Descriptive Statistics

In this section, we present summary statistics pertaining to the characteristics of bids and sold homes within our cleaned panel of 129,369 listings. Much of our estimation below will restrict to listings with at least two competing bidders in order to enable identification of listing fixed-effects; these statistics make no such restriction (i.e., a listing could involve just a single bidder).

We also present statistics pertaining to a novel measure of real estate market tightness, which we are uniquely positioned to offer due to the nature of our dataset.

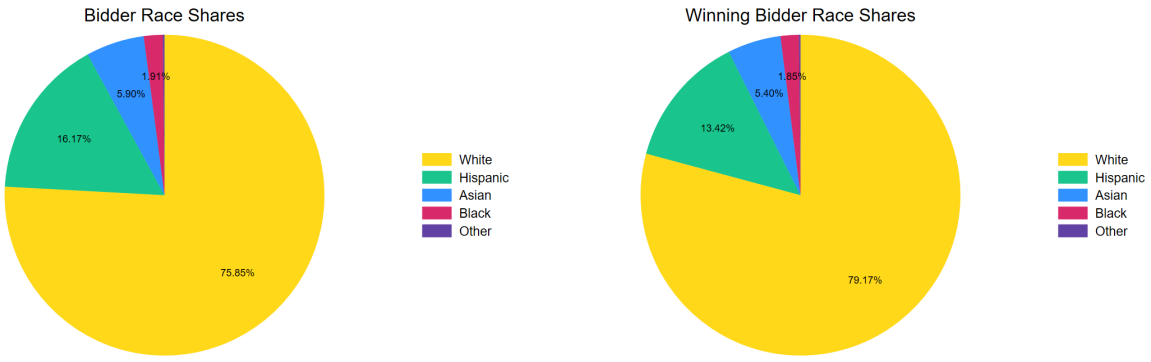
3.1 Statistics Describing Listings and Bidding Behavior

Table 1 decomposes our sample of bids and listings, and Table 2 provides demographic information about bidders and sellers.

Panel A shows the distribution of the total bid amount associated with an offer, as well as statistics about how these bids are financed. In particular, we are interested in "earnest money," which is equivalent to a deposit submitted with bids, and "cash at closing," referring to the portion of the bid that the bidder could provide outside of receiving a loan from a bank. We refer to the sum of these two forms of payment as "total cash." Bids that contain a higher share of total cash are bids that present less liquidity risk: the greater the share of a bid that a bidder can provide independently of a bank, the less risk there is that an

offer falls through because of an inability to secure mortgage financing. Note that the distribution of total cash share increases dramatically between the 75th and 90th percentiles, from nearly 30% to 100%. We also tabulate loan-to-value as the ratio of the loan amount written on a bid to the total purchase amount.[5]

Panel B shows the distribution of listing-wide metrics. The distribution of sale price (according to CoreLogic) is greater than the distribution of bid amount at every percentile shown, which is consistent with the idea that a winning offer likely contains the highest bid. The second row shows the difference between the total amount of the bid that we identify as the winner, and the corresponding sale price recorded in CoreLogic, which is a check on the quality of our bid data and our procedure for identifying the winning bidder. Across much of this distribution, the difference is zero. The difference is negative below the 10th percentile, indicating a sale price greater than what we identify as the winning bidder, and positive at the 90th percentile, reflecting the opposite – we cannot determine the extent to which these differences are due to bargaining outside of the CTMe platform versus error in our ability to identify the winning bidder. The third row shows the distribution of bidders who submit a bid before a sale. 99% of sales see 16 or fewer bidders, and the maximum number of prospective buyers in our data is 69. The final row reports the distribution of time-to-close, as defined in Section 2.7.



(a) Bidders (b) Winners

Figure 1: Racial Composition Among Bidders

Panel C shows the frequency of various offer characteristics. The top row shows that 15.56% of offers are "all-cash," meaning the total cash value is equal to 100% of the offer amount. The next four rows show the frequency of the four different types of contract provisions discussed in section 2.4.3, which can modulate how attractive an offer is to a seller without changing the offer price.

Table 2 provides demographic information about bidders, winners, and sellers. Panels D, E, and F

[5]In a simplified world, bids are only financed with earnest money, cash at closing, and loans, so bids at the 10th percentile of total cash share would fall at the 90th percentile of loan-to-value, and these fractions would sum to 1 (and the same for p25 and p75 and so on). In reality, due to more exotic forms of payment and occasional error when filling in these forms, some of these fractions are different from 1 by < 0.005.

Table 1: Bid and Listing Statistics

			p10	p25	p50	p75	p90	N
Panel A: Bid-level distributions								
	Offer Amount		259000	330000	435000	585000	785000	391369
	Cash at Closing Share		0.0187	0.0332	0.1548	0.2797	0.9848	391369
	Earnest Money Share		0.0081	0.0098	0.0113	0.0150	0.0208	391369
	Total Cash Share		0.0300	0.0473	0.1733	0.2982	1.0000	391369
	Loan-to-Value		0.0000	0.7000	0.8226	0.9515	0.9700	391369
Panel B: Listing-level distributions								
	Sale price (CoreLogic)		280000	355000	465000	629000	850000	129369
	(Winning offer − sale price)		0	0	0	0	5000	129369
	Number of bidders		1	1	2	4	6	129369
	Time to Close		19	26	31	37	45	89669
						Count	Share	N
Panel C: Bid-level counts (offer characteristics)								
	All-cash-offer	*Avoids financing risk*				60894	0.1556	391369
	Escalation clause	*Offer can increase with competition*				15005	0.0383	391369
	Waive inspection	*Reduces closing risk*				57690	0.1474	391369
	Appraisal gap	*Reduces financing risk*				15059	0.0385	391369
	Sale contingency	*Increases closing risk*				9926	0.0254	391369

Notes: Total cash share is computed as *(earnest money + cash at closing)/total offer price.* All-cash offers are offers where total cash share equals 100% of bid price. Loan-to-value ratio is computed as the ratio of the loan amount written on a bid to its total offer price. Details regarding contract provisions can be found in Section 2.4.2. Sample is not restricted by number of bidders. Time to close requires a merge with MLS data to compute, and is missing for some listings.

Table 2: Demographic Statistics

	Count	Share	N
Panel D: All bidders, by race			
Asian	23083	0.0590	391369
Black	7463	0.0191	391369
Hispanic	63281	0.1617	391369
Other	684	0.0017	391369
White	296858	0.7585	391369
Panel E: All-cash bidders, by race			
Asian	4350	0.0714	60894
Black	780	0.0128	60894
Hispanic	6254	0.1027	60894
Other	139	0.0023	60894
White	49371	0.8108	60894
Panel F: Winning bidders, by race			
Asian	6983	0.0540	129369
Black	2393	0.0185	129369
Hispanic	17366	0.1342	129369
Other	201	0.0016	129369
White	102426	0.7917	129369
Panel G: Seller race from HMDA			
Asian	2161	0.0363	59529
Black	1350	0.0227	59529
Hispanic	5371	0.0902	59529
Other	359	0.0060	59529
White	44535	0.7481	59529

Notes: Panels D-F show demographic decomposition among various subsamples, using most-likely race from CTMe data. Panel G shows seller race, as inferred from HMDA data. Seller race requires a merge with HMDA data to observe, and is missing for some listings.

classify offers into a single most likely race category according to the method described in 2.4.1. These panels show the frequency of offers submitted by bidders of each race, all-cash offers submitted by each race, and listing winners of each race. White bidders submit 75.85% of offers and comprise 79.17% of winners. Hispanic bidders submit 16.17% of offers and comprise 13.42% of winners. Asian bidders submit 5.90% of offers and comprise 5.40% of winners. Black bidders submit 1.90% of offers and comprise 1.85% of winners. The subsample of all-cash bidders (Panel E) contains a greater share of Asian and white bidders and a smaller share of Black and Hispanic bidders, relative to the overall sample of bidders (Panel D). In Panel G, we report the racial demography of sellers for the subset of listings where we can identify seller race from HMDA, according to the procedure described in Section 2.7.

3.2 Market Characteristics

Table 3 illustrates the different market conditions in which our sample of listings occur. Each observation in the distributions shown corresponds to a single listing. Note that metrics using MLS data contain slightly fewer observations because our MLS data for the Denver MSA runs out in the middle of 2022, and coverage can be thin in less-populated ZIP codes.

The first three rows depict statistics that are unique contributions of this project through the novelty of our bidding dataset. The first row describes the number of bidders competing for each listing, while the second describes the total number of offers submitted for each listing. More than half of our sample consists of listings where multiple bidders competed against one another, and more than 75% of listings saw more than one submitted offer. The third row computes the number of offers submitted in a given year in a three-digit ZIP code (e.g. "802" for ZIP code 80203) divided by the number of distinct properties that appeared in MLS data for the same ZIP3-year. This distribution is right-skewed — going from the 10th to 50th percentiles is a difference of .7 offers per listing, while going from the 50th to 90th percentiles is a difference of 1.23 offers per listing.

The next five rows depict more widely available market statistics. ZIP-Quarter TOM computes time-on-market (days) for listings appearing in a ZIP-code quarter.[6] ZIP white share, college rate, and poverty rate all come from the ACS5.[7] $PM_{2.5}$ is a fine particulate matter air pollution measure described in Section 2.8.

[6] We perform a routine that groups together duplicate listings on a single property if they occur within 90 days of each other. If a property lists on Jan. 1, is taken off the MLS service on Jan. 20, re-lists again on Mar. 1, and sells on May 1, the time-on-market for this listing would be 120.

[7] We group together individuals with a Bachelor's degree or graduate/professional degrees together to compute those with a college degree.

