Journal of Urban Economics 102 (2017) 91–105

Contents lists available at ScienceDirect

Journal of Urban Economics

journal homepage: www.elsevier.com/locate/jue

Racial and ethnic price differentials in the housing market[☆]

Patrick Bayer[a,*], Marcus Casey[b], Fernando Ferreira[c], Robert McMillan[d]

[a] Duke University, United States, and NBER
[b] University of Illinois at Chicago, United States
[c] The Wharton School, University of Pennsylvania, United States, and NBER
[d] University of Toronto, Canada, and NBER

ARTICLE INFO

Article history:
Received 26 January 2017
Revised 21 June 2017
Available online 28 July 2017

JEL codes:
H0
J15
K4
R2
R3

Keywords:
Price differentials
Housing market
Racial differences
Racial prejudice
Repeat-sales approach

ABSTRACT

Do minorities pay more than whites for similar housing? We revisit this important question using a rich new dataset that covers two million repeat-sales housing transactions drawn from four major metropolitan areas. Our analysis applies a repeat-sales framework, including house and neighborhood-by-time fixed effects to control for unobserved differences in the quality of homes and their associated neighborhoods. In contrast to most of the recent literature, we find that black and Hispanic homebuyers pay premia of around 2% on average across the four cities – differences not explained by variation in buyer income or access to credit. We also show black and Hispanic buyers pay more for housing regardless of the race or ethnicity of the seller, suggesting that the estimated premia are unlikely to be driven by a very direct form of racial prejudice. Our estimates have implications for the levels and persistence of racial differences in home ownership, the segregation of neighborhoods, and the dynamics of wealth accumulation.

© 2017 Published by Elsevier Inc.

1. Introduction

Whether blacks and Hispanics pay more than whites for similar houses remains a central concern in the literature studying discrimination in housing markets. For much of the 20th Century, centralized (institutional) forms of racism such as mortgage redlining and restrictive covenants[1] not only gave rise to substantial racial segregation in many American cities but also created housing shortages, and higher prices, in predominantly black neighborhoods within those cities (Hirsch, 1978). In line with this view of the effects of centralized discrimination, the pioneering empirical study of Kain and Quigley (1975) found that blacks paid more than whites for housing on opposite sides of boundaries dividing segregated neighborhoods in St. Louis in 1970, and Cutler et al. (1999) established that blacks paid especially high housing prices in the most segregated cities in 1950.

By the end of the 20th Century, several landmark legal changes helped to diminish the force of centralized discrimination significantly.[2] As a consequence, whether black and Hispanic households pay more for comparable housing is no longer primarily a matter of forced shortages and higher prices in segregated minority neighborhoods but is instead a question of whether *decentralized* discrimination – by real estate agents, landlords, or sellers – serves to drive up prices for black and Hispanic homebuyers. As Becker (1971) showed, any form of discrimination (whether centralized or decentralized) that effectively limits the set of housing choices available to minorities should raise the prices they pay in equilibrium. Yet, contrary to Becker's prediction, most recent empirical research has found the opposite result – that whites pay more for

☆ We would like to thank Jeff Brown, Hanming Fang, Robert Kaestner, Ron Laschever, Darren Lubotsky, Marie Mora, Ben Ost, Steve Rivkin, Jake Vigdor and seminar participants at the Harris School of Public Policy, the University of Illinois, New College of Florida, Purdue University, Sciences Po, the University of Windsor, and the AEA and IEA Meetings for helpful comments. Financial support is gratefully acknowledged from the Research Sponsor Program of the Zell/Lurie Real Estate Center at Wharton, SSHRC, and the National Science Foundation through grant SES-0721136. Any opinions, findings, conclusions, or recommendations expressed in this paper are those of the authors and do not necessarily reflect the views of the National Science Foundation. All remaining errors are our own.

* Corresponding author.
E-mail address: patrick.bayer@duke.edu (P. Bayer).

[1] See Arrow (1998) for a personal reflection.

[2] That the role of centralized racism had diminished by the end of the 20th Century is supported by the findings of Cutler et al. (1999), who show that the relationship between segregation and any racial price premia had effectively disappeared by 1990.

http://dx.doi.org/10.1016/j.jue.2017.07.004
0094-1190/© 2017 Published by Elsevier Inc.

comparable housing.[3] Properly measuring housing quality (which is only coarsely captured by observed attributes) remains an ongoing challenge in this literature, however, and it is unclear whether these findings persist after accounting for unobserved differences in housing quality.

Taking up that challenge, the primary goal of this paper is to revisit the extent of racial and ethnic price differentials in the housing market, making use of a new panel dataset that allows us to credibly isolate variation in sales prices for the same properties within neighborhoods. The dataset is based on all housing transactions that occurred in four major metropolitan areas – Chicago, Baltimore/Washington DC ("Maryland" for short), Los Angeles, and San Francisco – over a period spanning almost two decades, 1990–2008. For each transaction, we observe precise information about location and housing attributes, in addition to the sales price. We then merge information relating to the buyer's race, ethnicity, and income with these transactions data by matching comprehensive mortgage records, creating a rich dataset that includes over two million repeat-sale transactions. Further, for over three quarters of a million houses in our dataset that sell more than once, we can also establish the race and ethnicity of the seller.

With these data in hand, we estimate how the sales price varies with the race or ethnicity of the buyer using a repeat-sales approach. This controls explicitly for the unobserved quality of the individual house through the inclusion of house fixed effects, thereby isolating racial/ethnic differences in sales prices for comparable homes.[4] Our rich dataset also allows us to control for time-varying neighborhood attributes in a flexible way using neighborhood-by-time fixed effects.

In contrast to virtually all recent research, our baseline results indicate that black and Hispanic buyers pay a premium for comparable housing relative to white buyers in all four major city markets.[5] These premia range between 1 and 3%; the average estimated premium pooling across all four metropolitan areas is around 1.6% for black buyers and 1.7% for Hispanic buyers. Due to the large sample sizes involved, all of these estimates are very precise and significantly different from zero.

The novel features of our dataset allow us to assess whether these race/ethnic differentials might be attributable to other buyer attributes affecting the home sales process. For each transaction, we observe the buyer's income reported on the mortgage application as well as the down payment associated with the mortgage and the lender's name. When we include controls for buyer's income and down payment in the repeat-sales specification described above, the estimated premia for black and Hispanic buyers increase, to an average of 2.1% for blacks and 2.2% for Hispanics. When we add a complete set of lender fixed effects, which might proxy for the buyer's access to credit, the results in each city remain very similar, with the average premia remaining at around 2%. Thus it appears that the correlation of race and ethnicity with financial considerations leads, if anything, to an underestimate of the black and Hispanic premia in our initial analysis.[6]

Having established the existence of robust racial and ethnic differentials in purchase prices, we next assess how these differentials vary with the race and ethnicity of sellers. If racial animosity/prejudice is an important factor influencing seller behavior, we might expect sellers to favor buyers of their own race or ethnicity in a systematic way (although the strength of this prediction is mitigated by the presence of real estate agents and the potentially limited information sellers might have about the buyer). Against that, our results indicate that black and Hispanic buyers pay a similar premium when buying from sellers of each race and ethnicity, and that (if anything) these premia are greater when blacks buy from blacks and Hispanics from Hispanics. These results help to rule out an especially direct form of racial prejudice, or any explanation that would lead sellers to favor buyers of their own race/ethnicity, as the primary explanation for the racial price differentials we find.

Our analysis cannot isolate a single alternative explanation for the estimated price premia paid by black and Hispanic buyers in these markets. Statistical discrimination might be motivated, for example, by the correlation of race with search costs or experience in real estate bargaining; and further research is needed to determine the extent to which a particular channel predominates.[7] Regardless of the ultimate explanation, however, our results show that black and Hispanic buyers pay more than their white counterparts in almost every purchase setting, and that robust racial differences in the price paid to buy a home – around 2% on average in multiple large US markets – persist to the present day, long after the most overt forms of institutional discrimination have been eliminated. Price differences of this magnitude have implications more generally for the levels and persistence of racial differences in home ownership, neighborhood segregation, and the dynamics of wealth accumulation.

The remainder of the paper is organized as follows: Section 2 describes the dataset that we have constructed for our analysis. Section 3 introduces our primary research design and relates it to those used in previous studies seeking to infer racial price differentials in housing markets. We present our main estimates in Section 4, and in Section 5, explore whether racial prejudice might provide the primary explanation for these results. Section 6 concludes, drawing out some implications of our analysis.

2. Data

The dataset we have assembled combines information from proprietary transactions data collected by a real estate monitoring service, *DataQuick*, with publicly-available loan application registry information gathered under the Home Mortgage Disclosure Act (HMDA).[8] The transactions dataset includes a complete census of housing transactions and is available for the San Francisco and Los Angeles metropolitan areas from 1990 to 2008, and for the Chicago metropolitan area and Baltimore-Washington DC corridor in Maryland from 1997 to 2008. These cities are particularly useful for studying racial and ethnic price differences because each has a large and heterogeneous population of homebuyers. Data on the transaction price, date of sale, loan amount, lender name, and house location are provided for each transaction. In addition, and important for our analysis, each property is characterized by a unique identifier that makes it possible to track the longitudinal transactions history of each home.

We match demographic information to individual homes by using the HMDA application registry files. The HMDA legislation was enacted to monitor potential redlining and discriminatory lending behavior. An important feature of this legislation is the

[3] See, for example, Follain and Malpezzi (1981), Chambers (1992), and Kiel and Zabel (1996).

[4] The study by Myers (2004) uses a similar approach, and is perhaps the closest in the literature to our paper. We discuss that work in some detail in Section III, along with other relevant prior studies.

[5] In addition to Myers (2004), one exception is the study by Ihlanfeldt and Mayock (2009). In their preferred specifications, they find evidence of racial/ethnic price premia, in line with the estimated premia in our study.

[6] We obtain qualitatively similar results in a host of additional specifications designed to test the robustness of these main findings in the face of concerns about the reliability of the data.

[7] A recent alternative strand of literature, starting with Harding et al. (2003), seeks to control for differences between buyers and sellers in terms of their (unobserved) bargaining power. We describe findings from that line of research below.

[8] Detailed information about the HMDA legislation and the public-use data can be found at http://www.ffiec.gov.

Table 1
Summary characteristics of the repeat-sales transactions dataset.

