Does Strong Credit Overcome Racial Discrimination in Rental Housing?

Ayden Young[a]* Peter Christensen[a,b] Jeremy West[a]

[Job Market Paper]

[a] *University of California, Santa Cruz*
[b] *National Bureau of Economic Research*

October 2026

Abstract

Racial discrimination constrains Black and Hispanic renters' access to housing. We ask whether stronger economic credentials can help renters to overcome these racial barriers. Linking experimental measures of neighborhood-level discrimination to individual credit records for California renters, we find evidence that strong credit partially offsets racial barriers: the credit score gap between Black and Hispanic renters and white renters moving into the same Census block is smaller where discriminatory barriers are larger, consistent with landlords being more selective when considering Black and Hispanic applicants. The offsetting effect is pronounced only for renters with Very Good or Exceptional credit scores, and is stronger in lower-credit neighborhoods, where a high score more favorably stands out. Collectively, our findings support that while positive economic signals are inadequate to fully counteract racial discrimination in rental housing, there are meaningful offsetting benefits for those with sufficiently strong financial circumstances.

JEL: D83, J15, R31, G51

Keywords: racial gaps, credit scores, economic mobility, random-audit studies

*Corresponding author: Ayden Young, ayoung14@ucsc.edu. We gratefully acknowledge financial support for this project from the University of California, Santa Cruz Institute for Social Transformation and the UCSC Department of Economics Dissertation Research Grant program. We thank the California Policy Lab for hosting, documenting, and facilitating access to the University of California Consumer Credit Panel, and thank Ignacio Sarmiento-Barbieri for discussion about the correspondence study data that we use in this paper. This work benefited greatly from helpful comments by seminar participants at UCSC. Any errors and conclusions are our own and do not represent the position of the California Policy Lab or any credit bureau.

Introduction

We know a great deal about how discrimination limits access, but much less about what enables some to overcome it. Correspondence studies, which send landlords otherwise identical inquiries that differ only in the sender's racially distinctive name, have established that Black and Hispanic applicants receive fewer responses from landlords than white applicants (Christensen et al., 2021), and these barriers constrain access to high-opportunity neighborhoods with lasting consequences (Chetty et al., 2016; Christensen and Timmins, 2023). Yet landlords screen on a broader set of renter qualifications, and strength on these may compensate for a racial penalty. Whether such compensation is possible depends on the form discrimination takes: some landlords may exclude applicants regardless of their qualifications, while others may hold them to a higher bar that strong enough credentials can clear.

Evidence from the labor market suggests the second form is a meaningful channel. Correspondence studies document employer-level racial discrimination in callbacks (Bertrand and Mullainathan, 2004; Kline et al., 2022), and stronger credentials can offset part of it: Black applicants from elite universities receive as many employer responses as white applicants from less selective universities (Gaddis, 2015). Whether a similar offset operates in rental housing, and how it varies with the intensity of local discrimination, remains largely unexamined as correspondence studies typically end at the landlord's reply, before we observe who secures housing or what credentials they hold. This paper asks whether Black and Hispanic renters can offset part of the discrimination they face with stronger economic credentials, measured by credit scores. Credit is a natural measure of these credentials in this context as it is a common signal landlords screen on, racial gaps in credit are large, well-documented, and largely rooted in structural inequality rather than current housing-market behavior (Bhutta et al., 2025), and unlike rental histories or references, it is recorded at scale. We treat credit as an observable proxy for the broader set of renter qualifications landlords consider.

We answer this question by linking two sources of data. The first is the correspondence experiment of Christensen et al. (2021), in which fictitious rental inquiries with randomized racially distinctive names were sent to property managers across the fifty largest U.S. metropolitan areas, generating experimentally derived racial response-rate gaps at the individual listing level. Because inquirer race was randomly assigned, the response-rate gap measures landlord behavior toward Black and Hispanic applicants rather than any other differences among applicants. We aggregate these response rates to the census block level for six California metropolitan areas, yielding roughly 936 experimental blocks. We then merge this with the University of California Consumer Credit Panel (UC-CCP), which tracks quar-

terly credit scores and residential location for California residents with a credit history from 2011 to 2026. Linking these data lets us compare the credit profiles of approximately 500,000 Black, Hispanic, and white renters who move into the same census blocks (hereafter, movers).