Table 3: Listing-level distribution of market characteristics

	p10	p25	p50	p75	p90	N
Bidder count	1	1	2	4	6	129369
Offer count	1	2	3	6	10	129369
ZIP3-Year Mkt Tightness	1.97	2.08	2.25	2.45	2.77	129349
ZIP-Quarter TOM	43.31	51.99	62.05	76.76	92.82	109745
ZIP White Share	35.38	53.92	69.22	80.27	84.60	129369
ZIP College Rate	21.66	31.64	48.47	59.44	66.35	129369
ZIP Poverty Rate	2.96	3.90	6.83	10.18	14.67	129369
$PM_{2.5}$	4.44	4.62	4.98	5.63	7.25	129369

Notes: Observations are individual listings. Bidder count is the number of distinct individuals who submit a bid (see section 2.4.1). Offer count is the number of distinct offers submitted by all bidders. ZIP3-Year Market Tightness is the count of bids submitted in a 3-digit ZIP code within a given year, divided by the number of properties that appear in MLS data in the same ZIP3-year. ZIP-Quarter TOM is the average time-on-market (days) for listings that appear in MLS data in a five-digit ZIP code in a given quarter. Rows with MLS data have fewer observations because our MLS data for the Denver MSA runs out in the middle of 2022. ZIP White Share, College Rate, and Poverty Rate all come from the ACS5 (see section 2.8). $PM_{2.5}$ is a measure of fine particulate matter air pollution (see section 2.8).

4 Empirical Analysis

4.1 Empirical Framework

Our empirical analysis exploits a novel feature of the data — for each listing, we observe the full set of competing offers submitted, along with detailed characteristics of each offer, and the identity of the winning bidder. This structure allows us to examine racial disparities at the precise stage where seller discretion is exercised — offer evaluation — rather than inferring seller behavior indirectly from transaction prices alone.

We organize the analysis around three closely related empirical objects — bid, accepted bid, and acceptance probabilities. All three are estimated using the offer-level data and rely on within-listing comparison across competing offers. By comparing bidders competing for the same property at the same time, this approach differences out fixed housing quality, seller identity, neighborhood attributes, and local market conditions by construction. This allows us to improve upon existing empirical strategies in two fundamental ways.

First, a longstanding challenge in the housing discrimination literature is unobserved heterogeneity in housing quality, neighborhood amenities, and seller characteristics. Existing studies typically address this issue using repeat-sales designs or high-dimensional hedonic controls applied to transaction prices (e.g. Bayer et al. (2017), Box-Couillard and Christensen (2024)). While powerful, these approaches require strong assumptions — e.g., that property quality remains constant over time and can be differenced out with repeat sales, or that different sellers exhibit similar bargaining skills and level of patience in repeated transactions. Changing neighborhood attributes need to be adequately controlled for with, for example,

tract-by-year fixed effects. These assumptions can be difficult to verify empirically. Results may be biased if houses undergo unobserved renovation or deterioration, or if local housing trends occur at a finer level than the census tract.

Our setting allows for a more direct comparison. Because all offers associated with a listing pertain to the same listing, identification comes from differences across buyers competing for the same property listed by the same seller under identical market conditions. This design differences out all fixed attributes of the property, seller, and local market environment, substantially reducing concerns about unobserved heterogeneity.

Second, transaction-based approaches face difficulty disentangling racial disparities from differences in non-price dimensions of offer quality, such as financing reliability, liquidity, and contractual contingencies. Existing studies often rely on borrower characteristics - such as income or credit scores - to proxy for these factors. While informative, these measures may not fully capture variation in financing flexibility or the ability to waive contingencies. To the extent that such offer attributes are correlated with race but imperfectly observed, transaction-based estimates may conflate differential treatment with unobserved differences in offer quality.

Our data allow us to address this concern directly. We observe detailed offer terms, including cash share, inspection and appraisal contingencies, escalation clauses, earnest money, and closing-time flexibility — features that sellers commonly associate with transaction risk and deal certainty. Controlling for these dimensions enables a cleaner separation of racial disparities from differences in offer quality, allowing us to more precisely isolate seller behavior.

In the sections that follow, we present reduced-form evidence on how offer prices and acceptance probabilities vary with buyer race. First, we show that bidders of color tend to bid less than white bidders when competing for the same listing, and document how offer price trades off with various offer terms. Then, we provide reduced-form evidence that bidders of color are less likely to win a listing using a linear probability model that controls for offer amount, offer terms, and buyer's agents. We supplement these results with additional reduced-form evidence to reinforce our argument that latent differences in bidder quality or taste are not the central mechanism driving our results. We defer most of our discussion of racial premia to Section 5, when we estimate a logit choice model that quantifies the acceptance thresholds implicit in seller decisions.

4.2 Initial Offer Price

We begin with a set of regressions that characterize how offer prices vary with with the bidder's race and other attributes that might make a bid differentially attractive to sellers (e.g. all-cash offer status). To allow for the identification of listing fixed effects, we restrict the sample to listings with at least 2 bidders. We also winsorize the sample by excluding 652 listings with more than 19 competing bidders, which represents 0.5% of all listings and 0.7% of all multi-bidder listings. These restrictions trim our sample down to 328,794 bids submitted across 82,702 listings.

We first estimate how bids vary by bidder race using regressions of the form:

$$\ln(Bid_{ij}) = \sum_{r=1}^{4} \alpha_r \mathbb{1}\{Race_i = r\} + X'_{ij}\gamma + \theta_j + \varepsilon_{ij}, \tag{1}$$

where $\ln(\text{Bid}_{ij})$ denotes the offer price submitted by bidder i for listing j; $\mathbb{1}\{Race_i = r\}$ is a set of predicted race probabilities, Black, Hispanic, Asian, and Other, with White as the omitted category;[8] X_{ij} includes observed offer characteristics including all-cash payment, waiving inspection contingencies, escalation clauses, financing terms, appraisal-gap coverage, and sale contingencies; θ_j denotes a listing fixed effect.

Table 4 reports the estimates from Equation 1. Column (1) compares bids submitted for the same property listed by the same seller at the same time. Relative to white bidders, Asian bidders submit offers that are 0.83 percent lower, Black bidders 2.66 percent lower, and Hispanic bidders 0.56 percent lower.

Column (2) augments the specification with a rich set of offer-level controls, including indicators for all-cash payment, escalation clauses, inspection waivers, appraisal-gap coverage, and sale contingencies. We also add an indicator variable to control for offers with an LTV of 80 or above, indicating greater financing risk. These controls allow us to compare bids not only within property within listing, but across offers with comparable contractual structure, thereby isolating price differences net of financing and risk concerns.

These contractual terms are important determinants of offer price. All-cash offers are associated with offer prices that are approximately 2 percent lower, consistent with improved liquidity substituting for higher price (Han & Hong, 2024). Escalation clauses are associated with lower written offer prices, reflecting that the posted amount represents a floor designed to adjust upward in response to competing bids. Inspection waivers are likewise associated with lower offer prices – much like all-cash offers, these types of clauses can substitute improved closing liquidity for higher offer amounts. By contrast, appraisal-gap coverage is associated with higher offer amounts, suggesting that bidders combine price increases with liquidity assurances

[8] Probabilities enter into the regression as fractions. In the tables below, one can easily interpret the coefficient on $P(Asian)$ as the effect of switching a buyer's race from $P(White) = 100\%$ to $P(Asian) = 100\%$.

when competition is strong. Sale contingencies are also associated with higher offer amounts, consistent with bidders compensating sellers for additional closing risk.

Importantly, differences in financing terms and offer structure explain only a modest share of the observed bid disparities, particularly for Black buyers. Column (2) shows that Asian bidders submit offers that are 0.46 percent lower, Black bidders 2.32 percent lower, and Hispanic bidders 0.37 percent lower than comparable white bidders. Even conditional on detailed contract characteristics and within-listing comparisons, bidders of color, and especially Black bidders, submit lower offer prices for the same home.

4.3 Winning Probability

We now turn to the question of the seller's decision regarding which offers to accept and close. We estimate the following equation:

$$\mathbb{1}(Win_{ij}) = \sum_{r=1}^{4} \alpha_r \mathbb{1}\{Race_i = r\} + \beta \ln(Bid_{ij}) + X'_{ij}\gamma + \theta_j + \varepsilon_{ij}. \tag{2}$$

Table 5 reports estimates from Equation 2, where the dependent variable indicates whether an offer ultimately wins the listing. As before, the specification includes listing fixed effects and detailed offer characteristics, so identification comes from within-listing comparisons across competing offers with similar terms. Property characteristics, seller identity, neighborhood attributes, and local market conditions are differenced out by construction. The coefficients capture how winning probabilities vary across bidders competing for the same property listed at the same time. In the second column, we extend these results by adding fixed effects for the bidder's agent to explore alternative explanations for our main findings.

The coefficient on $\ln(Bid_{ij})$ is positive and precisely estimated, indicating that, else equal, a 1 percent increase in the bid raises the probability of winning by approximately 1.4 percentage points. Thus offer price remains the dominant margin of competition within listings.

Conditional on bid level and contractual terms, however, systematic racial disparities emerge. Relative to white bidders, Asian, Black, and Hispanic bidders are 8.9 percentage points, 8.9 percentage points, and 8.5 percentage points less likely to have their offers accepted, respectively. In other words, even when submitting the same price for the same property under the same contractual conditions, minority bidders face significantly lower probabilities of winning.

Turning to financing conditions, the main effect of submitting an all-cash offer is negative and statistically significant. In results not shown here, this effect is statistically insignificant when omitting the control for high LTV. Together, these results suggest that beyond a certain threshold, sellers are not concerned about a bidder's ability to finance an offer, and may even view all-cash offers as a signal of poor bidder quality.