	Chicago	San Francisco	Maryland	Los Angeles	**All-City**
Number of observations	382,389	535,286	278,221	925,622	2,121,518
Repeat sales (proportions)					
Sold twice	0.75	0.57	0.72	0.55	0.61
Sold three times	0.22	0.32	0.26	0.36	0.31
Sold four or more times	0.03	0.11	0.02	0.09	0.08
Buyer race (proportions)					
White	0.75	0.61	0.62	0.64	0.65
Black	0.09	0.04	0.25	0.06	0.09
Hispanic	0.08	0.06	0.05	0.11	0.08
Asian and other	0.08	0.29	0.08	0.19	0.18
Transaction type (proportions)					
White-to-white	0.60	0.47	0.52	0.43	0.48
White-to-black	0.04	0.02	0.13	0.02	0.04
White-to-Hispanic	0.04	0.07	0.05	0.12	0.08
Black-to-black	0.03	0.01	0.08	0.01	0.02
Black-to-white	0.02	0.01	0.06	0.02	0.02
Black-to-Hispanic	<0.01	0.03	<0.01	0.01	<0.02
Hispanic-to-Hispanic	0.02	0.03	0.01	0.10	0.06
Hispanic-to-white	0.05	0.04	0.04	0.06	0.05
Hispanic-to-black	0.01	0.01	0.01	0.01	0.01
Other types	0.18	0.31	0.09	0.22	0.22
Transaction statistics ($)					
Mean transaction price	220,737	380,258	227,944	306,974	299,557
Median transaction price	182,500	312,500	190,000	239,000	240,935
Mean income	86,558	120,489	88,188	106,434	104,005
Median income	67,000	99,000	71,000	81,000	81,707
Median down payment	27,500	78,000	22,000	48,500	48,683

Notes: This table provides summary statistics for the main dataset used in the analysis. Data for Chicago and Maryland span the years 1997–2007, and for Los Angeles and San Francisco, the period 1990–2007. 'Transaction type' refers to the race/ethnicity of those involved in the transaction (e.g. 'White-to-black' refers to a white seller transacting with a black buyer). The transaction price is the actual recorded closing price of the home. 'Income' refers to reported income on the mortgage application. The bottom two panels of the table characterize transactions for all houses that sell a minimum of two times in the four metropolitan areas.
(*Sources*: *DataQuick* and HMDA data.)

requirement that race/ethnicity and other pertinent demographic information be recorded for each mortgage application. Hence, we are able to retrieve from these data the race and income of the buyer, transactions date, and the census tract of the home in question for linkage purposes.

Our housing transactions data for each city are then matched with the HMDA dataset using a sequential procedure on the basis of four key variables: (i) census tract, (ii) loan amount, (iii) transaction year, and (iv) lender name.[9] Overall, the match rate is high, as shown in Data Appendix Table 1 – around 67% overall across the four cities, with the percentage of homes matched in the four cities ranging between 64% (Los Angeles) to 71% (Chicago). A fraction of homes in the transactions dataset – around 13% overall – could not be matched because they were purchased with cash. Among houses purchased with a home loan, match rates are much higher. Around two thirds of such homes can be matched uniquely on the basis of the four key matching variables (so-called "Perfect matches" in the Data appendix table), with city-level 'perfect' match rates ranging between 64 and 72%. Empirical results presented below correspond to a dataset based on all matched homes, though our baseline results are shown to be robust when restricting to perfect matches.

An observation in the matched dataset is a transaction involving a particular house at a given time. It includes the full set of information provided for each transaction record by *DataQuick*,[10] along with demographic and economic information about the buyer drawn from HMDA. Further, tract-level Census data are merged into the *DataQuick*-HMDA matched dataset to augment the neighborhood-level characteristics available for our empirical analysis.

The main estimating equation, described below, will include house fixed effects. Thus, the effective sample for our study, which uses a repeat-sales approach, consists of houses that sell at least twice during the period.[11] Our repeat-sales sample includes over quarter of a million transactions for Maryland, over a third of a million for Chicago, over half a million transactions for San Francisco, and close to one million for Los Angeles, giving an overall sample size of more than two million.

Table 1 presents descriptive statistics for the observations in our repeat-sales sample relating to the attributes of buyers and sellers, as well as housing transactions.[12] The table's columns report statistics separately for each of the four metropolitan areas and for the "All-City" sample. The top panel in the table gives the number of observations in the repeat-sales dataset and the frequency with which homes in each area sold two, three, or four or more times, respectively – the three categories are exhaustive. In all four cities, houses that sold exactly twice make up easily the majority of observations. The second panel reports the distribution

[9] See the Data appendix for a detailed discussion of this procedure, including various steps taken to clean the raw data. We also provide further details regarding the resulting match quality.

[10] Data Appendix Table A2 provides descriptive statistics comparing observable characteristics of homes that could be matched (versus unmatched homes) – because of data constraints, just for Los Angeles and San Francisco. The table shows reasonable balance across matched and unmatched samples.

[11] Our main sample excludes properties that have had major renovations based on a flag in the *DataQuick* assessor data. (See the Data appendix for more details.)

[12] Appendix Table A1 compares demographic and house characteristics for the repeat-sales and overall transaction datasets, showing the repeat-sales sample to be representative of the overall transactions sample, both in terms of buyer demographic characteristics and statistics on housing transactions.

of buyer race and ethnicity. While the majority of buyers in each city are white, there is considerable cross-city variation in buyer race/ethnicity. For example, while only 4 percent of buyers are black in San Francisco and 6% in Los Angeles, blacks constitute 25% of buyers in Maryland. The fraction of Hispanic buyers is contrastingly low there, at 5%, and higher – at 11% – in Los Angeles, while the fraction of Asian (including 'Other' race/ethnicity) buyers is by far the highest in San Francisco, at 29%.

For homes that sell multiple times, we are able to characterize the race and ethnicity of the seller, beginning with the second transaction observed in the dataset.[13] The third panel of Table 1 reports transaction type – the distribution of sales in each metropolitan area jointly by the race/ethnicity of the buyer and seller for all transactions where both buyer and seller race/ethnicity are recorded. In each metropolitan area, white-to-white transactions are easily the most common pairing – between 43 and 60% of the total across the four cities. There is, however, considerable variation in the composition of the remaining transactions in each metropolitan area. Maryland has the highest fraction of inter-racial transactions involving blacks and whites: white-to-black and black-to-white transactions make up 13 and 6% of repeat sales, respectively. In Chicago, white-to-black sales account for 4% of the total while black-to-white transactions are less common (2% of all sales). Transactions between whites and Hispanics, running in each direction, are reasonably common in Chicago, Maryland, and San Francisco and are still higher in Los Angeles.

The bottom panel of Table 1 reports financial statistics relating to housing transactions by metropolitan area for the repeat-sales sample.[14] It shows that home prices are significantly higher, not surprisingly, in the two Californian metropolitan areas, as are the average incomes and median down payments of homebuyers.

3. Research design: identification of racial and ethnic price differentials

The research design that we implement makes use of the unique structure of our dataset to shed light on racial and ethnic differences in the price paid to purchase comparable housing. Our primary approach can be summarized using the following regression equation, which relates the log price of transaction i for house j in neighborhood n at time t to the buyer's race/ethnicity as well as a set of controls:

$$ln\left(p_{ijnt}\right) = raceeth\_buyer_{it}\gamma + X_{ijt}\beta + \mu_j + \theta_{nt} + \varepsilon_{ijnt}. \tag{1}$$

Here, $raceeth\_buyer_{it}$ is a vector of indicators for the race/ethnicity of the buyer at time t, with white being the omitted category; X_{ijt} denotes a vector of observable house buyer characteristics for transaction i involving house j at time t; μ_j is a house-specific fixed effect, and θ_{nt} denotes a set of neighborhood-by-time fixed effects. The coefficient γ measures the average premium, if any, paid by black and Hispanic buyers respectively relative to whites.[15]

Our goal is to measure any difference in the prices that buyers of different races/ethnicities pay for comparable housing. As we make clear in what follows, the inclusion of house fixed effects, μ_j, in Eq. (1) ensures that these parameters are identified by comparing prices for houses for which buyer race or ethnicity changes over successive transactions to those for which buyer race and ethnicity remains constant. The inclusion of neighborhood-by-time fixed effects ensures that these comparisons are made within the same neighborhood during the same time period.

To see how the effect of buyer's race/ethnicity is identified in this model, it is helpful to re-write Eq. (1) by differencing observations for a consecutive pair of transactions (i, t) and (i', t') involving house j:

$$\begin{aligned} ln\left(p_{ijnt}\right) - ln\left(p_{i'jnt'}\right) &= (raceeth\_buyer_{it} - raceeth\_buyer_{i't'})\gamma \\ &\quad + \left(X_{ijt} - X_{i'jt'}\right)\beta + (\theta_{nt} - \theta_{nt'}) + \omega_{ijnt} \end{aligned} \tag{2}$$

Notice that the house fixed effect drops out of Eq. (2) and that γ now multiplies the difference in buyer race for transactions i and i'. This term is, of course, equal to zero if the buyer's race does not change between transactions, and thus those transactions contribute to the identification of the θ-parameters, which characterize the pattern of price appreciation in neighborhood n. When buyer race/ethnicity *does* change over successive transactions, the vector of differenced buyer race/ethnicity variables is non-zero and the γ vector is identified by comparing the price appreciation of these transactions relative to the baseline rate of appreciation in the neighborhood.

By way of a concrete example, consider a pair of neighboring houses that both sold initially to white buyers in 1999 and which both sell again in 2006, though this time to a black and white buyer, respectively. In this instance, the estimated racial difference in prices would be identified by the difference in price appreciation between the two sales. If, for instance, the price paid by the white buyer in 2006 implied an appreciation rate of 50% and the price paid by the black buyer implied an appreciation rate of 55%, we would infer that the black buyer paid a 5% premium. Averaging these differences across the full set of comparable houses in the dataset provides the basis for identifying the buyer race/ethnicity parameters in Eqs. (1) and (2).

A potential concern with the research design captured by Eq. (1), and one that applies to the entire existing literature, is that the homes purchased by buyers of different races may have experienced different rates of appreciation within the same neighborhood or undergone differential amounts of renovation or maintenance during the holding period of the previous owner. If buyers of one race, for example, tend to buy houses that are more likely to have been improved over the previous holding period, this would bias the buyer race parameters. In essence, the appreciation associated with these improvements would be mistakenly attributed to the race of the buyer. Because white buyers have significantly higher levels of income and wealth than black and Hispanic buyers, we would generally expect whites to be systematically more likely to buy improved or especially well-maintained houses – that is, improved or well-maintained in ways not observed by the econometrician.[16] If this were indeed the case, the estimated γ parameters would *understate* the actual premia paid by blacks and Hispanics relative to whites. We address this potential concern below by using the detailed assessor records associated with each housing transaction in the dataset.[17]

[13] The seller's race/ethnicity is not observed for the first transaction in our repeat-sales sample; for the second transaction, it is equal to the buyer's race/ethnicity for the first transaction (and so on for the third etc.).