Because high-barrier neighborhoods differ from low-barrier ones in many ways, we use a difference-in-differences design where we compare Black and Hispanic movers to white movers within the same census block, and ask how that within-block credit score gap varies across blocks with different intensities of experimentally measured discrimination. Block fixed effects absorb stable neighborhood characteristics common to all renters. Our main finding is that the credit score gap between Black and Hispanic movers and white movers in the same block is smaller where the own-race barrier is larger. In blocks where landlords responded to no Hispanic inquiries, the Hispanic–white gap is about 7 points smaller than in blocks where they responded to all of them. The corresponding estimate for Black movers is about 3.5 points and also statistically significant; although it is about half as large, we cannot reject that the two estimates are equal. The pattern is concentrated in the upper tail of the credit distribution, among renters with Very Good or Exceptional credit, with little change in the middle. It is also stronger in neighborhoods with lower average credit, where a given score stands out more favorably from the local pool.

These patterns suggest that strong credit partially offsets the racial penalty in housing access. We also consider an alternative sorting-based explanation: Black and Hispanic renters who move into higher-barrier blocks might be positively selected relative to their white neighbors for reasons unrelated to landlord screening. Block fixed effects absorb any sorting that shifts credit equally for all groups within a block. Any remaining sorting would have to be race-specific, operate within the same block, survive controls for occupation, education, and household composition, and concentrate in the upper tail of the credit distribution. We cannot rule out such sorting, but given how demanding the joint condition is, we view the evidence as more consistent with landlord screening.

This paper makes three contributions. First, we provide the first evidence that strong economic credentials partially offset racial discrimination in the rental housing market. A large correspondence literature documents discrimination at the inquiry stage (Yinger, 1995; Hanson and Hawley, 2011; Ewens et al., 2014; Christensen et al., 2021; Christensen and Timmins, 2022), but these studies typically end at the landlord's reply. By linking experimental discrimination measures to the credit of renters who secure housing, we observe what happens after the landlords screen. The closest precedent, Berkovec et al. (1994), tested for a similar mechanism in mortgage lending by inferring borrower quality from default and found little support; we observe credit directly, and in a market where the screener is a landlord rather than a lender (Bhutta and Hizmo, 2021; Dobbie et al., 2021; Bartlett et al., 2022). Second,

we show that the offset requires a sufficiently strong signal: it is concentrated among renters with Very Good or Exceptional credit, with little change in the middle of the distribution. Credentials do not offset discrimination at the margin; they must be strong enough to clear a higher bar. Third, we show that the offset is stronger where the local credit pool is weaker, where a given score stands out more. The value of a credential in overcoming discrimination thus depends on how far it sets an applicant apart from competing applicants, a logic that extends to other screening markets, where the same signal may be more or less distinctive depending on the applicant pool. These findings suggest that restricting the use of credit in tenant screening could close one of the channels through which Black and Hispanic renters overcome discriminatory barriers, echoing evidence that removing screening signals can harm the groups such policies aim to protect (Agan and Starr, 2018; Doleac and Hansen, 2020). We discuss this implication in the conclusion.