Table 4: Bid Price Determinants

	(1) log(Offer Price)	(2) log(Offer Price)
Prob(Asian)	-0.0083***	-0.0046***
	(0.0007)	(0.0007)
Prob(Black)	-0.0266***	-0.0232***
	(0.0012)	(0.0011)
Prob(Hispanic)	-0.0056***	-0.0037***
	(0.0005)	(0.0005)
Prob(Other)	-0.0818***	-0.0456***
	(0.0039)	(0.0038)
All-Cash Offer		-0.0490***
		(0.0005)
Escalation Clause		-0.0041***
		(0.0007)
Waive Inspection		-0.0077***
		(0.0004)
Appraisal Gap		0.0096***
		(0.0007)
Sale Contingency		0.0015
		(0.0009)
LTV80+		0.0001
		(0.0004)
Constant	13.0148***	13.0207***
	(0.0003)	(0.0004)
Listing FEs	YES	YES
Observations	328794	328794
R^2	0.9817	0.9827
F-statistic	310.1192	1699.8817

Standard errors in parentheses

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

Notes: Outcome variable in both columns is the log of the total purchase amount on a submitted bid. Race variables are entered as a number between 0 and 1; coefficients on a given race category can be interpreted as the effect of switching a bidder's predicted race from 100% probability of being white (the omitted category) to 100% probability of being the given race category. Column 2 adds controls for offer characteristics, as defined in Section 2.

For Asian and Black bidders, the interaction coefficients between all-cash offers and race are not statistically significant, suggesting that cash status does not eliminate the winning gap. For Hispanic bidders, the interaction is negative and significant (-0.0716), implying that the winning disadvantage persists and may even widen among cash offers. Taken together, these patterns indicate that differential financing risk across races is unlikely to be the primary driver of the observed disparities.

Other contractual terms behave as expected. Both waiving inspection and covering an appraisal gap significantly increase the likelihood of winning, while including a sale contingency reduces it. The negative coefficient on escalation clauses likely reflects selection — escalation clauses tend to be used in highly competitive environments where baseline winning probabilities are lower.

In Column (2), we explore how the results change when we add fixed effects for the bidder's agent, using the identifiers included in the CTMe data. While differences in the probability of winning shrink for all bidders of color, the differences remain statistically and economically significant, around 3 to 4 percentage points. This suggests that an alternative explanation for our data — bidders of color having access to lower quality agents — cannot fully explain the disparities in outcomes we see in Column (1). In future work, we plan to explore additional dimensions of this concern by quantifying measures of agent experience.

Taken together, the estimates from Table 5 imply that minority bidders must increase their offers by substantial amounts — approximately 6.5 percent for Black bidders — to offset the observed acceptance gap. However, this bid-equivalent measure is derived from a linear probability model and therefore reflects a reduced-form trade-off in probability space rather than a structural feature of seller behavior. Because sellers choose among competing offers, acceptance probabilities depend on the full set of available bids and are inherently nonlinear. To obtain a structurally interpretable measure of the racial penalty — expressed as a price-equivalent trade-off in seller utility — we estimate a conditional logit choice model that explicitly models the seller's selection among competing bids in Section 5.

4.4 Winning Probability by Rank and Race

Building on the evidence above, we now ask the question slightly differently: among the top-ranked bidders within a listing, how does the probability of winning vary by race? By focusing on bidders who are already highly competitive on price, this specification isolates the margin on which seller discretion is most likely to operate. To examine this, we estimate the following within-listing model:

$$\mathbb{1}(Win_{ij}) = \sum_{r=1}^{4} \alpha_r \mathbb{1}\{Race_i = r\} + \sum_{t=1}^{6} \beta_t \{Rank_{ij} = t\} + \sum_{t=1}^{6} \sum_{r=1}^{4} \delta_{rt} \mathbb{1}\{Race_i = r\} \times \{Rank_{ij} = t\} + X'_{ij}\gamma + \theta_j + \varepsilon_{ij}, \tag{3}$$

Table 5: Determinants of Winning Bids

	(1) Baseline	(2) Baseline + Buyer Agent FE
log(Offer Price)	1.4099***	1.1728***
	(0.0137)	(0.0154)
Prob(Asian)	-0.0891***	-0.0327***
	(0.0054)	(0.0068)
Prob(Black)	-0.0890***	-0.0338***
	(0.0082)	(0.0090)
Prob(Hispanic)	-0.0851***	-0.0426***
	(0.0038)	(0.0046)
Prob(Other)	-0.1238***	-0.1112***
	(0.0301)	(0.0323)
All-Cash Offer	-0.0351***	-0.0295***
	(0.0063)	(0.0068)
Cash x P(Asian)	0.0230*	0.0133
	(0.0131)	(0.0150)
Cash x P(Black)	-0.0255	-0.0684***
	(0.0228)	(0.0260)
Cash x P(Hispanic)	-0.0716***	-0.0143
	(0.0103)	(0.0133)
Cash x P(Other)	-0.1161**	0.0938
	(0.0576)	(0.0639)
Escalation Clause	-0.0798***	-0.1028***
	(0.0047)	(0.0052)
Waive Inspection	0.0484***	0.0624***
	(0.0027)	(0.0033)
Appraisal Gap	0.0524***	0.0521***
	(0.0049)	(0.0055)
Sale Contingency	-0.0270***	-0.0402***
	(0.0061)	(0.0066)
LTV80+	-0.0399***	-0.0394***
	(0.0027)	(0.0028)
Constant	-18.0163***	-14.9553***
	(0.1783)	(0.2001)
Listing FEs	YES	YES
Buyer Agent FEs	NO	YES
Observations	328794	323030
R^2	0.1586	0.2596
F-statistic	929.9447	483.0103

Standard errors in parentheses: * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Notes: Outcome variable in both columns is an indicator variable for whether a bid won the listing for which it was submitted. Race variables are entered as a number between 0 and 1; coefficients on a given race category can be interpreted as the effect of switching a bidder's predicted race from 100% probability of being white (the omitted category) to 100% probability of being the given race category. Variables Cash × P(Race) are interactions between all-cash offer and the race variable. Column 2 adds fixed effects for buyer's agents, as identified from CTMe data.

where $Rank_{ij}$ is the within-listing rank of the bid submitted by buyer i for listing j ($t = 1$ indicates highest offer, winsorized at 6). Because identification comes entirely from comparisons within the same listing, this specification holds constant the property, the seller, the timing, and the full competitive environment. The key coefficients are the rank-by-race interactions, which reveals whether a minority bidder is less likely to win than a White bidder submitting an offer of the same rank.

Table 6 presents the estimates of equation 3. Panel A shows a sharply monotonic pattern for white bidders — relative to bids ranked sixth or lower, a top-ranked bid is 46.5 percentage points more likely to win, a second-ranked bid 21.59 percentage points more likely, and a third-ranked bid 9.62 percentage points more likely. As expected, bid rank is the dominant determinant of winning. Importantly, even the top-ranked bid does not win with certainty. This creates room for non-price considerations to influence seller choice. Our data are uniquely suited to document this margin, allowing us to disentangle the relative importance of price, contract structure, and bidder characteristics. In particular, beyond those observable and contractually specified features, bidder race also enters the seller's selection process. Panels B-E show the racial disparity among bidders of the same rank with the same contract terms. Among top rank bids, Black bidders are 12.92 percentage points less likely to win than comparable white bidders, Hispanic bidders are 19.26 percentage points less likely; and Asian bidders are 9.54 percentage points less likely. Even when a minority bidder submits the highest offer with the same contract terms, their probability of acceptance is substantially lower.

Importantly, these penalties decline as rank falls. For lower-ranked bids, racial gaps compress and eventually become statistically significant. The decline in racial penalties as rank falls is economically intuitive. When a bid is clearly inferior on price, the seller forgoes substantial revenue by selecting it, leaving little scope for non-price considerations to affect the outcome. In contrast, among top-ranked bids where multiple bids are close, sellers possess greater discretion. The largest racial disparities therefore emerge precisely where the cost of exercising discretion is lowest. This pattern aligns with Becker (1957)'s insight that discriminatory behavior is most visible when the cost of discrimination is low.

4.5 Outcome Test

In this subsection, we provide additional evidence that differences in bidder outcomes by race cannot be rationalized by latent differences in bidder quality. In the sections above, we have included controls for offer terms and financing structure to control for ways in which different bids might present different closing risk to sellers. Sellers may still be concerned that bidders of color present additional closing risk: if bidders of color tend to take longer to close transactions once their offers have been accepted, or are disproportionately

Table 6: Winning Bids by Rank and Race

Panel A: Baseline Ranks	
Rank 1	0.4650***
	(0.0066)
Rank 2	0.2159***
	(0.0067)
Rank 3	0.0962***
	(0.0073)
Rank 4	0.0529***
	(0.0081)
Rank 5	0.0220**
	(0.0091)

Panel B: $P(\text{Asian})\times$ **Rank**	
Prob(Asian)	-0.0165
	(0.0113)
P(Asian) x Rank 1	-0.0954***
	(0.0143)
P(Asian) x Rank 2	-0.0883***
	(0.0146)
P(Asian) x Rank 3	-0.0302*
	(0.0158)
P(Asian) x Rank 4	-0.0251
	(0.0175)
P(Asian) x Rank 5	-0.0049
	(0.0202)
Cash x P(Asian)	0.0209
	(0.0128)

Panel C: $P(\text{Black})\times$ **Rank**	
Prob(Black)	-0.0079
	(0.0177)
P(Black) x Rank 1	-0.1292***
	(0.0224)
P(Black) x Rank 2	-0.0912***
	(0.0231)
P(Black) x Rank 3	-0.0476*
	(0.0252)
P(Black) x Rank 4	-0.0606**
	(0.0284)
P(Black) x Rank 5	-0.0094
	(0.0323)
Cash x P(Black)	-0.0890***
	(0.0222)