[14] Appendix Table A2 offers a further comparison of house characteristics, for homes in the repeat-sales sample that sell twice versus three or more times.

[15] We also control for Asian and 'Other' races, though interpreting the corresponding coefficient is less clear because of the broad set of ethnicities falling under that umbrella. Those results are available upon request from the authors.

[16] Another possibility is that wealthier white sellers may be more likely to have improved their properties, leading to an upward bias in the estimates of the premia that white sellers receive in the analysis below, conditioning on seller race/ethnicity.

[17] In particular, the records for most homes include an assessor flag that indicates whether there has been a major renovation of the home and when such a renovation occurred. We drop homes involving major renovations from the main sample. Doing so has a negligible impact on the parameter estimates, indicating that there is little within-neighborhood correlation between race/ethnicity and major housing improvements/maintenance to confound our estimates. (Estimates based on a sample that retains such homes are available on request.)

3.1. Relation to prior studies of price differentials in the housing market

To estimate price differentials in the housing market, ideally one would (as noted in the Introduction) like to compare the price that buyers of different races/ethnicities would pay for identical properties. Much of the previous literature has pursued this strategy by comparing the prices of homes purchased in the same neighborhood at around the same time. The pioneering studies that developed and implemented this approach (King and Mieszkowski, 1973; Kain and Quigley, 1975; Yinger, 1978) found that minorities paid a premium for comparable housing. In contrast, subsequent analyses based on larger samples, notably by Follain and Malpezzi (1981), Chambers (1992) and Kiel and Zabel (1996), found statistically significant *discounts* for black buyers relative to whites. These latter studies were still constrained by the size and nature of the datasets available to researchers at the time, limiting the researchers' ability to control for unobserved differences in house quality within neighborhoods. Given their higher average levels of income and wealth, it would not be surprising if white buyers generally purchased higher quality housing within the same neighborhood, potentially leading to a spurious finding that white buyers pay a 'premium' for housing. By including house fixed effects in Eq. (1), we address this issue directly.

In this regard, our approach is related to specifications reported in Myers (2004) that include house fixed effects, estimated with data from three waves of the American Housing Survey, and supplemented with special neighbors' samples. Myers finds evidence that blacks pay premia for their homes of around 10% relative to whites, although the estimates are significant only at the 10% level.[18] Another recent study, by Ihlanfeldt and Mayock (2009), examines price differentials by race, first estimating a 'traditional' specification that controls for detailed structural characteristics and block-group fixed effects using data from 20 metropolitan areas in Florida. As in much of the prior literature using that general approach, Ihlanfeldt and Mayock find evidence of price discounts for non-white buyers, and argue that the correlation between race and unobserved property characteristics may bias the estimated race coefficients downwards.[19] It is precisely this potential bias that the inclusion of house fixed effects in our analysis is designed to eliminate.

Relative to Myers (2004), our analysis offers two main advantages. The first is that we have a *complete* census of housing transactions in major metropolitan areas over a long period of time, including many houses that sell to buyers of different races or ethnicities. Large numbers of such 'switches' in the race/ethnicity of the buyer are needed to identify racial/ethnic differences in pricing when house fixed effects are included in the analysis. We effectively have multiple orders of magnitude more data, and this allows us to estimate racial/ethnic price differences far more precisely and to characterize how they vary across sellers, neighborhoods, and market conditions. The second main advantage of our research design involves the inclusion of neighborhood-by-time fixed effects that effectively control for time-varying factors in each neighborhood that might influence prices. Again, it is the sheer size and scope of our dataset that allows us to control for such neighborhood dynamics in this especially flexible way.

4. Main results

We now present the main results of our analysis. For contrast, we begin by reporting estimates of versions of Eq. (1) that leave out the house and/or neighborhood-by-time fixed effects. Comparing our final model (1) with such restricted specifications provides a clear demonstration of the importance of the additional controls when estimating racial/ethnic differences in purchase prices for comparable housing. We then add controls that proxy for other aspects of the buyer's financial position – income, down-payment, and lender fixed effects – to examine whether the differences estimated in our baseline specification might be proxying for these financial considerations, which are correlated with race/ethnicity, rather than the buyer's race/ethnicity itself.

4.1. Baseline results

Table 2 reports baseline results for black-white and Hispanic-white price differentials from the estimation of Eq. (1). The table reports results for each metropolitan area and for all four cities combined, organized into two panels: Panel A presents the relevant black premium (or discount) coefficient and Panel B presents the Hispanic coefficient for each specification. In each case, the coefficient measures the difference in transaction price relative to a non-Hispanic white buyer.

Each panel reports results from four specifications in turn. The first controls for neither house nor neighbor-by-time fixed effects. Not surprisingly, the parameter estimates in this case are uniformly negative and large in magnitude, as black and Hispanic buyers purchase substantially less expensive properties than white buyers in each metropolitan area.

The second column of each panel reports parameter estimates for a specification that includes neighborhood-by-time fixed effects but does not include house fixed effects. This specification compares more closely to the specifications typically reported in much of the existing literature and as found there, the parameter estimates reported in column (2) remain negative in almost every instance, the precise zero for Los Angeles being the exception. Pooling across the cities, the estimated black-white and Hispanic-white price differentials for the full sample are −0.015 and −0.003, respectively, giving no indication that black or Hispanic buyers pay more for *comparable* housing than their white counterparts.

The specification reported in the third column of Table 2 includes house fixed effects, while the one reported in the fourth column controls simultaneously for both house and neighborhood-by-time fixed effects. Comparing the parameter estimates reported in column (4) – our preferred specification – in each panel with those in columns (1) and (2), it is immediately obvious that the inclusion of house fixed effects flips the sign on the estimated price differentials from negative to positive in *every* sample for both black and Hispanic buyers. The estimated price differentials for black and Hispanic buyers in fourth column for the All-City sample are 0.016 and 0.017, respectively.[20] There is some variation in the estimated price differentials across cities. The estimated black-white price differential is highest in Chicago (2.9%), and the Baltimore area (1.6%), while estimates range between 1.1–1.3% in

[18] Myers' sample consists of just under 22,000 observations. She notes that a larger sample would increase the precision of fixed effects estimates reported in the paper, as more changes in the race/ethnicity of home owners would then be likely.

[19] Because of the likely bias associated with the traditional approach, Ihlanfeldt and Mayock (2009) then implement as their preferred estimator an approach put forward in Harding et al. (2003) designed to account for differences in buyer and seller bargaining power. Applying that approach using their Florida data, they find evidence of racial price premia; our findings below are consistent with their estimates. In Section 6, we discuss findings reported in the Harding *et al.* paper, as they are relevant when interpreting our main results.

[20] Due to computational constraints associated with the size of the overall sample and the two-dimensional nature of the fixed-effect models, we construct the pooled all-city estimates and associated standard errors for our preferred specifications by combining the individual city estimates using the variance-weighted least squares approach described in Becker and Wu (2007).

Table 2
Racial/ethnic housing price differentials – baseline results.

	Panel A: Black-white differential				*Panel B: Hispanic-white differential*				
Location	(1)	(2)	(3)	(4)	(1)	(2)	(3)	(4)	Obs.
Chicago	−0.491	−0.007	0.014	0.029	−0.231	−0.013	−0.010	0.015	382,389
	[0.002]	[0.006]	[0.005]	[0.005]	[0.002]	[0.005]	[0.005]	[0.003]	
San Francisco	−0.392	−0.021	−0.001	0.011	−0.241	−0.014	0.034	0.028	535,286
	[0.003]	[0.004]	[0.004]	[0.003]	[0.002]	[0.004]	[0.003]	[0.003]	
Maryland	−0.294	−0.025	0.012	0.016	−0.12	−0.004	0.037	0.024	278,221
	[0.002]	[0.006]	[0.005]	[0.003]	[0.003]	[0.006]	[0.006]	[0.005]	
Los Angeles	−0.381	−0.011	−0.002	0.013	−0.294	0.002	0.006	0.011	925,622
	[0.002]	[0.003]	[0.002]	[0.002]	[0.002]	[0.002]	[0.002]	[0.002]	
All-City	**−0.389**	**−0.015**	**0.001**	**0.016**	**−0.238**	**−0.003**	**0.010**	**0.017**	**2,121,518**
	[0.004]	**[0.005]**	**[0.004]**	**[0.003]**	**[0.005]**	**[0.004]**	**[0.005]**	**[0.003]**	
Additional controls included:									
Time	Yes	No	Yes	No	Yes	No	Yes	No	
House fixed effects	No	No	Yes	Yes	No	No	Yes	Yes	
Tract x time fixed effects	No	Yes	No	Yes	No	Yes	No	Yes	

Notes: Cluster-robust standard errors in brackets. The estimates for each city are derived from a regression of log transaction price on race indicators and other variables. The specifications reported show the impact of including a set of house fixed effects and/or a set of neighborhood-by-time fixed effects on the estimated differentials.

Table 3
Racial/ethnic housing price differentials – adding buyer and lender controls.

	Panel A: Black-white differential			*Panel B: Hispanic-white differential*			
Location	(1)	(2)	(3)	(1)	(2)	(3)	Obs.
Chicago	0.029	0.035	0.034	0.015	0.023	0.023	382,389
	[0.005]	[0.005]	[0.006]	[0.003]	[0.003]	[0.005]	
San Francisco	0.011	0.021	0.019	0.028	0.032	0.034	535,286
	[0.003]	[0.004]	[0.005]	[0.003]	[0.004]	[0.005]	
Maryland	0.016	0.026	0.024	0.024	0.032	0.029	278,221
	[0.004]	[0.004]	[0.006]	[0.005]	[0.005]	[0.006]	
Los Angeles	0.013	0.017	0.016	0.011	0.015	0.013	925,622
	[0.002]	[0.002]	[0.002]	[0.002]	[0.002]	[0.002]	
All-City	**0.016**	**0.021**	**0.020**	**0.017**	**0.022**	**0.019**	**2,121,518**
	[0.003]	**[0.003]**	**[0.004]**	**[0.003]**	**[0.003]**	**[0.004]**	
Additional controls included:							
House fixed effects	Yes	Yes	Yes	Yes	Yes	Yes	
Tract x time fixed effects	Yes	Yes	Yes	Yes	Yes	Yes	
Buyer attributes	No	Yes	Yes	No	Yes	Yes	
Lender fixed effects	No	No	Yes	No	No	Yes	

Notes: Cluster-robust standard errors in brackets. The estimates are derived from a regression of log transaction price on race indicators and other variables. The 'Buyer attributes' include buyer income and buyer down payment percentage and squares of these variables. 'Lender fixed effects' consist of indicators for large individual lenders in each market and one indicator for small lenders. The All-City estimates are constructed using the procedure described in Becker and Wu (2007). The column (1) estimates in each panel are taken from the corresponding column (4) of Table 2.

the other two cities. The estimated Hispanic-white price differential is in the 2.4–2.8% range for Maryland and San Francisco, 1.5% for Chicago, and 1.1% for the Los Angeles area.