Predictions

The correspondence experiment of Christensen et al. (2021) establishes that Black and Hispanic renters face substantially lower response rates from landlords than otherwise identical white inquiries. We treat this documented response-rate gap as evidence that Black and Hispanic renters face a higher effective screening threshold in the rental application process, regardless of whether that threshold reflects taste-based (Becker, 1957) or statistical discrimination (Phelps, 1972; Arrow, 1973), or landlords' expectations of incumbent white tenants' preferences, as in the literature on customer discrimination (Holzer and Ihlanfeldt, 1998; Yinger, 1995). Standard models of discrimination imply that when a higher threshold governs entry, only applicants with sufficiently strong credentials clear it (Becker, 1957; Phelps, 1972), and labor market evidence shows that stronger credentials can offset part of a racial penalty (Gaddis, 2015). Three predictions follow from this threshold structure.

P1: The credit score gap narrows where the barrier is larger. If Black and Hispanic renters face a higher effective screening threshold and strong credit can partially offset it, the within-block credit score gap between Black and Hispanic renters and white renters should be smaller in blocks where the own-race barrier is larger. This is the housing market analog of the credential offset studied in labor markets.

P2: The offset is concentrated in the upper tail of the credit distribution. A threshold mechanism implies that credit matters discontinuously. Renters below the effective threshold cannot clear it regardless of marginal improvements, while those above it can. The credential offset should therefore appear in the upper tiers of the credit distribution, among renters with Very Good or Exceptional credit, rather than distributing uniformly across the distribution.

Renters with Poor credit should be underrepresented in higher-barrier blocks relative to white renters, reflecting the access disadvantage that the threshold imposes on those furthest from clearing it.

P3: The offset is stronger in lower-credit neighborhoods. If landlords evaluate applicant credit relative to the local applicant pool rather than on an absolute scale, a given credit score represents a larger deviation above the local mean in neighborhoods where the average credit score is lower. The credential signal is therefore more powerful in lower-credit neighborhoods, and the offset should be stronger where the surrounding population has lower average credit.

Data

Experimental Discrimination Measures

We test these predictions using two sources of data. The first source is block-level discrimination measures constructed from the correspondence experiment of Christensen et al. (2021). The experiment sent automated inquiries from fictitious renters randomizing perceived inquirer race via names associated with Black, Hispanic, and white racial identities to rental property managers across the fifty largest U.S. metropolitan areas between February and July 2020. In each market, rental listings were sampled through stratified random sampling each week across the experimental window, a design intended to balance the listings sample against market-specific dynamics and seasonality. For each experimental listing k in census block j, and for each group g, the response indicator r^g_{jk} measures whether the property manager responded to the inquiry. We aggregate these response measures to the census block level and define:

$$b^g_j = 1 - \bar{r}^g_j \tag{1}$$

where $\bar{r}^g_j$ is the block-level average of r^g_{jk} for all listings k in census block j, and b^g_j is the own-race barrier for group g in block j, so that higher barrier values correspond to more intense levels of discrimination. The response measures are aggregated to the census block level, the finest geography at which we can link the credit panel. The California sample includes six metropolitan areas: Los Angeles, San Diego, and Riverside (Southern California) and San Francisco, San Jose, and Sacramento (Northern California). These six markets provide an average of 936 experimental census blocks across the three census vintages. The barrier measure varies substantially across blocks and exhibits systematic geographic structure: discrimination against Hispanic renters is most pronounced in Southern California markets such as San Diego and comparatively low in Northern California markets

such as Sacramento and San Jose. This cross-block variation is what our empirical strategy exploits.

Credit Panel Data

The second is the University of California's Consumer Credit Panel (UC-CCP), a large credit panel derived from anonymized records of a nationwide consumer reporting agency, the type of source commonly used in tenant screening. It covers California residents from the first quarter of 2011 to the first quarter of 2026. The panel data provide individual-level credit scores, financial characteristics, BISG race probabilities, and geocoded census-block codes for the population of California residents with a formal credit history. The panel tracks individuals at a quarterly frequency, allowing us to observe credit scores over time and whether individuals move residences between quarters.

The Bayesian Improved Surname Geocoding (BISG) probabilities combine surname and geographic location to estimate race probabilities (Elliott et al., 2009), and are provided across a range of racial groups in the UC-CCP. We assign race to an individual based on the highest provided BISG probability, and then restrict to individuals classified as white, Black, or Hispanic.