Panel D: $P(\text{Hispanic})\times$ **Rank**	
Prob(Hispanic)	0.0147*
	(0.0078)
P(Hispanic) x Rank 1	-0.1926***
	(0.0098)
P(Hispanic) x Rank 2	-0.1108***
	(0.0101)
P(Hispanic) x Rank 3	-0.0528***
	(0.0110)
P(Hispanic) x Rank 4	-0.0371***
	(0.0122)
P(Hispanic) x Rank 5	-0.0142
	(0.0138)
Cash x P(Hispanic)	-0.1133***
	(0.0101)

Panel E: $P(\text{Other})\times$ **Rank**	
Prob(Other)	0.0488
	(0.0579)
P(Other) x Rank 1	-0.1792**
	(0.0767)
P(Other) x Rank 2	-0.2875***
	(0.0749)
P(Other) x Rank 3	-0.1469*
	(0.0816)
P(Other) x Rank 4	-0.0985
	(0.0924)
P(Other) x Rank 5	-0.0767
	(0.1020)
Cash x P(Other)	-0.1900***
	(0.0572)

Panel F: Summary	
All-Cash Offer	-0.0220***
	(0.0061)
Escalation Clause	-0.0562***
	(0.0046)
Waive Inspection	0.0343***
	(0.0026)
Appraisal Gap	0.0347***
	(0.0047)
Sale Contingency	-0.0315***
	(0.0059)
Listing Fixed Effects	YES
R^2	0.212
Observations	328794

Standard errors in parentheses
* $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

Notes: Outcome variable in all regressions is an indicator for whether a bid was the winning bid for a given listing. Panel A shows the main effect associated with a given rank, coefficients are relative to the omitted category of bids ranked 6th or higher. Panels B-E show estimates for the main effect associated with each race, as well as interactions between each race and each rank. Race enters into the regression as a vector of probabilities, as in previous tables. Panel F shows the effects of offer terms and model statistics.

likely to trigger a legal clause to cancel a deal entirely, differences in acceptance probability reflect efforts by sellers to avoid these costs rather than race-based premia. In Table 7, we show that such expectations are unfounded, if they exist.

Each pair of columns shares an outcome variable reflecting some costly outcome related to closing: "Time-to-close" is a continuous measure of how much time elapses between the date an offer is signed and the date that the property is transacted (see Section 2.7); TTC30, TTC45, and TTC60 are all indicator variables for time-to-close of 30, 45, and 60 days (roughly the median, 90th percentile, and 99th percentile of listings in our sample). Using only observations corresponding to winning bids, the first column in each pair regresses the outcome variable on the vector of race probabilities and log(price), and the second column in each pair adds controls for offer terms to account for the varying liquidity and financing risks associated with an offer's contract structure. We include fixed effects for property, Census-tract year, and 20-point LTV bins in all specifications.

Across all specifications, Black and Hispanic bidders win a listing, they do not present meaningful differences in terms of long closing risk relative to white buyers; Asian winners tend to take 2 days longer to close a deal and are roughly 7 percentage points more likely to require 30 or 45 days to close a deal. For each race, these estimates vary little with the inclusion of offer terms, which provide different mechanisms for speeding up or slowing down a deal's closing. The coefficients on these offer terms all conform to expectations: when a bidder presents an all-cash offer, waives inspection, or is willing to cover an appraisal gap, deals close faster.

5 A Discrete Choice Model of Seller Selection

The reduced-form evidence establishes that bidders of color bid less *ex ante*, but even when making the most competitive offers, face higher thresholds for seller acceptance. This is difficult to reconcile with explanations based on stronger preference or thinner markets for minority buyers. Instead, they point to disparities arising at the seller's acceptance stage: not all bids are treated equally. Consistent with this, we find that minority buyers, particularly Black and Hispanic buyers, are less likely to win even when submitting comparable offers, and in many cases, top-ranked bids.

However, these results do not speak directly to the extensive literature we discuss in our introduction, showing that conditional on winning, buyers of color pay higher prices for similar homes. We are not able to replicate such a result for the Denver MSA using transactions data from 2014-2024. In Appendix A.2, we follow Bayer et al. (2017) and implement a repeat-sales methodology: using a dataset of CoreLogic transactions records merged with HMDA data, we regress sales price on buyer race, using property and

Table 7: Outcome Test

	(1) Time to Close	(2) Time to Close	(3) 1(TTC ⩾ 30)	(4) 1(TTC ⩾ 30)	(5) 1(TTC ⩾ 45)	(6) 1(TTC ⩾ 45)	(7) 1(TTC ⩾ 60)	(8) 1(TTC ⩾ 60)
P(Asian)	2.0797**	2.0225**	0.0813*	0.0788*	0.0611**	0.0622**	0.0130	0.0130
	(0.8828)	(0.8805)	(0.0445)	(0.0444)	(0.0274)	(0.0274)	(0.0112)	(0.0112)
P(Black)	-0.2697	-0.2549	0.0179	0.0187	0.0421	0.0439	0.0220	0.0224
	(1.3013)	(1.2972)	(0.0656)	(0.0655)	(0.0404)	(0.0404)	(0.0164)	(0.0165)
P(Hispanic)	0.4107	0.4376	0.0213	0.0231	0.0318	0.0315	0.0024	0.0022
	(0.6400)	(0.6380)	(0.0323)	(0.0322)	(0.0199)	(0.0199)	(0.0081)	(0.0081)
P(Other)	4.3635	3.9107	0.4777*	0.4613*	0.0916	0.0908	0.0685	0.0670
	(5.2556)	(5.2403)	(0.2650)	(0.2645)	(0.1632)	(0.1633)	(0.0664)	(0.0665)
Log price	3.2831*	2.9020	0.1385	0.1172	0.0262	0.0287	0.0002	0.0012
	(1.9807)	(1.9799)	(0.0999)	(0.0999)	(0.0615)	(0.0617)	(0.0250)	(0.0251)
Cash offer		-9.3527***		-0.4363***		-0.0577		0.0047
		(2.4832)		(0.1254)		(0.0774)		(0.0315)
Escalation clause		-0.1517		-0.0449		-0.0172		-0.0003
		(0.9405)		(0.0475)		(0.0293)		(0.0119)
Inspection waived		-1.3285***		-0.0681***		0.0113		0.0024
		(0.4952)		(0.0250)		(0.0154)		(0.0063)
Appraisal gap		-2.1565***		-0.0385		-0.0378		-0.0195*
		(0.8060)		(0.0407)		(0.0251)		(0.0102)
Sale contingency		0.6562		0.0562		0.0401		0.0027
		(1.0502)		(0.0530)		(0.0327)		(0.0133)
Constant	-10.9243	-5.0268	-1.2268	-0.9081	-0.2605	-0.2886	0.0044	-0.0090
	(25.6131)	(25.6069)	(1.2917)	(1.2926)	(0.7952)	(0.7977)	(0.3238)	(0.3249)
Property FEs	YES	YES	YES	YES	YES	YES	YES	YES
Tract-Year FEs	YES	YES	YES	YES	YES	YES	YES	YES
LTV Bins	YES	YES	YES	YES	YES	YES	YES	YES
Observations	11168	11168	11168	11168	11168	11168	11168	11168
R^2	0.7273	0.7296	0.7193	0.7211	0.7049	0.7053	0.6802	0.6805

Standard errors in parentheses

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

Notes: Outcome variables: Cols 1-2, continuous measure of time-to-close (defined as time between when buyer and seller agree to a contract and time when property transacts, measured in days, see Section 2.7); Cols 3-4, time-to-close of 30 days or more; Cols 5-6, time-to-close of 45 days or more; Cols 7-8, time-to-close of 60 days or more. The first column in each pair only controls for race and the log of bid price, the second column in each pair adds controls for offer terms. Fixed effects for property, Census tract-year combination, and 20-point LTV bins included in all regressions. Sample is restricted to winning bids where time-to-close is between [1, 364] days (drops 1% of winning bids).

tract-year fixed effects and controlling for buyer characteristics. When we apply this estimation strategy to data on home sales in the Denver MSA between 2000 and 2007 (matching their sample period), we find estimates consistent with their main finding. However, these premia do not appear when we run the same estimation using data from 2014-2021 (or a pre-Covid subsample of Jan. 2014–Feb. 2020).

The transaction-level evidence restricts the analysis by only studying winning bids, but using our bid-level data, we can more directly investigate whether buyers of color face bid-equivalent penalties. The existence of such penalties can jointly explain our findings that bidders of color tend to bid less and are less likely to win when making the same offer as a white competitor; this would also be wholly consistent with transaction prices that do not show racial premia – for example, bidders of color could make their offers more attractive using non-price margins. While the reduced-form estimates quantify racial gaps in linear probability and price space, in practice, sellers choose one offer from a set of competing bids. To express the observed disparities as economically interpretable trade-offs in seller utility, we therefore estimate a conditional logit model that embeds the seller's discrete choice problem.

5.1 Seller Utility and Choice Probabilities

We model the seller as choosing the offer that maximizes latent utility. For listing j, seller utility from accepting offer i is given by

$$U_{ij} = \sum_{r=1}^{4} \alpha_r \mathbb{1}\{Race_i = r\} + \beta \ln(Bid_{ij}) + X'_{ij}\gamma + \theta_j + \varepsilon_{ij}. \tag{4}$$

Under the assumption that ε_{ij} follows a Type 1 extreme value distribution, the probability that offer i wins listing j is

$$P_{ij} = \frac{exp(V_{ij})}{\sum_{k \in I_j} exp(V_{kj})}, \quad V_{ij} = \sum_{r=1}^{4} \alpha_r \mathbb{1}\{Race_i = r\} + \beta \ln(Bid_{ij}) + X'_{ij}\gamma + \theta_j \tag{5}$$

Identification continues to rely exclusively on within-listing variations. All property attributes, seller characteristics, and local market conditions are absorbed by θ_j. The estimates therefore compare competing offers with the same contract terms submitted for the same property at the same time.