4.2. Controlling for other buyer attributes

The baseline results presented in Table 2 provide strong statistical evidence that black and Hispanic buyers pay more than whites do for comparable housing. It could be, however, that race is simply correlated with other buyer attributes that affect the home sales process. To assess this possibility, Table 3 presents results from additional specifications that control for various aspects of the buyer's financial position, including buyer income, the down payment percentage, and a set of lender fixed effects. As mentioned previously, the buyer's financial position might be correlated with the sales price for a number of reasons, including the ability to secure mortgage financing and differences in search costs. Because some lenders specialize in high-priced or subprime loans, the inclusion of lender fixed effects proxies for the buyer's credit-worthiness.

For reference, the first column in each panel of Table 3 repeats the specification in column (4) of Table 2, which controls for house and neighborhood-by-time fixed effects. The specification reported in column (2) adds controls for measures of the buyer's income and down payment percentage as well as second-order polynomial terms in these buyer attributes.[21] The coefficients for both the black-white and Hispanic-white differentials increase in a relatively uniform way for each city with the inclusion of these controls. As shown in the bottom row, the black-white gaps increase from 1.6 to 2.1% for the full sample, while the Hispanic-white gaps increase from 1.7 to 2.2%. Column (3) reports results for a specification that adds lender fixed effects to the specification in column (2). The inclusion of these lender controls has a negligible impact on the estimates in each sample, and the 'All-City' estimated price premia for blacks and Hispanics are similar and very precisely-estimated: 2% for blacks and 1.9% for Hispanics.[22] Taken as a whole, the results

[21] It is worth noting that the buyer's down payment percentage may be endogenous if, for example, black and Hispanic buyers tend to put down more than comparably wealthy whites to overcome potential discrimination on the part of sellers.

[22] Lender fixed effects do not provide a perfect proxy for the credit worthiness of the borrower, due to the substantial amount of within-lender variation and the possibility that minority borrowers are positively selected within-lender due to discrimination in the mortgage market. A good deal of useful information is

Table 4
Robustness of premia across different subsamples.

	Panel A: Black-white differential				Panel B: Hispanic-white differential			
	(1) Baseline	(2) Perfect matches only	(3) Pre-2004 transactions only	(4) Dropping all subprime	(1) Baseline	(2) Perfect matches only	(3) Pre-2004 transactions only	(4) Dropping all subprime
All-city estimate	0.021 [0.003]	0.024 [0.005]	0.019 [0.003]	0.023 [0.005]	0.022 [0.003]	0.023 [0.005]	0.018 [0.003]	0.017 [0.003]
Additional controls included:								
House fixed effects	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Tract x time fixed effects	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Buyer attributes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Observations	2,121,518	970,696	1,508,724	1,685,975	2,121,518	970,696	1,508,724	1,685,975

Notes: Cluster-robust standard errors in brackets. Estimates are derived from a regression of log transaction price on race indicators, individual controls, and various fixed effects. The relevant subsample is listed in the column heading. Estimates in column (1) of each panel are drawn from the corresponding column (2) of Table 3. Column (2) retains only transactions that can be matched perfectly (according to our matching procedure - see Data appendix). Column (3) uses only pre-2004 transactions, before the subprime boom. Column (4) drops all transactions originated by subprime lenders, as defined by the HUD Subprime and Manufactured lender list (see http://www.huduser.org/portal/datasets/manu.html).

presented in Table 3 provide clear and consistent evidence that the correlation of race and financial considerations does not lead to an over-estimate of the magnitude of racial price differentials.[23]

4.3. Robustness

Table 4 presents estimates for three additional specifications designed to explore the robustness of our main price-premium findings. The first column in each panel shows the baseline results (all-city estimates that include house fixed effects, tract-by-time fixed effects, and buyer attributes) from column (2) of Table 3. The second column restricts the sample to observations with ‘perfect’ *DataQuick*-HMDA matches. As explained in the Data section, these are the matches for which we are very confident about the race/ethnicity of buyers and sellers, given that we are able to match observations in *DataQuick* with HMDA uniquely, on the basis of census tract, year, lender's name and loan amount. Since about 55–60% of the overall housing transactions data are ‘high-quality’ matches with HMDA, the number of repeat-sales observations used in the specifications reported in column (2) drops, by more than 50%. Yet the estimates are still very precisely estimated, and are even larger in magnitude compared with the full repeat-sales sample: blacks and Hispanics now pay approximately 2.3 or 2.4% premia for comparable houses.

The next two columns in each panel shed light on whether our estimates are dependent on the composition of buyers and, in particular, whether they might be driven by the set of buyers drawn into the housing market by relatively easy access to credit in the housing boom of the mid-2000s. In the third column of each panel, we restrict the sample to pre-2004 transactions, which occurred before the peak years of the last housing cycle when subprime mortgage activity began to pick up. Even though point estimates are slightly smaller than the baseline, namely 0.019 for blacks and 0.018 for Hispanics, they are not statistically different from those based on the full sample. In column (4), we drop all subprime loans.[24] Again, the estimated magnitudes are quite similar to those in the baseline model, increasing slightly for the black-white gap and shrinking slightly for the Hispanic-white gap. Overall, these final two specifications indicate that the estimated price premium by race/ethnicity is not a recent occurrence driven completely by either the frenzy associated with the recent housing boom or subprime lending.[25]

Table 5
Heterogeneity in racial/ethnic price differentials based on initial neighborhood racial composition.

	All-city estimates Neighborhood percent white		
	Baseline: >0.0	> 0.5	>0.8
Black	0.021 [0.005]	0.028 [0.005]	0.029 [0.009]
Hispanic	0.022 [0.003]	0.018 [0.003]	0.010 [0.005]
Other buyer attributes	Yes	Yes	Yes
House fixed effects	Yes	Yes	Yes
Tract x time fixed effects	Yes	Yes	Yes
Observations	2,121,518	1,342,550	498,097

Notes: Cluster-robust standard errors in brackets. Estimates are derived from a regression of log transaction price on race indicators, a set of house fixed effects and a set of neighborhood-by-time fixed effects. ‘Other buyer attributes’ include income quartile indicators and a down payment percentage polynomial. ‘Baseline’ estimates in column (1) of each panel are drawn from the corresponding column (2) of Table 3. ‘> 0.5’ refers to a sample that includes neighborhoods initially greater than 50% non-Hispanic white. ‘> 0.8’ refers to neighborhoods initially greater than 80% non-Hispanic white.

4.4. Neighborhood demographic composition

Table 5 examines how the premium paid varies with the initial non-Hispanic white composition across neighborhoods, again reporting results for the All-City sample. Turning first to the pattern for black buyers, shown in the first row, the black premium remains high across the different categories of neighborhood (based on initial percent white), rising somewhat with the white percentage. The fact that the premium paid by black homebuyers remains

available in the across-lender variation, however, as substantial sorting occurs across lenders on the basis of credit-worthiness. In essence, our analysis considers whether such variation can explain the estimated racial and ethnic price differentials. If we had found that controlling for lender fixed effects cut these differentials significantly, for instance, this would be highly suggestive that the racial and ethnic price differentials were due primarily to issues related to credit-worthiness. But this is not what we find.

[23] A potential explanation for this finding is that individuals with higher levels of income and wealth have higher search costs due to having a higher value of time. As a result, conditioning on income and wealth – both highly correlated with race – would lead to larger estimates of racial/ethnic price premia.

[24] We define subprime loans based on the HUD classification of subprime lenders. See Ferreira and Gyourko (2015) for more information about this classification and other proxies for subprime borrowing.

[25] In an additional specification, we included controls for the order of transactions over the observation period. Including these controls had a small impact: the estimated premia were both qualitatively and quantitatively similar.

high suggests that it is likely a consequence of a market-wide mechanism.

A rather different pattern emerges for Hispanic buyers. In particular, the estimated Hispanic premium diminishes as we restrict attention to neighborhoods with higher fractions of non-Hispanic whites. One possible explanation for this declining pattern is that the Hispanic buyers in predominantly white neighborhoods may be systematically different from those who purchase homes in more racially diverse neighborhoods. Recent immigrants, for example, may be less likely to buy homes in predominantly white neighborhoods.

5. Heterogeneity by race and ethnicity of seller

Next, we present a series of results that characterize how the black and Hispanic price differentials we have estimated vary by the race and ethnicity of the seller.[26] We begin by decomposing the estimated premia by the race/ethnicity of both the buyer and seller.[27] To that end, we expand Eq. (1) to include a full set of interactions between buyer and seller race/ethnicity:

$$ln\left(p_{ijnt}\right) = (raceeth\_seller \times raceeth\_buyer)_{it}\lambda + X_{ijt}\beta + \mu_j + \theta_{nt} + \varepsilon_{ijnt}, \quad (3)$$

reporting all of the estimated premia, λ, relative to white-to-white transactions. Again, for expositional purposes, it is helpful to consider a differenced version of Eq. (3), comparing transactions for the same house at time t versus time t':

$$ln\left(p_{ijnt}\right) - ln\left(p_{i'jnt'}\right) = (raceeth\_s \times raceeth\_b_{it} - raceeth\_s \times raceeth\_b_{i't'})\lambda + \left(X_{ijt} - X_{i'jt'}\right)\beta + (\theta_{nt} - \theta_{nt'}) + \omega_{ijnt}. \quad (4)$$

Because *seller's* race is unobserved for the first transaction involving each house in our repeat-sales sample (as noted above), we consider three alternative approaches to estimating Eq. (4). All three approaches yield quite similar point estimates for the key parameters of interest, λ.

Our first approach involves restricting the sample to observations from homes that sell at least three times, dropping the first transaction for each of these houses (given that the seller's race is unknown) from the analysis. These results are reported as 'Approach 1' in Table 6. As one might expect, this approach reduces the number of observations in the sample substantially, entirely eliminating houses that sell only twice and cutting in half the number of observations in Eq. (4) for houses that sell three times. As a result, the neighborhood-by-time effects, θ, are not very precisely estimated, leading to less precise estimates of our main coefficients, λ.