5.2 Baseline Results

Table 8 reports the conditional logit estimates in the left column. The coefficient on $\ln(Bid_{ij})$ is large and precisely estimated (15.32), confirming that offer price is the dominant determinant of acceptance. At the same time, the race coefficients are negative and highly statistically significant, -0.342 for Black bidders, -0.365 for Hispanic bidders, -0.338 for Asian bidders, and -0.653 for Other bidders. These coefficients

represent direct shifts in seller utility associated with accepting an offer from each group, holding price and contract terms fixed.

To interpret magnitudes, we translate the estimated utility coefficients into economically interpretable race premia, measured as the additional percentage of the offer price that a minority buyer must bid — holding observable offer terms fixed — to achieve the same probability of acceptance as a white buyer. Specifically, the bid-equivalent penalty for race r is given by:

$$\Delta_r = -\frac{\alpha_r}{\beta} \tag{6}$$

The imputed race penalties are reported in the right column of Table 8. Black bidders must increase their offer by approximately 2.23 percent to offset the estimated utility penalty; Hispanic bidders by 2.38 percent; and Asian bidders by 2.20 percent. In a housing market with median home prices near $500,000, a 2-2.5 percent penalty corresponds to $10,000-$12,500. Importantly, this penalty arises conditional on identical bid levels and contractual characteristics within the same listing, reflecting differential treatment at the margin of seller selection rather than differences in offer quality.

These estimates clarify the mechanism suggested by the reduced-form evidence. Holding the level of the bid constant, bidders of color face a lower acceptance probability *ceteris paribus*. Bidders of color submit lower initial offers on average, yet conditional on winning, they need to pay more than bidders from other racial groups. The logit framework shows that this pattern is consistent with a negative seller utility component attached to offers from bidders of color.

The remaining coefficients are consistent with our expectation. Waiving the inspection contingency increases the probability of acceptance by an amount equivalent to raising the bid by approximately 1.5-2.0 percent, reflecting sellers' valuation of reduced execution risk. Similarly, including an appraisal gap clause is valued at roughly 1.0-1.5 percent of the bid, consistent with sellers placing substantial weight on protection against appraisal shortfalls.

Earnest money also plays an important role. A one-percentage-point increase in the earnest money share is equivalent to increasing the bid by approximately 0.2-0.3 percent, indicating that sellers interpret larger deposits as a credible signal of buyer commitment. In contrast, sale contingencies substantially reduce acceptance probabilities: including a sale contingency is equivalent to lowering the bid by roughly 1.0-1.5 percent, underscoring sellers' aversion to deal uncertainty and delay. Similarly, offers with higher cash fractions are valued more highly by sellers, reflecting the premium placed on reducing financing risk.

Taken together, these magnitudes confirm that sellers evaluate offers along multiple dimensions, trading off price against execution risk and contractual flexibility. Because these economically relevant contract

terms are explicitly observed and controlled for in all specifications, the remaining coefficients can be interpreted as capturing differential treatment beyond offer quality.

Even conditional on price and all observable contract terms, buyer race remains a statistically and economically significant predictor of seller choice. The baseline estimates imply that, holding offer characteristics fixed, sellers behave as if minority buyers must increase their bids by 2-2.5% to be treated equivalently to White buyers. Notably, these race-based bid premia are comparable in size to the implicit price effects of non-price contract terms, such as waiving an inspection contingency, offering an appraisal gap, or substantially increasing earnest money. In other words, buyer race enters seller decision-making with weight similar to some of the most consequential dimensions of offer quality.

Because these comparisons are made within the same listing and conditional on a rich set of contract terms, the estimated bid premia cannot be attributed to differences in housing quality, neighborhood selection, seller heterogeneity, or buyer financing conditions or deal risk. Instead, they reflect systematic differences in the bid amounts sellers require from buyers of different races when evaluating otherwise comparable offers.

5.3 Heterogeneity in the Bid-Equivalent Racial Penalty

The baseline conditional logit estimates establish that minority buyers face 2-2.5% bid-equivalent penalties even after conditioning on listing fixed effects and detailed contractual terms. We now examine how these penalties vary across market environments and neighborhood characteristics. Doing so sheds light on the economic forces underlying seller discretion and helps distinguish between alternative mechanisms.

To do so, we extend the conditional logit model in Section 5.1 to allow seller preferences to vary smoothly with a continuous neighborhood characteristic z_j (e.g., percent White, house value, median income, PM2.5 exposure). This approach preserves the within-listing identification strategy while permitting the parameters governing seller utility to vary flexibly across local conditions.

Local Logit Specification

For a given value z, seller utility from accepting offer i in listing j is specified as

$$U_{ij}(z) = \sum_{r=1}^{4} \alpha_r(z)\, \mathbb{1}\{Race_i = r\} + \beta(z) \ln(Bid_{ij}) + X_{ij}'\gamma(z) + \theta_j + \varepsilon_{ij}. \tag{7}$$

As in Section 5.1, ε_{ij} is assumed to follow a Type I extreme value distribution. The probability that offer i wins listing j at neighborhood characteristic z is therefore

Table 8: Structural Estimation Baselines

	Conditional Logit Estimates	$-\beta_{\text{race}}/\beta_{\ln(\text{bid})}$
Buyer race		
Asian	-0.342***	0.0225***
	(0.046)	(0.0018)
Black	-0.367***	0.0223***
	(0.019)	(0.0030)
Hispanic	-0.344***	0.0239***
	(0.028)	(0.0013)
Other	-0.675***	0.0441***
	(0.151)	(0.0099)
Offer characteristics		
ln(Bid)	15.331***	
	(0.092)	
% Earnest Money	0.146***	
	(0.009)	
Escalation Clause	-0.443***	
	(0.040)	
Waived Inspection	0.414***	
	(0.021)	
Appraisal Gap	0.536***	
	(0.082)	
Sale Contingency	-0.131***	
	(0.037)	
Cash Share (Q2)	0.302***	
	(0.018)	
Cash Share (Q3)	0.309***	
	(0.021)	
Cash Share (Q4)	0.112***	
	(0.032)	
LTV > 80%	-0.100***	
	(0.024)	

Notes: All estimates are statistically significant ($p < 0.001$). Conditional logit estimates: Each column reports conditional logit estimates of offer acceptance with listing fixed effects and full offer controls. The rows for Asian, Black, Hispanic, and Other correspond to the variables $Pr(Asian)$, $Pr(Black)$, $Pr(Hispanic)$, and $Pr(Other)$ from above. Standard errors are in parentheses (t-stats omitted). Cash share variables are indicators for quartiles of the cash-to-price ratio (Q1 omitted). The bid coefficient corresponds to $\ln(\text{Bid})$. *Delta method:* Entries are delta-method estimates of the race premium, defined as $-\beta_{\text{race}}/\beta_{\ln(\text{bid})}$. Standard errors (delta method) are reported in parentheses. All specifications include listing fixed effects and full offer controls.

$$P_{ij}(z) = \frac{\exp\left(V_{ij}(z)\right)}{\sum_{k \in I_j} \exp\left(V_{kj}(z)\right)}, \tag{8}$$

where

$$V_{ij}(z) = \sum_{r=1}^{4} \alpha_r(z)\, \mathbb{1}\{Race_i = r\} + \beta(z) \ln(Bid_{ij}) + X_{ij}'\gamma(z) + \theta_j. \tag{9}$$

The key departure from the baseline model is that the preference parameters $\alpha_r(z)$, $\beta(z)$, and $\gamma(z)$ are allowed to vary with z. This flexibility allows seller responses to race and price to differ systematically across neighborhoods.

Local Likelihood Estimation

We estimate the model using a locally weighted maximum likelihood approach. For each value z, we maximize

$$\mathcal{L}(z) = \sum_{j=1}^{J} K_h(z_j - z) \left[\ln P_{i_j^* j}(z)\right], \tag{10}$$

where i_j^* denotes the winning offer in listing j, and $K_h(\cdot)$ is a Gaussian kernel weight:

$$K_h(z_j - z) = \frac{1}{\sigma h \sqrt{2\pi}} \exp\left(-\frac{(z_j - z)^2}{2h^2\sigma_z^2}\right). \tag{11}$$

Here, σ_z is the standard deviation of z, and h is a bandwidth parameter. Listings closer to z receive greater weight, allowing the estimated preference parameters to vary smoothly across the distribution of neighborhood characteristics.

Identification

Identification continues to rely exclusively on within-listing variation. Listing fixed effects θ_j absorb all property characteristics, seller identity, and time-specific local market conditions. The local weighting affects how listings contribute to the likelihood at each z, but not the source of identification. The estimates therefore compare competing offers for the same property, while allowing seller preferences to differ across neighborhood environments.

Bid-Equivalent Racial Penalty

At each value of z, we compute the bid-equivalent racial penalty as

$$\Delta_r(z) = -\frac{\alpha_r(z)}{\beta(z)}, \tag{12}$$

which measures the percentage increase in bid required for a bidder of race r to simply offset the disutility of being race i in the eye of the seller at neighborhood characteristic z.

5.4 Heterogeneity Estimates

For ease of interpretation, we aggregate Asian, Black, Hispanic, and Other minority groups into a single category: Nonwhite. Using the local logit estimates, Figures 2-3 plot how the bid-equivalent Nonwhite penalty varies across dimensions of market competition, neighborhood amenities, and socioeconomic characteristics.

Market Structure We begin by examining how bid-equivalent race penalties vary with competition. We distinguish between two related but conceptually distinct dimensions: market-level competitiveness and listing-level bidding intensity.

Market-level competition is captured using two complementary ZIP-year measures: market tightness and average time-on-the-market (TOM). In search-and-matching models of housing markets (e.g., Genesove and Han (2012)), market tightness - defined as the ratio of potential buyers to potential sellers - governs the arrival rate of offers to sellers and thus the expected time on the market. In equilibrium, tighter markets are characterized by higher buyer arrival rates, shorter time on the market, stronger seller outside options, and higher reservation values - these are necessary conditions that expand the scope for seller discretion. Slack markets exhibit the opposite.

This inverse relationship between tightness and TOM is a core implication of search theory and provides a natural framework for examining seller discretion. As the pool of potential buyers expands relative to available listings, the probability of sale in any given period rises and the opportunity cost of rejecting a current offer falls. Sellers therefore gain greater latitude - not only in bargaining over price, but also in selecting among competing buyers.

Leveraging our unique bidding and listing data, we construct market tightness, as the ratio of total bids - both accepted and rejected - to total listings, including those that ultimately fail to transact. This measure captures the full scope of demand and supply conditions in the local market. We complement it with average TOM, defined as the number of days a listing remains active, including those that do not sell. By incorporating unsuccessful listings and rejected offers, market tightness and TOM reflect the intensity of buyer search activity and the breadth of sellers' outside options more comprehensively than the traditional

measures based solely on completed sales.

The empirical patterns closely mirror theoretical predictions. Panel (a) in Figure 2 shows that Nonwhite bid-equivalent race penalty increases nearly linearly with market tightness. As tightness rises from 2.8 to 4.2, the implied penalty increases from 2.2 percent to 2.4 percent of the bid. Panel (b) reveals a complementary pattern along the TOM dimension. In neighborhoods where homes sell quickly - around 60 days on average - the racial penalty is approximately 2.5-2.6 percent. As average TOM lengthens toward 90 days, the penalty declines to roughly 1.9 percent.

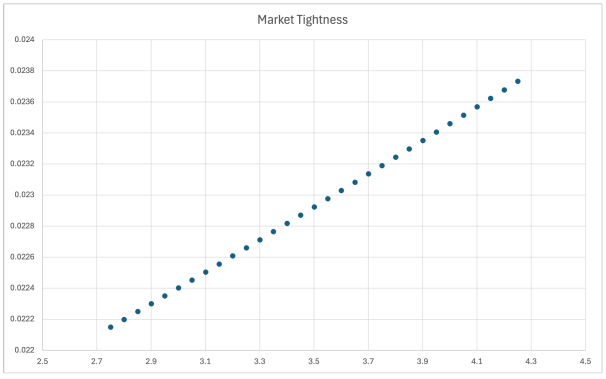
These gradients provide direct evidence linking the economics of discrimination (Becker, 1957) to a search-frictional environment. In tighter markets and fast turnover environments where buyer arrival rates are high and the opportunity cost rejecting a current offer is low, sellers possess stronger outside options and greater latitude to exercise discretion among otherwise comparable offers. This amplifies race-based disparities. Conversely, when demand is weak and properties linger on the market, the cost of rejecting a current offer rises and the scope for discretionary screening narrows. Racial disparities are therefore greatest precisely in environments where seller market power is strongest: when the value of continued search is highest and the option to wait for a preferred buyer is most valuable.

We next turn to listing-level competition, measured by the number of distinct bidders and the number of submitted offers for a given property. Whereas tightness and TOM capture aggregate search conditions and equilibrium scarcity, listing-level measures reflect the intensity of competition within a specific listing, holding constant the broader market environment.

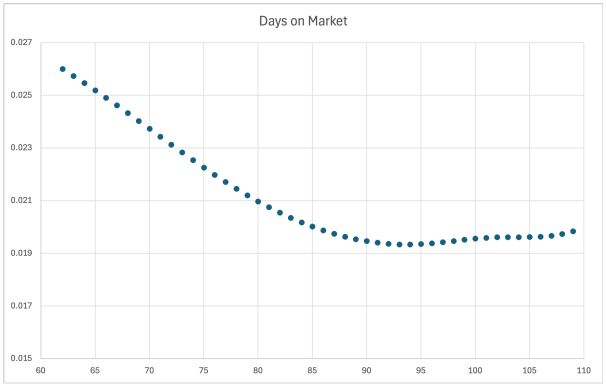
Panels (c)-(d) of Figure 2 reveal a similar pattern at this more granular level. As the number of bidders increases from two to ten, the implied bid-equivalent penalty rises from roughly 2.1 percent to 2.7 percent, reaching approximately 3.4 percent in listings with fifteen bidders. The economic magnitude is substantial: in highly contested listings, minority bidders must submit 3.4 percent higher offers to appear the same as white bidders in the eyes of the seller. Using the total number of offers instead of bidders yields similar but slightly attenuated patterns, consistent with the fact that multiple offers may originate from the same buyer and thus less precisely capture competitive breadth.

Taken together, across market tightness, neighborhood-level transaction speed, and listing-level bidding intensity, a consistent pattern emerges: racial penalties expand when sellers possess stronger outside options. In environments where competition is intense, the cost of continued search falls and the scope for discretionary screening widens. This pattern is consistent with Becker (1957)'s insight that discrimination is more likely to persist when market power reduces its economic cost. Here, stronger seller market power translates directly into larger racial disparities in acceptance thresholds.

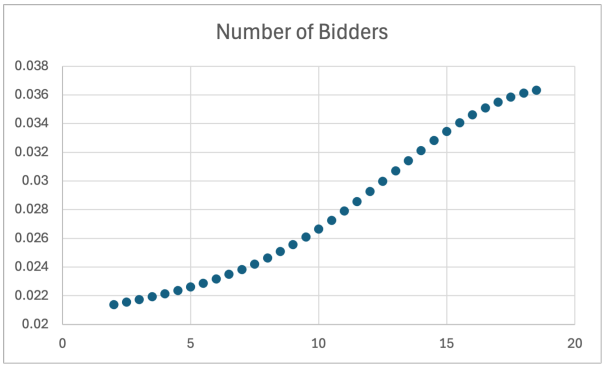
Figure 2: Race Penalty by Competitiveness



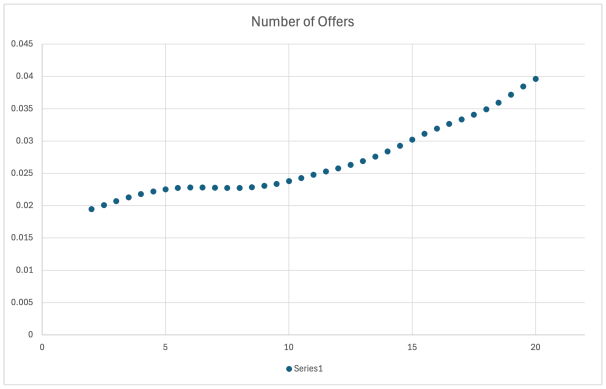
(a) Market Tightness



(b) Days on Market



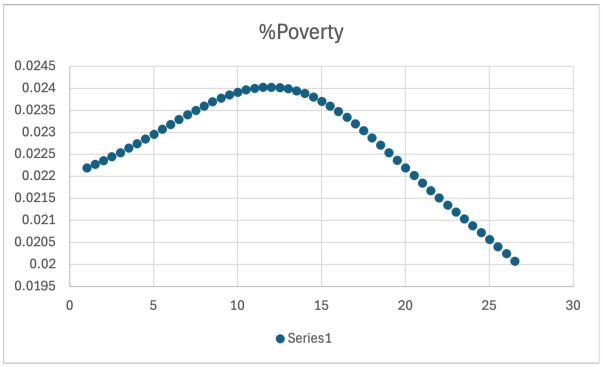
(c) Number of Bidders



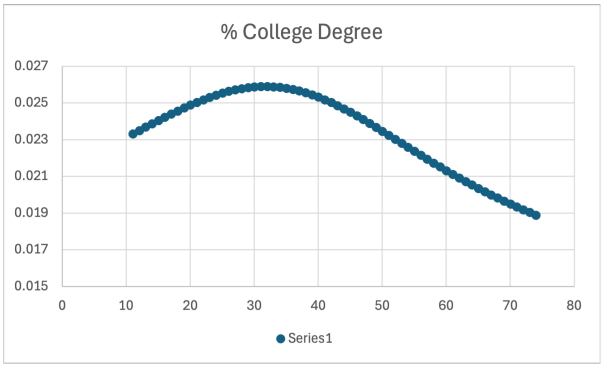
(d) Number of Offers

The figures plot the estimated bid-equivalent racial penalty across measures of market and transaction-level competitiveness. Panel (a) reports the implied penalty as a function of market tightness, measured by the ratio of total bids to total listings at the ZIP3 $\times$ year level. Panel (b) presents the penalty by days on market, capturing listing-level demand pressure. Panel (c) reports the relationship by the number of unique bidders at the transaction level, while Panel (d) shows the corresponding pattern by the number of unique offers. Each point reflects the model-implied price-equivalent racial penalty, conditional on listing fixed effects and contractual characteristics.