In order to estimate the neighborhood-by-time effects more precisely, our two alternative approaches make use of the full repeat-sales sample to identify the pattern of price appreciation at the neighborhood level, while simultaneously making sure that observations that are missing the seller's race are not used to pin down the estimates of λ. On that basis, our second approach creates a new category for seller's race that equals one if it is unobserved. We then simply include interactions of this unobserved seller race category with the appropriate buyer race for the first transaction observed for each house. We treat the coefficients on interactions that involve sellers of unknown race and buyers of each race as nuisance parameters. The advantage of this approach over the first is that it allows all of the consecutive pairs of transactions to help identify the neighborhood-by-time fixed effects; these results are shown as 'Approach 2' in Table 6.

For our third approach, we estimate Eq. (4) for all consecutive pairs of transactions for each house j in the repeat-sales sample, but for the first pair of transactions for house j, we simply set the buyer-seller race difference (i.e., the first term on the right hand side of Eq. (4)) to zero. This is our preferred approach because houses that sell only twice, for example, are helpful in obtaining a much more precise estimate of price appreciation patterns by neighborhood, but do not contribute directly to the identification of λ. The results from this preferred approach are shown as 'Approach 3' in Table 6.

Each reported coefficient in Table 6 represents the price differential associated with each buyer-seller race/ethnicity combination relative to a white-to-white transaction for the All-City sample. The buyer's race/ethnicity is indicated in the row heading and the seller's race/ethnicity is indicated in the column heading. Thus, comparing results across rows shows how the sales price varies with buyer's race/ethnicity for sellers of the race/ethnicity given in the column heading. Comparing the row-average of the price differentials for buyers of different races is, in essence, the focus of Tables 2 and 3. In contrast, when examining the results presented in Table 6, we focus on comparing *across* columns. Doing so shows how the sales price for each type of buyer varies with the seller's race/ethnicity.

The results in the table reveal a clear, consistent pattern across the three approaches, implying that black and Hispanic buyers pay a premium that is at least as large when purchasing from black and Hispanic sellers (relative to white sellers). These results were foreshadowed to some extent by those presented in Table 5, which showed that the Hispanic price premia in particular were no larger in predominantly white neighborhoods than in those with a smaller share of non-Hispanic whites.

Examining the results in more detail for our preferred specification – Approach 3 – a black buyer, on average, pays a premium of

Table 6
Heterogeneity in racial/ethnic price differentials based on seller's race.

	White seller	Black seller	Hispanic seller	Obs.
Approach 1				**557,815**
Black buyer	0.017	0.033	0.026	
	[0.005]	[0.01]	[0.009]	
Hispanic buyer	0.018	0.022	0.030	
	[0.004]	[0.008]	[0.009]	
White buyer	–	0.015	−0.001	
		[0.010]	[0.005]	
Approach 2				**2,121,518**
Black buyer	0.015	0.033	0.021	
	[0.006]	[0.018]	[0.011]	
Hispanic buyer	0.020	0.032	0.033	
	[0.005]	[0.016]	[0.005]	
White buyer	–	−0.002	0.006	
		[0.010]	[0.005]	
Approach 3				**2,121,518**
Black buyer	0.019	0.038	0.025	
	[0.005]	[0.014]	[0.011]	
Hispanic buyer	0.018	0.027	0.033	
	[0.003]	[0.008]	[0.005]	
White buyer	–	0.008	−0.001	
		[0.010]	[0.005]	

Notes: Cluster-robust standard errors in brackets. Estimates are derived from a regression of log transaction prices on race indicators and other variables, including buyer income and down-payment percentage. Each regression includes a set of house fixed effects and a set of neighborhood-by-time fixed effects. All comparisons are relative to a white-to-white transaction. Estimates are reported for three alternative specifications, described as 'Approach 1,' 'Approach 2,' and 'Approach 3' in the table.

[26] We use the term "seller" to refer to the actual owner or the owner's selling agent. Due to data limitations, we are not able to distinguish when the seller employs the services of an agent.

[27] This test bears some resemblance to the test for racial profiling proposed by Anwar and Fang (2006) and Close and Mason (2007).

around 1.9% when buying from a white seller, compared to 3.8% and 2.5% when buying from black and Hispanic sellers, respectively. Similarly, Hispanic buyers pay a premium of 1.8% when purchasing from white sellers (relative to white buyers), but slightly more when buying from black and Hispanic sellers. The results for white buyers also paint a similar picture, revealing that purchase price is not a function of the seller's race and ethnicity. None of the differences across seller race and ethnicity are significantly different from zero for buyers of each race and ethnicity in each of the three specifications.

Taken as a whole, the results presented in Table 6 provide strong evidence that the price differentials for black and Hispanic buyers reported in Tables 2 and 3 are at least as large for sellers of the same race or ethnicity. Put slightly differently, they also imply that, conditional on buyer attributes, existing homeowners of each race and ethnicity do equally well (on average) when they sell their homes.

6. Mechanisms and interpretation

The magnitude and robust nature of the estimated racial/ethnic premia constitute strong evidence that black and Hispanic households pay more than their white counterparts when purchasing comparable homes. The existence of these price differentials naturally raises the question of whether they result from a form of discrimination on the part of sellers or arise due some other aspect of the home buying process that differs by buyer race/ethnicity. We take up this question in the current section by discussing the kinds of mechanism that might be consistent with the complete set of results reported in Sections 4 and 5.

The notion that discriminatory behavior is ongoing in housing markets is uncontroversial. Audit studies provide compelling evidence to this effect, offering a particularly powerful strategy for detecting signs of discriminatory behavior in housing. Generally speaking, such studies aim to test for discrimination by sending individuals of different races or ethnicities, matched as well as possible on other characteristics, to inquire about housing units either for sale or rent. Using this approach, Yinger (1986) and Ondrich et al. (2003) find direct evidence of the discriminatory treatment of minority buyers and renters along a number of dimensions. In particular, they uncover evidence of statistical discrimination likely stemming from uncertainty about black potential buyers' ability to put forward successful bids. They find, for example, that agents do not increase their effort in response to higher sales prices and that there is substantial steering of blacks to homes with fewer features than they request. These findings suggest that agents may be skeptical of the ability of blacks to purchase more expensive homes. If such beliefs are pervasive, one can imagine a situation where, conditional on choosing a particular house, blacks may need to submit higher bids than observationally equivalent whites in order to be taken seriously.

While audit studies provide a powerful way to identify the pervasiveness of exclusion and steering in the market, they yield little evidence about transactions that are actually consummated.[28,29] This matters both because audit studies may miss aspects of discrimination that occur through price negotiation/determination and because it is difficult to gauge the ultimate impact of the observed exclusion or steering on housing outcomes (Goldberg, 1996; Yinger, 1998).[30] Conversely, a limitation of studies such as ours that focus on the pricing of observed transactions is that they miss exactly the kinds of effects that steering and exclusion have on the choice of homes captured by audit studies. It is worth noting that sellers may observe potential buyer race/ethnicity imperfectly. While sellers certainly have numerous channels through which they can learn about the race of potential buyers, including meeting them in person, making inferences based on their name or current address, or through communication with the sellers' agents, there is still the possibility that sellers are unaware of the racial/ethnic background of potential buyers. Notwithstanding these concerns, in combination, the two approaches for detecting signs of discriminatory behavior in consumer markets complement each other neatly.

One potential explanation for the estimated price differentials is that they reflect differences in financial standing and access to credit when comparing black and Hispanic homebuyers with their white counterparts. It is indeed possible that sellers may have access to more information than simply the offer made by the buyer. In all of the areas under study, a typical financing-contingent offer includes some information regarding intended down payment. In addition, depending on the degree of competition for the home, there are likely to be cases where buyers reveal more information about their financial standing than would typically be included in a mortgage application, including job and salary information as well as the size of any 401k and other brokerage accounts, to assure sellers that they have a high probability of obtaining the requisite financing.[31] That possibility duly noted, the findings in Table 3 cast doubt on the importance of such financial factors. As already discussed, when we condition on key proxies for financial standing – income, down payment, and lender fixed effects – the differentials either increase or are unaffected.

Just as financial factors are unlikely to provide the principal explanation for these differentials, there is little evidence, in line with the conclusions of Cutler et al. (1999) and Clapp and Ross (2004), supporting the view that racial animus or prejudice exercised exclusively and directly by white homeowners serves as a primary explanation.[32] The results in Table 6 suggest that Hispanic and black sellers selling to non-white buyers obtain premia as large or larger than white sellers obtain. Such evidence does not, of course, rule out the possibility that all sellers, regardless of race of ethnicity, discriminate against black and Hispanic buyers. Moreover, while the arms-length nature of agent/broker-facilitated transactions may limit the information the homeowner has about a potential buyer's race and ethnicity, it also raises the possibility of discrimination on the part of the seller's real estate agent, which might occur regardless of the race or ethnicity of the seller they represent.[33]

[28] A recent paper by Hanson et al. (2016) does use a matched–pair email correspondence experiment combined with HMDA data to present evidence on mortgage pricing discrimination by race.

[29] Another limitation of audit studies is that they require all participants to behave identically, despite the fact that participants are likely aware of the goals of the study. To address this concern, Hanson and Hawley (2011) use email communication involved with online apartment listings, signaling race through the use of names that are predominantly used by either white or black parents. Another interesting recent paper by Ewens et al. (2014) uses a similar approach to distinguish between prejudice and statistical discrimination as explanations for call-back behavior. Their results point to statistical discrimination rather than prejudice being the most likely explanation for landlord behavior – broadly consistent with our findings.

[30] Ross and Turner (2005) report evidence that steering of blacks and Hispanics, in contrast to the declines in direct discrimination that have been observed, has increased in recent years

[31] Information obtained in private conversation with a real estate agent.

[32] Neither Cutler et al. (1999) nor Clapp and Ross (2004) find evidence of black premia in modern housing markets, concluding there was little evidence of centralized, systematic exclusion of blacks from access to housing markets on the basis of race.

[33] It is worth noting that the one-to-one nature of housing transactions means that racial discrimination based on prejudice can certainly survive in equilibrium, with prejudiced sellers simply foregoing some of their potential profits by refusing to sell to black or Hispanic buyers. The one-time nature of home sales does distinguish interactions in the owner- versus renter-occupied part of the housing market, where landlords and tenants enter into a longer-term relationship.

Another possible explanation for the robust premia we find is that non-white buyers pay higher prices because of differences in bargaining power or search costs. On this theme, Harding et al. (2003) develop a method that, under certain assumptions, allows them to infer the relative bargaining power of buyers and sellers involved in housing transactions using differences in their observable characteristics. Using data from the American Housing Survey, they find no evidence that blacks have lower bargaining power in the housing market, which suggests that scope for price discrimination against blacks may be limited.[34] Moreover, their estimates indicate that income and education – proxies for unobserved wealth – are actually negatively related to bargaining power in the housing market. This surprising result can be rationalized based on the notion that wealthier people may not bargain as aggressively because of diminishing marginal utility of wealth. If true, this finding would suggest that blacks and Hispanics, who tend have less income and wealth, should have the incentive to bargain more aggressively. Yet against that, blacks and Hispanics also tend to be first-time buyers, which they argue is associated with lower bargaining power. As an alternative explanation, higher income or wealthier buyers may face a higher opportunity cost of time and thus are willing to accept somewhat higher prices.

Search costs for non-white buyers might be higher for several reasons. If, for example, a substantial portion of sellers were historically unwilling to sell to minority buyers (see e.g., Courant, 1978), such buyers might then incorporate these expectations in their search behavior by performing more intensive searches in areas amenable to their race and perhaps avoiding homes for sale in areas that were not.[35] Alternatively, ongoing steering on the part of real estate agents may contribute to increasing search costs facing blacks and Hispanics, especially those searching in predominantly white neighborhoods (Turner et al., 2013; Ross and Turner, 2005). In both cases, the effective cost of continuing to search for a suitable home would be higher for non-white buyers, and this might lead them to settle for a higher transaction price in order to secure a property of their liking. The results showing that non-white buyers appear pay a premium with sellers of all races are consistent with this potential explanation.

7. Summary and conclusion

This paper has presented new estimates of the extent of any price differentials in the housing market on the basis of buyer race/ethnicity. Combining a repeat-sales approach with a rich new panel dataset that enabled us to control for unobserved housing and neighborhood quality more fully than in prior work, our analysis has documented statistically significant premia in the prices paid by black and Hispanic buyers in each of the four major metropolitan areas we studied – a result that is both new in the context of the traditional approach to estimating racial price premia and complementary to the best-known findings in the recent literature. Conditioning on flexible functions of income as well as lender fixed effects, the results indicated that the estimated racial/ethnic premia remained positive and even increased in size, implying that the correlation of race with these financial factors serves to lower the premia in our baseline specification.

Having established that there are robust differences in the prices paid by buyers of different races/ethnicities, we then considered whether systematic cross-racial bias by sellers could explain these differences. Here, we found that the premia paid by black and Hispanic buyers are higher when buying from black and Hispanic sellers than from whites. Indeed, pooling across all the cities in our sample, we found the premia to be highest for blacks when purchasing homes from blacks, and for Hispanics when purchasing from Hispanics. While these results provide evidence against racial bias on the part of sellers as the primary explanation for the estimated racial price premia, we cannot rule out that prejudice may lead to the exclusion of minority buyers from purchasing certain properties in the first place – the steering channel assessed in various audit studies.

No matter what the ultimate reason for the price premia, our results imply that systematic, robust racial/ethnic differences in the price paid to buy a home – on the order of 2% on average in multiple major US markets – persist to the present day, long after many of the most overt forms of institutional discrimination have been eliminated. Considering the average purchase price paid by a black homebuyer in the overall sample is approximately $180,000, this translates into an average premium of nearly $4,000 per transaction – a substantial amount relative to average household incomes of blacks in these major cities.

These price differentials are likely to have important implications for a range of other social and economic outcomes, including the evolution of racial differences in wealth, home ownership rates, and location decisions. Faced with what amounts to a substantial transaction tax with each home purchase, one would naturally expect home ownership rates to be lower and the benefits of ownership for wealth accumulation to be systematically diminished for minority households. Moreover, to the extent that these differentials represent price discrimination, the added cost may alter black and Hispanic household location decisions (i.e., leading to the choice of neighborhoods with more rental properties or lower prices).[36] Hence, existing residential segregation may be reinforced, which has further important consequences for educational and labor market outcomes in the longer term.

Data appendix

This data appendix is subdivided into three sections. The first describes the main data sources we use to build our new dataset; the second lays out the merge procedure we follow in some detail; and the third explains the sample restrictions we impose to arrive at the final dataset, along with the construction of several key variables. In the process, we provide a numerical description of the merge's success, as well as a comparison between the 'overall' merged dataset and the final repeat-sales sample that is used in our empirical analysis.

1. Data Sources

The dataset we have assembled combines information from three sources: proprietary housing transactions data collected by *DataQuick*; publicly-available loan application registry information contained in the Home Mortgage Disclosure Act (HMDA) data files; and neighborhood demographic information from the decennial Census.

The *DataQuick* data provide very detailed information about individual housing transactions: the transaction price, date of sale, loan amount, lender name, and location of the home for each transaction, along with assessor information. In addition, each property is characterized by a unique identifier that makes it

Interestingly, Early et al. (2016) estimate racial price differentials in the rental market of a similar magnitude to the ones we report here.

[34] In the application in their paper, the authors concede that their ability to detect racial price discrimination is limited by the relatively small number of mixed-race transactions in their American Housing Survey data.

[35] Consistent with this, we find evidence that blacks and Hispanics are disproportionately more likely to purchase homes from other blacks and Hispanics within a given neighborhood (in results not reported in the paper).

[36] See Yinger (1995) and Cutler and Glaeser (1997) for discussions of the consequences of segregation for educational and labor market outcomes. It is particularly noteworthy that the estimated premia are largest for black households in Chicago and Maryland – metropolitan areas that remain largely segregated along racial lines to this day.

possible to track the longitudinal transactions history of each home.

The HMDA data provide demographic information for each mortgage applicant, as required under the HMDA legislation in order to monitor potential redlining and discriminatory lending behavior. From these data, we are able to retrieve the race/ethnicity and income of the buyer, along with the transactions date and the census tract of the home in question for linkage purposes. (Further information is available at http://www.ffiec.gov.)

Tract-level neighborhood racial compositions are drawn from the 1990 and 2000 Censuses.

2. Merge Description

As noted in the main text, the *DataQuick* housing transaction files do not contain any information about buyers beyond their names. In order to make use of household demographic characteristics and self-reported income levels, we match housing transactions and assessor information available in the *DataQuick* files with detailed loan application data in the Home Mortgage Disclosure Act (HMDA) files. Before setting out the merge procedure, we describe how the main data components are prepared.

Data Preparation

The *DataQuick* data consist of transactions and assessor information for each transaction during the observation period.

We perform pre-merge cleaning of the raw *DataQuick* data: identifying transactions with 0 and missing loan amounts for elimination; condensing three-way sales to ensure a unique transaction; dropping clear data entry errors; dropping properties purchased by business entities; and dropping transactions with a nominal price of 0. In addition, we create two rounded loan amount variables: (1) loan amounts rounded to the nearest ten-dollars amount, and (2) loan amounts rounded to the nearest five-dollar increment. Since there are sometimes small discrepancies between the recorded loan amounts in HMDA and *DataQuick* for homes otherwise matched well using our other keys, these separate variables are used to further refine the match.

The HMDA data files are constructed using county-level information from each of our chosen MSAs. We create loan application registry files separately by year for each MSA. Each loan application has a number of key pieces of information: applicant (and co-applicant, if any) name(s); the gender of each applicant; loan amount; loan year; and the census tract where the home is located.

An important issue for our study concerns maintaining consistent coding of race and ethnicity variables for the applicant and co-applicant across the entire sample period. In 2004, in line with changes instituted in the 2000 US Census, HMDA changed the manner in which the race and ethnicity of borrowers were recorded, in two ways. First, instead of treating Hispanic as a 'race,' HMDA moved to classifying race and Hispanic ethnicity separately, therefore allowing for people of any race to claim Hispanic ancestry. In addition, the Asian racial category was split into two categories: (1) Asian and (2) Native Hawaiian and other Pacific Islander. To ensure uniformity of the race/ethnicity variables across the entire sample period, we recode the post-2004 race/ethnicity variables to be consistent with their coding prior to 2004. Thus, any applicant or co-applicant reporting Hispanic ancestry is assigned to the Hispanic category. Similarly, any applicant or co-applicant reporting Asian, Native Hawaiian, or Pacific Islander ancestry and no Hispanic ancestry is placed in the Asian category.

Having implemented consistent coding of race/ethnicity, we drop all loan applications that were not originated. We then append these files for use in the merge described below.

Merge Algorithm

The merge process begins by joining the HMDA and *DataQuick* files, based on the full 11-digit census tract number, county, and loan amount, separately by year. Then, a crosswalk is merged into this joined dataset to account for small lenders who use larger lenders when reporting to HMDA.

1. An initial 'join' based on census tract, year, and loan amount yields all reasonable matches.
2. We then proceed with the primary match algorithm, which ranks the resulting matches based on flagcodes, and retains the lowest flagcode (representing the *best* match quality) for each transaction. After initial assignment of these flagcodes, the algorithm then carries out further refinements.
3. All separate year files are appended. The program then ensures no unique mortgage is originated in two years. Next, if multiple mortgages are matched to the same house, the program first identifies the best match based on the criteria that we used to classify match quality, then retains this match and eliminates the rest. If the duplicate mortgages are of the same quality, we then attempt to determine if they are indeed duplicates or rather the consequence of data problems. To that end, we randomly choose one home and compare the other on the basis of applicant demographics discussed below. If any differences in these demographics are found, then the extra match is nullified. If not, then we randomly choose which match to use since it is a true duplicate report. Observations surviving that procedure are considered to be unmatched.
4. In a final step, we assign a match quality designation to each home. We classify these matches into the following categories: 'Perfect matches' refer to unique initial matches on the basis of the four matching keys: census tract, year, loan amount, and lender name. 'High-quality matches' include perfect matches (so defined) plus matches that differ very slightly, either on the basis of small discrepancies in loan amounts or 6-digit census tract identifiers. 'Medium-quality matches' are duplicate matches in which randomization was used to select the house we consider matched. And 'Low-quality matches' consist of all other matches in the dataset, based on our best guess given the closeness of our matching keys. Unmatched homes consist of cash-only transactions, duplicate houses in which the demographic information in the HMDA files does not match, and houses with errors in the matching keys or missing matching keys.

The counties used to construct our sample are as follows: 'Chicago' consists of homes sold in Cook, DuPage, Lake, and Will Counties; 'Los Angeles' refers to LA and Orange Counties; 'Maryland' comprises Anne Arundel, Baltimore, Calvert, and Howard Counties along with Baltimore City, and in the DC area, Prince George's, Montgomery, and Fredrick Counties; and 'San Francisco' consists of Alameda, Contra Costa, Marin, San Francisco, San Mateo, and Santa Clara Counties.