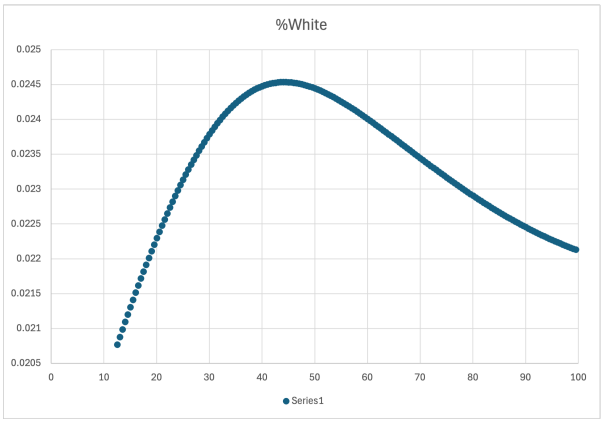
Figure 3: Race Penalty by Neighborhood Characteristics



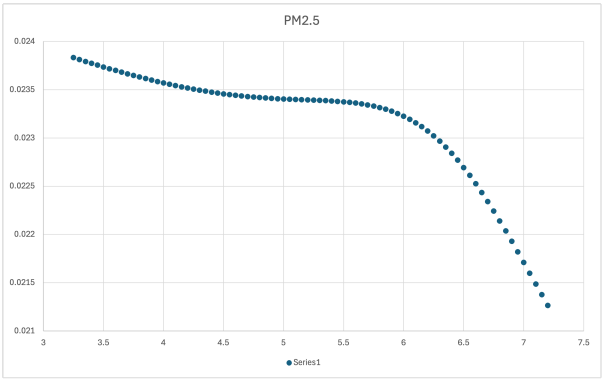
(a) Poverty Rate (%)



(b) College Degree and Higher (%)



(c) White Population (%)



(d) PM2.5 (Pollution)

The figure plots the model-implied bid-equivalent racial penalty across neighborhood characteristics. Panel (a) varies the poverty rate at the ZIP5 X Year level. Panel (b) reports the penalty across the share of residents with a college degree or higher. Panel (c) varies the share of White residents. Panel (d) plots the penalty against ambient PM2.5 concentrations at the ZIP5 X Year level. Each point represents the estimated price-equivalent racial penalty from the conditional logit model, holding listing fixed effects and detailed contractual characteristics constant.

Neighborhood Characteristics We next examine how race-based acceptance penalties vary with neighborhood characteristics. Figures 3 plots the bid-equivalent Nonwhite penalty against socioeconomic status (SES) (Panels a-b), racial composition (Panel c), and environmental quality (Panel d).

Across poverty rates, college attainment, and the share of White residents, the racial penalty exhibits an inverted-U shape. It is modest in very high-SES neighborhoods, rises in middle-SES areas, and declines again in the lowest-SES neighborhoods. A parallel non-monotonicity appears with respect to racial composition: penalties are smallest where the White share is either very low or very high and largest in racially mixed neighborhoods.

These gradients point to a common mechanism. At the upper end of the SES distribution - neighborhoods with high incomes, high college attainment, and elevated property values - high listing prices and substantial down payments already screen buyers on observable financial capacity, reducing uncertainty about buyer type and diminishing the marginal value of discretionary racial screening. At the opposite extreme - very low-SES or heavily minority neighborhoods - sellers may anticipate a buyer pool that is already diverse, limiting the effectiveness of screening in shaping neighborhood composition.

The largest penalties arise between these extremes: moderate poverty, moderate education, racially mixed population - precisely the neighborhoods at the margin of demographic and socioeconomic transition. Two complementary mechanisms help explain why penalties concentrate in these locations.

First, these areas often coincide with improving amenities. Panel (d) shows that the racial penalty declines monotonically with PM2.5 concentration, indicating that discrimination is strongest in cleaner neighborhoods and attenuates in more polluted ones. In the Denver context, lower PM2.5 levels characterize residential neighborhoods with improving environmental quality and amenity value, while higher exposure is concentrated in dense, historically industrial corridors. Minority buyers therefore face greater barriers precisely in cleaner, higher-amenity locations. Importantly, many such neighborhoods, especially those with intermediate SES, remain priced below fully gentrified markets, attracting strong demand while staying broadly accessible. This combination intensifies competition and strengthens sellersâ outside options, expanding their latitude to screen among otherwise comparable buyers. Discrimination thus shapes not only who transacts, but also where.

Second, transitional neighborhoods often lie near demographic tipping points. Panel (c) shows that the penalty peaks in neighborhoods that are racially mixed - around 40-50% White population - and declines toward both extremes, closely mirroring the tipping-point dynamics documented by Card et al. (2008). When neighborhood composition is perceived as fragile, the identity of the marginal buyer may be viewed as consequential for future trajectory. Incentives to screen are therefore strongest where demographic outcomes are contested but not yet settled. In more homogeneous neighborhoods - whether predominantly

White or predominantly minority - the scope for influencing composition through screening diminishes, and penalties attenuate. Our finding here complements prior evidence from Hanson and Hawley (2011) and Gold et al. (2025), who document that rental discrimination is most pronounced in neighborhoods with intermediate White population shares.

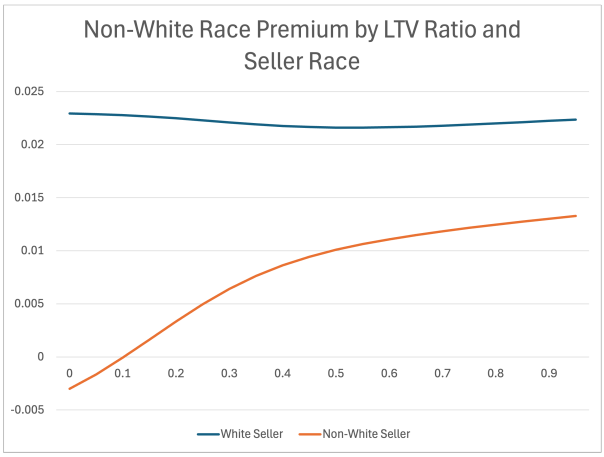
Taken together, heterogeneity across market competition, socioeconomic status, racial composition, and pollution levels converges on a unified mechanism centered on seller discretion. Racial disparities are not spatially uniform; they concentrate where sellers face abundant choices and view the identity of the marginal buyer as consequential, and hence have greater latitude and incentives in screening among competing offers. These conditions are most prevalent in transitional neighborhoods - areas with improving amenities, moderate yet rising prices, and evolving demographic composition. These are precisely the neighborhoods that offer critical pathways to upward mobility for minority households. Yet our evidence indicates that it is in these spaces of opportunity that access is most constrained.

Seller Race By linking home bidding records with historical mortgage applications recorded under the Home Mortgage Disclosure Act (HMDA) via CoreLogic, we successfully recovered the self-reported race for 18,703 sellers within our dataset. This allows us to explore how the estimated racial premium varies with seller race.

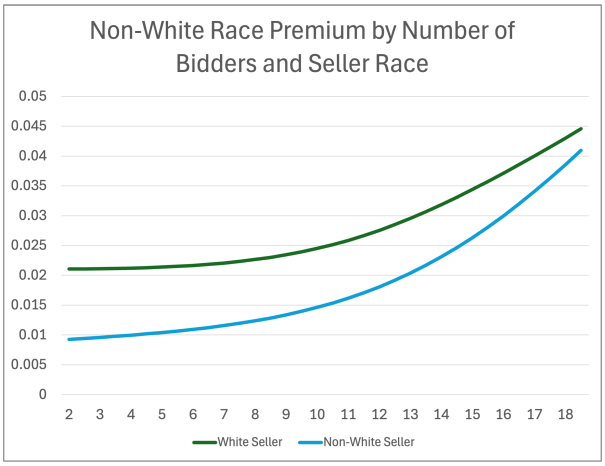
Figure 4

indicates a substantial higher racial premium required by white sellers relative to non-white sellers. Specifically, the estimates reveal that while racial premium increases with LTV for non-white sellers, it is invariant to LTV for white sellers. This suggests that the differential treatment of bidders of color by white sellers is less driven by financial concerns. Moreover, in environments with a high volume of bidders - where the "cost" of discrimination is relatively low for buyers - the price gap between white and non-white sellers tends to narrow, again suggesting the discretion amplifies even for non-white sellers when its cost is low.

Figure 4: Race Penalty by Seller Race



(a) Seller Race by LTV



(b) Seller Race by # Bidders

T

References

Ambrose, B. W., Conklin, J. N., & Lopez, L. A. (2021). Does borrower and broker race affect the cost of mortgage credit? *The Review of Financial Studies*, *34*(2), 790–826.

Bayer, P., Casey, M., Ferreira, F., & McMillan, R. (2017). Racial and ethnic price differentials in the housing market. *Journal of Urban Economics*, *102*, 91–105. https://doi.org/10.1016/j.jue.2017.07.004

Becker, G. S. (1957). *The Economics of Discrimination*. The University of Chicago Press.

Box-Couillard, S., & Christensen, P. (2024, August). Racial Housing Price Differentials and Neighborhood Segregation. https://doi.org/10.3386/w32815

Card, D., Mas, A., & Rothstein, J. (2008). Tipping and the Dynamics of Segregation*. *The Quarterly Journal of Economics*, *123*(1), 177–218. https://doi.org/10.1162/qjec.2008.123.1.177

Christensen, P., Sarmiento-Barbieri, I., & Timmins, C. (2022). Housing Discrimination and Pollution Exposures in the United States. *Review of Economics and Statistics*, *forthcoming*.

Christensen, P., & Timmins, C. (2022). Sorting or steering: The effects of housing discrimination on neighborhood choice. *Journal of Political Economy*, *130*(8), 2110–2163.

Dymski, G. A. (2006). Discrimination in the credit and housing markets: Findings and challenges. *Handbook on the Economics of Discrimination*, *215*, 220.

Galster, G., & Godfrey, E. (2005). By words and deeds: Racial steering by real estate agents in the us in 2000. *Journal of the American planning association*, *71*(3), 251–268.

Genesove, D., & Han, L. (2012). Search and matching in the housing market. *Journal of Urban Economics*, *72*(1), 31–45. https://doi.org/10.1016/j.jue.2012.01.002

Gold, D., Han, L., & Timmins, C. (2025). *Hidden Biases: Selective Advertising in the Rental Housing Market* (Working Paper).

Guryan, J., & Charles, K. K. (2013). Taste-Based or Statistical Discrimination: The Economics of Discrimination Returns to its Roots. *The Economic Journal*, *123*(572).