In Data Appendix Table 1, we present summary statistics that describe the success of the match among these counties. The Overall panel of the table shows that an average of 67% of all sales transactions across the four cities were matched (the range being 64 to 71%). In the bottom panel, labeled Match Quality, we describe the quality of the matches. The majority of these matches are so-called perfect matches – those unique matches achieved solely on the basis of our four matching keys. Because there is less certainty about the quality of the matches beyond these perfect matches, we provide a robustness test (see Table 4) where we estimate our baseline regression model using only these perfect matches.

To explore potential selection biases induced by the match procedure, we compare observable features of unmatched and matched houses in Data Appendix Table 2. Here, we focus on Los Angeles and San Francisco, because many house-specific characteristics are not available for Chicago and Maryland in our data.

Data Appendix Table 1
Performance of the merge procedure.

	Chicago	San Francisco	Maryland	Los Angeles	All-City
Total observations	1,180,387	1,442,806	1,019,135	2,527,439	6,169,767
Overall					
Matched	0.71	0.66	0.68	0.64	0.67
Cash only	0.12	0.11	0.14	0.14	0.13
Match Quality					
Perfect matches	0.70	0.72	0.71	0.64	0.67
High-quality matches	0.81	0.86	0.85	0.84	0.84
Medium-quality matches	0.02	0.02	0.02	0.02	0.02
Low-quality matches	0.17	0.12	0.14	0.15	0.14

Notes: The table summarizes, by MSA and aggregated overall, the HMDA-*DataQuick* merge for those counties used in the study, as discussed in the text. Total observations consist of all possible homes in the *DataQuick* file for those counties. The Matched row in the Overall panel refers to the fraction of the total observations that were successfully matched using the algorithm described in the Data appendix; the Cash only row consists of the fraction of total observations that were unmatched as a consequence of not appearing in HMDA because no mortgage was originated. The Match Quality panel summarizes the relative quality of the matches among those houses that were successfully matched. The matches are defined as follows: Perfect Matches refer to houses that can be merged uniquely on the basis of the four matching variables (as described in the Data appendix); High-Quality Matches refer to Perfect Matches plus almost-unique matches differing only slightly on the basis of loan amount or 6-digit census tract; Medium-Quality Matches are duplicate matches in which randomization was used to select those houses considered to be matched; and Low-Quality Matches consist of all other matches. Unmatched homes constitute the remainder, consisting of cash transactions, duplicate houses in which the demographic information in the HMDA files do not match, and houses with errors in the matching keys.

Data Appendix Table 2
Comparison of matched and unmatched homes – LA and SF.

	Los Angeles		San Francisco	
	Matched	Unmatched	Matched	Unmatched
Housing characteristic				
Year built	1964	1966	1958	1956
Square footage	1521.3	1515.4	1666.3	1718.8
Lot size	8612.4	9120.7	17,182.2	16,764.0
Total rooms	4.8	4.2	6.4	6.4
Bedrooms	2.8	2.7	3.0	3.0
Bathrooms	2.0	2.0	2.1	2.1
Observations	1,628,525	898,812	954,115	488,616

Notes: This table provides a comparison of the observable characteristics of matched and unmatched homes that appear in the Los Angeles and San Francisco datasets for those homes that had non-missing characteristics. We do not include the comparisons for the Chicago and Maryland MSAs as many of these characteristics are not available for the majority of homes there.

Specifically, we compare houses in the matched versus unmatched categories on the basis of rooms, bathrooms, and square footage. It is clear that no large discrepancies emerge when making these comparisons.

3. Sample Construction

In this section, we describe the construction of key variables used in the analysis and sample restrictions applied to these matched data in order to arrive at the repeat-sales sample used in the main empirical analysis.

The final dataset for each city used in the empirical analysis is constructed from the matched *DataQuick*-HMDA data. Our sample consists of data for San Francisco and Los Angeles metropolitan areas from 1988–2007, and for the Chicago metropolitan area and Baltimore-Washington DC corridor in Maryland from 1997–2007. Within these sample years, we focus on transactions involving single-family homes, condominiums and townhouses purchased using a mortgage. (We discuss other sample restrictions below.)

Variable Construction

We provide a brief description of the way several key variables are constructed:

a. Race Variables: The HMDA files provide the race and ethnicity of the 'Applicant' and the 'Co-Applicant.' The race of the household is assigned using these variables. We form four groups: 'Black,' 'Hispanic,' 'White,' and 'Asian and Other.' As Hispanic households can technically be of any race, every household listing Hispanic ethnicity is assigned to that group. If both applicant and – where relevant – co-applicant indicate white race and non-Hispanic ethnicity, then the household is assigned to the white category. Any households that include members of non-white minorities and/or Hispanic ethnicity are assigned to non-white racial or ethnic groups. If at least one applicant is black and non-Hispanic, then they are assigned to the black category. If either applicant or co-applicant reports that they are Hispanic, then we assign them to the Hispanic ethnic category. All non-black or non-Hispanic households who are not considered white are assigned to the Asian and Other category. For interracial and inter-ethnic couples consisting of an Asian and a white person, we treat such couples as Asian households for the purposes of the analysis.[37]

b. Race-to-Race Matches: Race-to-race matches for the buyer-seller analysis are constructed on the following lines: When a house is involved in a repeat sale, we assign the race of the buyer of that home to the seller when the home sells next. On that basis, for every house in the repeat-sale sample that has complete information, we are able to determine the race/ethnicity of the buyer and seller, save for the first transaction. In that case, we do not know the race of the seller, so we cannot classify its transaction type fully. We implemented several approaches (Approaches 1–3) to deal with this problem, discussed in more detail in the paper. Under the first approach, we exclude the homes with no race information from the analysis. Our second approach creates an 'unobserved' category for homes that have no associated initial information. A third strategy (Approach 3) assigns the race/ethnicity of the missing initial seller to be equal to the race/ethnicity of the first observed buyer, thereby setting the racial/ethnic difference on this first

[37] Treating these interracial couples as white has little impact to the qualitative nature of the results for blacks and Hispanics as most interracial couples in the data are white and Asian.

Table A1
Comparison of repeat-sales sample with overall transactions sample.

	Repeat-sales sample			
	Chicago	San Francisco	Maryland	Los Angeles
Buyer race (proportions)				
White	0.74	0.61	0.62	0.64
Black	0.09	0.04	0.25	0.06
Hispanic	0.08	0.06	0.05	0.11
Asian and other	0.09	0.29	0.08	0.19
Transaction statistics ($)				
Mean transaction price	220,737	380,258	227,944	306,974
Median transaction price	182,500	312,500	190,000	239,000
Mean income	86,558	120,489	88,188	106,434
Median income	67,000	99,000	71,000	81,000
Median down-payment	27,500	78,000	22,000	48,500
Observations	382, 389	535, 286	278, 221	925, 622
	Overall Sample			
	Chicago	San Francisco	Maryland	Los Angeles
Buyer race (proportions)				
White	0.74	0.61	0.60	0.64
Black	0.10	0.04	0.26	0.06
Hispanic	0.08	0.07	0.07	0.10
Asian and other	0.08	0.28	0.07	0.20
Transaction statistics ($)				
Mean transaction price	225,761	394,285	235,986	310,514
Median transaction price	188,000	324,500	186,000	239,000
Mean income	86, 697	122,432	92,324	106,508
Median income	68,000	101,000	72,000	80,000
Median down-payment	36,000	75,499	33,000	48,000
Observations	1,065,154	1,271,820	793,900	1,998,526

Notes: This table compares characteristics of the repeat-sales sample (i.e., transactions involving houses that sell a minimum of twice) with the overall sample of transactions for the four metropolitan areas. Statistics for the repeat-sales sample are reported in the upper panel and for the overall sample, in the lower panel. The transaction price is the actual recorded closing price of the home.'Income' refers to reported income on the mortgage application. (Sources: *DataQuick* and HMDA data.)

Table A2
Comparing characteristics of homes involved in two versus three or more repeat-sales transactions.

	Chicago		San Francisco		Maryland		Los Angeles	
	2 Only	3 or More	2 Only	3 or More	2 Only	3 or More	2 Only	3 or More
Mean price	222,027	214,981	388,152	365,154	231,013	214,356	310,845	299,167
Median price	184,000	176,500	320,000	300,000	193,852	180,000	240,000	234,500
Square footage	1945	1743	1445	1335	1364	1320	1407	1223
Baths	2	2	2	2	2	2	2	2
Bedrooms	–	–	3	3	-	–	3	3
Observations	228,039	51,115	341,251	178,350	166,787	37,666	527,960	261,791

Notes: The table compares observable characteristics of homes that are sold only twice versus three or more times in the repeat-sales sample for homes that had non-missing characteristics. The number of bedrooms is not available in the Chicago and Maryland datasets. In addition, the number of baths is only available for less 30% of the sample in both the Chicago and Maryland datasets.

transaction to be zero. An additional approach not reported here assigns "white" to all homes with unobserved race/ ethnicity.

c. Down payment Variables: Down payments are calculated in the following manner: We sum up the total mortgage amounts reported in the loan application registry. Then we take the difference between the total amount of mortgage taken out on the house and the transaction price. We classify households as making a 'standard' down payment if their down payment constitutes at least 20% of the purchase price.

d. Lender Fixed Effects: These are constructed using the lender code variable provided in the data. To conserve parameters, lenders in the top 50th percentile of loans originated are considered individually, while the rest are amalgamated into a single category, as many of these lenders originated fewer than 20 loans in the dataset.

e. Census Variables: Neighborhood characteristics are drawn from the 1990 and 2000 Censuses. We use 1990-to-2000 crosswalks to connect tracts across time in a stable way.

Sample Restrictions

In order to arrive at the final dataset used in our empirical analysis, we impose a number of sample restrictions on the matched data for each city.

These include:

1. Outlier Prices: We trim the data based on outlier prices, dropping transactions above the 99th percentile and below the 1st percentile.

Table A3
Heterogeneity in racial/ethnic price differentials based on buyer income quartile.