Han, L., & Hong, S.-H. (2024). Cash is king? Understanding financing risk in housing markets. *Review of Finance*, *28*(6), 2083–2118. https://doi.org/10.1093/rof/rfae025

Hanson, A., & Hawley, Z. (2011). Do landlords discriminate in the rental housing market? Evidence from an internet field experiment in US cities. *Journal of Urban Economics*, *70*(2), 99–114. https://doi.org/10.1016/j.jue.2011.02.003

Ihlanfeldt, K., & Mayock, T. (2009). Price discrimination in the housing market. *Journal of Urban Economics*, *66*(2), 125–140. https://doi.org/10.1016/j.jue.2009.05.004

Myers, C. K. (2004). Discrimination and neighborhood effects: Understanding racial differentials in US housing prices. *Journal of Urban Economics*, *56*(2), 279–302. https://doi.org/10.1016/j.jue.2004.03.006

Ondrich, J., Ross, S., & Yinger, J. (2003). Now You See It, Now You Don't: Why Do Real Estate Agents Withhold Available Houses from Black Customers? *Review of Economics and Statistics*, *85*(4), 854–873.

Ondrich, J., Stricker, A., & Yinger, J. (1998). Do Real Estate Brokers Choose to Discriminate? Evidence from the 1989 Housing Discrimination Study. *Southern Economic Journal*, 880–901.

van Donkelaar, A., Hammer, M. S., Bindle, L., Brauer, M., Brook, J. R., Garay, M. J., Hsu, N. C., Kalashnikova, O. V., Kahn, R. A., Lee, C., Levy, R. C., Lyapustin, A., Sayer, A. M., & Martin, R. V. (2021). Monthly Global Estimates of Fine Particulate Matter and Their Uncertainty. *Environmental Science & Technology*, *55*(22), 15287–15300. https://doi.org/10.1021/acs.est.1c05309

Yinger, J. (1995). *Closed doors, opportunities lost: The continuing costs of housing discrimination*. Russell Sage Foundation.

Zhao, B., Ondrich, J., & Yinger, J. (2006). Why Do Real Estate Brokers Continue to Discriminate? Evidence from the 2000 Housing Discrimination Study. *Journal of Urban Economics*, *59*(3), 394–419.

A. APPENDIX

A.1 Validating Race Imputation

As described in Section 2, we use HMDA data to identify buyer race for 76,623 listings out of the 129,369 for which we have CTMe data. In this section, we describe how we use this information to cross-validate the imputed race measures from MRI.

We classify HMDA buyer races into Asian (4.85%), Black (2.64%), Hispanic (13.01%), White (68.21%), other races (0.63%), and missing (10.66%). We keep only listings where the buyer race falls into one of the first four categories. Then, we estimate a multinomial logit to see how well the MRI-predicted race of the bidder we identify as the winner matches with the buyer's self-reported race in HMDA. The estimating equation is:

$$\Pr(HMDA\ Race_i = j) = \frac{\exp(\alpha_j + \mathbf{X}_i\boldsymbol{\beta}_j)}{1 + \sum_{k\neq 4}\exp(\alpha_k + \mathbf{X}_i\boldsymbol{\beta}_k)} \quad \text{for } j \neq 4 \tag{13}$$

where $Pr(HMDA\ Race_i)$ is the probability that a buyer i is represented in HMDA as one of four races j (Asian (1), Black (2), Hispanic (3), White (4)); $\mathbf{X}_i$ is the vector of predicted race probabilities from MRI.

We convert the estimates from this model into Table 9 below. Each row contains the estimates associated with a 100% probability that the homebuyer is a given race (according to MRI). Within a row, the value in each column gives the probability that the homebuyer is identified as a given race in HMDA. We are most interested in how often a buyer would appear in HMDA as a given race, given a 100% probability that MRI predicts they are the corresponding race. These estimates are: 82.3% for Asian buyers, 79.1% for Black buyers, 85.2% for Hispanic buyers, and 95.6% for white buyers. The model's pseudo-R^2 is 0.41.

Table 9: Race Validation

	Predicted Probabilities: Pr(HMDA Race | 100% MRI prediction)			
	HMDA: Asian	HMDA: Black	HMDA: Hisp	HMDA: White
P(Asi)=1	0.823	0.008	0.036	0.133
P(Bla)=1	0.027	0.791	0.005	0.176
P(His)=1	0.012	0.009	0.852	0.127
P(Oth)=1	0.067	0.214	0.032	0.688
P(Whi)=1	0.011	0.006	0.027	0.956
N	*67967*	*67967*	*67967*	*67967*

In results not shown here, we re-estimate the model using only the most likely race from MRI as the predictive variable. The corresponding estimates are 69% (Asian), 33% (Black), 78% (Hispanic), 91% (White).

The pseudo-R^2 is 0.38 – the global metric of fit falls relatively little because the dataset is predominantly comprised of white individuals, but the ability to predict a specific minority race is considerably worse.

A.2 Repeat-Sales Method Results

In this section, we present the results of a quick exercise to evaluate how sales premia would differ by race if we adopted a repeat sales framework using only transaction-level data about winning bidders. We merge HMDA data to CoreLogic mortgage data going back to the year 2000, then merge this information onto CoreLogic transactions data. We estimate the following equation, comparable to Equation 1 in Bayer et al. (2017):

$$ln(p_{ijnt}) = race_{ijt} + X_{ijt}\beta + \mu_j + \theta_{nt} + \epsilon_{ijnt} \tag{14}$$

In our data, we see an individual i buy house j in census tract n in year t. We code the race of homebuyer i using HMDA data, following the scheme from the Data Appendix in Bayer et al. (2017) (this scheme is different from the one used in Appendix A.1, since bidding data from CTMe is not used here). Property j is identified consistently over time using the proprietary property key in CoreLogic. CoreLogic also provides us with census tracts n and year of sale t.

Race is coded as dummy variables for the buyer being Black, Hispanic, Asian/Other, or white (omitted category). Buyer characteristics X_{ijt} include income, down payment percentage, and their squares, all sourced from HMDA data. Fixed effects μ_j control for time-invariant characteristics of the property, and θ_{nt} control for tract-year shocks.

Table 10 shows three different specifications estimated over three different sample periods. In each triplet of columns, the first includes only property and year fixed effects, the second includes property and tract-year fixed effects, and the third adds the buyer characteristics in X_{ijt}.

The first triplet of columns is estimated using home sales data between 2000-2007, which overlaps with the sample period in Bayer et al. (2017). Columns (1) and (2) are comparable to the Columns (3) and (4) in each panel of Table 2 in that paper, and Column (3) is comparable to Column (2) in Table 3. We find comparable price premia for Black and Hispanic buyers in Column (2): 1.85% for Black buyers and 1.77% for Hispanic buyers, compared to multi-city estimates of 1.6% and 1.7% in Bayer et al. (2017). These premia grow with the inclusion of buyer fixed effects, again mirroring the trend in our comparison paper.

The second triplet of columns repeats the estimation using home sales data between 2000-2013, and finds the race premia generally stable. But when we run the same estimation using sales data from 2014-2021, our estimates of racial premia fall by orders of magnitude, with large standard errors. In results not

shown, we run another triplet of specifications restricted to sales between 2014 and March 1, 2020, in order to rule out confounding effects due to the Covid pandemic and associated housing market conditions. The results are generally similar and show no evidence of racial premia.

Table 10: Replication: Repeat Sales Methodology

	2000-2007			2000-2013			2014-2021		
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
Black	0.0163***	0.0185***	0.0207***	0.0229***	0.0190***	0.0204***	0.0006	0.0018	0.0030
	(0.0046)	(0.0044)	(0.0045)	(0.0036)	(0.0031)	(0.0032)	(0.0025)	(0.0023)	(0.0022)
Hispanic	0.0224***	0.0177***	0.0197***	0.0214***	0.0103***	0.0122***	0.0012	-0.0005	0.0004
	(0.0025)	(0.0024)	(0.0025)	(0.0019)	(0.0017)	(0.0017)	(0.0013)	(0.0012)	(0.0012)
Asian and Other	-0.0006	0.0014	0.0018	-0.0019	-0.0006	-0.0003	0.0014	0.0030***	0.0026**
	(0.0022)	(0.0021)	(0.0021)	(0.0016)	(0.0014)	(0.0014)	(0.0011)	(0.0011)	(0.0010)
Income			0.0001***			0.0002***			0.0002***
			(0.0000)			(0.0000)			(0.0000)
Income^2			-0.0000**			-0.0000***			-0.0000***
			(0.0000)			(0.0000)			(0.0000)
Down payment pct.			0.0270***			0.0230***			0.0171***
			(0.0024)			(0.0019)			(0.0065)
DPP^2			-0.0052***			-0.0001			0.0275**
			(0.0013)			(0.0009)			(0.0108)
Constant	12.3484***	12.3475***	12.3280***	12.3818***	12.3833***	12.3540***	12.8010***	12.8002***	12.7605***
	(0.0007)	(0.0007)	(0.0020)	(0.0005)	(0.0004)	(0.0014)	(0.0004)	(0.0004)	(0.0013)
N	47449	46933	41953	120054	119420	111002	97953	97660	94507
Year	YES	NO	NO	YES	NO	NO	YES	NO	NO
Prop	YES	YES	YES	YES	YES	YES	YES	YES	YES
Tract-Year	NO	YES	YES	NO	YES	YES	NO	YES	YES
Buyer Chars.	NO	NO	YES	NO	NO	YES	NO	NO	YES

Standard errors in parentheses

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

Notes: All regressions estimated using CoreLogic transactions data merged with HMDA data. Outcome variable for all regressions is log of sales price. Each grouping of three columns contains estimates using different subsamples of the data, differing by time period: 2000-2007 (Cols. 1-3), 2000-2013 (Cols. 4-6), 2014-2021 (Cols. 7-9). In each group of three, the first column includes fixed effects for property and year, the second column includes fixed effects for property and Census tract-year combination, and the third column adds buyer characteristic controls to the second column. Unlike tables for most of the rest of the paper, race is defined not probabilistically, but as a numeric indicator. Race codification is designed to match that of Bayer et al. (2017).