	All-city estimates	
	Black	Hispanic
Main	0.038 [0.007]	0.022 [0.005]
x Income Quartile 2	−0.01 [0.005]	0.005 [0.002]
x Income Quartile 3	−0.02 [0.005]	−0.005 [0.003]
x Income Quartile 4	-0.024 [0.006]	-0.016 [0.004]
Other controls	Yes	Yes
House fixed effects	Yes	Yes
Tract x time fixed effects	Yes	
Observations	2,121,518	2,121,518

Notes: Cluster-robust standard errors in brackets. Estimates are derived from a regression of log transaction price on race indicators, a set of house fixed effects and a set of neighborhood-by-time fixed effects. 'Other controls' include income quartile indicators and a down payment percentage polynomial. Income quartile indicators are assigned using the city-specific buyer income distribution.

2. Arms-length Transactions: These are transactions conducted by two independent parties. They exclude, for example, transactions between parents and children, the concern being that if the parties are not independent, the price may not reflect the fair market value. All transactions not listed as 'arms-length' using the relevant flag in the *DataQuick* data are excluded.
3. Land Sales: To ensure all transactions actually involve homes, we eliminate all sales that consist of unimproved land with no structure.
4. Construction Year: We eliminate any home whose construction year is recorded as being after the transaction date.
5. Annualized Appreciation: We drop houses that experienced annualized appreciation in excess of 100%.[38]
6. 'Major Improvements': We drop houses that have undergone major improvements as indicated by flags for 'major improvements' in the *DataQuick* data. To deal with issues related to possible unobserved renovation, we remove any house whose loan amount is greater than the actual transaction price, as such additional funds are likely to be used for improvements to the house. Further, the cuts based on house price appreciation already mentioned help mitigate concerns about unobserved major improvements.
7. Income: Household income is reported by households on their mortgage application. We dropped observations on the basis of potential data entry errors such as negative or implausibly large income.

Final dataset

Having implemented the merge procedure, data construction, and sample restriction steps described above, we arrive at the final dataset used in our analysis.

A summary comparison of the repeat-sales sample with the overall transactions sample is presented in Appendix Table A1. It is apparent that the repeat-sales sample is highly representative of the overall sample, both in terms of the racial composition of buyers and transaction prices.

Among the set of homes that sold repeatedly, there is a potential concern that those selling only twice during the sample period might differ from those that sold more than twice. To shed light on this issue, Appendix Table A2 presents summary statistics comparing homes that sold twice with homes that sold three or more times. The table makes clear that there are only small discrepancies in the transaction prices and square footage across the two groups of homes, and they are very similar in terms of the number of beds and bathrooms, suggesting that any bias issues are not likely to be serious in practice.

Table A4
Heterogeneity in racial/ethnic price differentials - pre-2004 transactions.

	White seller	Black seller	Hispanic seller	Obs.
Approach 1				**108,322**
Black buyer	0.023 [0.012]	0.041 [0.028]	0.031 [0.021]	
Hispanic buyer	0.010 [0.007]	0.007 [0.021]	0.024 [0.012]	
White buyer	–	0.012 [0.010]	−0.005 [0.008]	
Approach 2				**1,508,721**
Black buyer	0.016 [0.010]	0.023 [0.008]	0.083 [0.038]	
Hispanic buyer	0.031 [0.007]	0.032 [0.008]	0.031 [0.011]	
White buyer	–	0.037 [0.027]	0.002 [0.005]	
Approach 3				
Black buyer	0.014 [0.013]	0.026 [0.016]	0.016 [0.010]	**1,508,721**
Hispanic buyer	0.011 [0.008]	0.016 [0.014]	0.018 [0.010]	
White buyer	–	−0.001 [0.014]	0.004 [0.006]	

Notes: Cluster-robust standard errors in brackets. Estimates are derived from a regression of log transaction prices on race indicators and other variables, including buyer income and down-payment percentage. Each regression includes a set of neighborhood-by-time fixed effects. All comparisons are relative to a white-to-white transaction. Estimates are reported for three alternative specifications, described as 'Approach 1,' 'Approach 2,' and 'Approach 3' in the text.

Table A5
Heterogeneity in racial/ethnic price differentials – no house FEs.

	White seller	Black seller	Hispanic seller	Obs.
Approach 1				**557,815**
Black buyer	−0.014 [0.006]	−0.041 [0.014]	−0.069 [0.009]	
Hispanic buyer	−0.046 [0.005]	−0.049 [0.008]	−0.057 [0.008]	
White buyer	–	−0.047 [0.006]	−0.089 [0.005]	
Approach 2				**2,121,518**
Black buyer	−0.021 [0.004]	−0.028 [0.006]	−0.066 [0.005]	
Hispanic buyer	−0.043 [0.003]	−0.040 [0.007]	−0.046 [0.005]	
White buyer	–	−0.036 [0.004]	−0.061 [0.005]	
Approach 3				**2,121,518**
Black buyer	−0.007 [0.010]	−0.021 [0.014]	−0.023 [0.014]	
Hispanic buyer	−0.018 [0.011]	−0.019 [0.008]	−0.018 [0.015]	
White buyer	–	−0.042 [0.008]	−0.069 [0.019]	

Notes: Cluster-robust standard errors in brackets. Estimates are derived from a regression of log transaction prices on race indicators and other variables, including buyer income and down-payment percentage. Each regression includes a set of neighborhood-by-time fixed effects. All comparisons are relative to a white-to-white transaction. Estimates are reported for three alternative specifications, described as 'Approach 1,' 'Approach 2,' and 'Approach 3' in the text.

[38] In an earlier version of the paper, we trimmed the data based on annualized appreciation in excess of 200%.

Table A6
Probability of purchasing from a white seller – matched sample.

			Obs.
Chicago			
Black	0.19	0.31	384,114
Hispanic	0.29	0.34	
White	0.92	0.89	
San Francisco			535,286
Black	0.20	0.23	
Hispanic	0.15	0.19	
White	0.89	0.86	
Maryland			279,304
Black	0.19	0.25	
Hispanic	0.15	0.27	
White	0.92	0.88	
Los Angeles			926,713
Black	0.37	0.23	
Hispanic	0.36	0.20	
White	0.90	0.85	
Tract FEs?	No	Yes	

Notes: Probability of purchasing a house from a white seller conditional on being a black, Hispanic or white buyer – estimates from a linear probability model.

References

Anwar, S., Fang, H., 2006. An alternative test of racial prejudice in motor vehicle searches: theory and evidence. Am. Econ. Rev. 96 (1), 127–151.

Arrow, K.J., 1998. What has economics to say about racial discrimination? J. Econ. Perspect. 12 (2), 91–100.

Becker, B.J., Wu, M., 2007. Synthesis of regression slopes in meta–analysis,. Stat. Sci. 33 (3), 414–429.

Becker, G., 1971. The Economics of Discrimination. University of Chicago Press, Chicago.

Chambers, D.N., 1992. The racial house price differential and racially transitional neighborhoods. J. Urban Econ. 32, 214–232.

Clapp, J.M., Ross, S.L., 2004. Schools and housing markets: an examination of school segregation and performance in connecticut. Econ. J. 114 (November), F435–F440.

Close, B., Mason, P., 2007. Searching for efficient enforcement: officer characteristics and racially-biased policing. Rev. Law Econ. 3 (2).

Courant, P., 1978. Racial prejudice in a search model of the urban housing market. J. Urban Econ. 5 (3), 329–345.

Cutler, D.M., Glaeser, E., 1997. Are Ghettos good or bad? Q. J. Econ. 112 (3), 827–872.

Cutler, D.M., Glaeser, E., Vigdor, J., 1999. The rise and decline of the American Ghetto. J. Polit. Econ. 107 (3), 455–506.

Early, D., Carillo, P., Olsen, E., 2016. Racial Rent Differences in the Housing Market. Southwestern University Mimeo.

Ewens, M., Tomlin, B., Wang, C., 2014. Statistical discrimination or prejudice? A large sample field experiment. Rev. Econ. Stat. 94 (1), 119–134.

Ferreira, F. and J. Gyourko (2015), “A new look at the U.S. foreclosure crisis: panel data evidence of prime and subprime borrowers from 1997 to 2012,” NBER Working Paper #21261.

Follain, J.R., Malpezzi, S., 1981. Another look at racial differences in housing prices. Urban Stud. 18, 195–203.

Goldberg, P.K., 1996. Dealer price discrimination in new car purchases: evidence from the consumer expenditure survey. J. Polit. Econ. 104 (3), 622–654.

Hanson, A., Hawley, Z., 2011. Do landlords discriminate in the rental housing market? Evidence from an internet field experiment in U.S. cities. J. Urban Econ. 70 (2), 99–114.

Hanson, A., Hawley, Z., Martin, H., Liu, B., 2016. Discrimination in mortgage lending: evidence from a correspondence experiment. J. Urban Econ. 92, 48–65.

Harding, J.P., Rosenthal, S., Sirmans, C.F., 2003. Estimating bargaining power in the market for existing homes. Rev. Econ. Stat. 85 (1), 178–188.

Hirsch, A., 1978. Making the Second Ghetto: Race and Housing in Chicago 1940–1970. University of Chicago Press, Chicago, IL.

Ihlanfeldt, K., Mayock, T., 2009. Price discrimination in the housing market. J. Urban Econ. 66:, 125–140.

Kain, J.F., Quigley, J.M., 1975. Housing Markets and Racial Discrimination: A Microeconomic Analysis. NBER, Cambridge, MA.

Kiel, K.A., Zabel, J., 1996. House price differentials in US cities: household and neighborhood racial effects. J. Hous. Econ. 5, 171–189.

King, A., Mieszkowski, P., 1973. Racial discrimination, segregation, and the price of housing. J. Polit. Econ. 81, 590–606.

Myers, C.K., 2004. Discrimination and neighborhood effects: understanding racial differentials in US housing prices. J. Urban Econ. 56, 279–306.

Ondrich, J., Ross, S., Yinger, J., 2003. Now you see it, now you don't: why do agents withhold available houses from black customers? Rev. Econ. Stat. 85 (4), 854–873.

Ross, S.L., Turner, M.A., 2005. Housing discrimination in metropolitan America: explaining changes between 1989 and 2000. Soc. Probl. 52((2), 152–180.

Turner, M.A., Santos, R., Levy, D.K., Wissoker, D., Aranda, C., Pitingolo, R., 2013. Housing Discrimination Against Racial and Ethnic Minorities 2012. Urban Institute.

Yinger, J., 1978. The black-white price differential in housing: some further evidence. Land Econ. 54, 187–206.

Yinger, J., 1986. Measuring discrimination with fair housing audits: caught in the act. Am. Econ. Rev. 76, 881–893.

Yinger, J., 1995. Closed Doors, Opportunities Lost: The Continuing Cost of Housing Discrimination. Russell Sage Foundation, New York.

Yinger, J., 1998. Evidence on discrimination in consumer markets. J. Econ. Perspect. 12 (2), 23–40.