Since the credit panel period spans 2011 to 2026, individual residence is recorded under different census block vintages over time with the 2000 census vintage applying to the earliest part of the panel, 2010 vintage to the middle, and the 2020 vintage applies to later quarters. Census block boundaries are redrawn at each decennial census, so a given physical location may correspond to different block identifiers across vintages. Because the discrimination measure is constructed from geocoded experimental listings, we re-aggregate the listing-level response indicators to the block definition of the appropriate vintage, so that b_j^g is computed on vintage-consistent boundaries rather than mapped mechanically across them. We then assign each individual-quarter to the block definition corresponding to its vintage and link it to the discrimination measure computed on that same vintage.

Sample Construction

Our analysis sample consists of movers, defined as individuals observed in an experimental census block in a given quarter who were not observed in that block in the prior quarter. We impose minimal additional restrictions. We include living individuals regardless of age, and we retain primary record holders, associated borrowers, and other household members. Race is assigned as described above, and we restrict to individuals classified as white, Black, or Hispanic. An individual may contribute multiple mover observations if they move into

experimental blocks in more than one quarter, and we treat all such observations symmetrically: each mover observation corresponds to a distinct instance of clearing the landlord screening process, which is the relevant event for our mechanism. Nearly all movers in our sample, 96 percent, reside in census blocks where movers from all three racial groups are present, indicating that imbalanced racial coverage across blocks is unlikely to drive our results. Table 1 reports summary statistics by racial group. We exclude roughly 7,300 observations with recorded credit score values below 300, which fall outside the range of the score and appear to be non-score codes.

Empirical Strategy

Descriptive Motivation

We begin by documenting the raw relationship between credit quality and the racial composition of movers who move into experimental blocks. We regress an indicator for mover race on credit score, the block-level barrier measure, and their interaction. For individual i, of racial group g, in census block j in quarter t, we have:

$$\mathbf{1}[\text{Race}_{ijt}^{g}] = \beta_1 \text{Credit}_{it} + \beta_2 b_j^g + \beta_3 \big(\text{Credit}_{it} \times b_j^g\big) + \varepsilon_{ijt} \tag{2}$$

This reduced-form specification is descriptive rather than causal: it pools variation across blocks and does not isolate the within-block relationship that our main empirical strategy targets. A causal interpretation of β_3 would require that the own-race barrier be uncorrelated with unobserved factors that determine racial composition, conditional on credit and time effects. In this pooled specification, with no block level fixed-effects, that condition fails. A high-barrier block is not otherwise comparable to a low-barrier block as own-race barrier is systematically related to block desirability, price, and racial composition, any of which can independently influence who moves in. Thus, β_3 conflates the relationship of interest with cross-block differences that are correlated with the own-race barrier. For example, if high-credit Black and Hispanic renters systematically prefer different blocks than high-credit white renters, and those preferences happen to be correlated with the local discrimination environment, β_3 would capture that cross-block differential sorting rather than the relationship of interest. It nonetheless establishes the basic pattern motivating the analysis, reported in Table **??**, that higher-credit movers who move into higher-barrier blocks are disproportionately more likely to be Black and Hispanic. We turn next to our main specification that absorbs these block-level confounders directly.

Main Specification

The descriptive pattern motivates a specification that isolates the within-block relationship between the own-race barrier and renter credit scores. We use credit score as the outcome rather than the race indicators for two reasons. Credit score is the signal landlords observe and screen on during the application process, making it the natural margin on which elevated credentials would appear if discrimination raises the bar for Black and Hispanic applicants. It is also measured in continuous, interpretable units that allow us to quantify the magnitude of the credit score gap directly. Our primary specification thus estimates how the credit score gap between Black and Hispanic movers and white movers, within census blocks, varies with the intensity of local discrimination. For individual i in census block j at time t, we estimate:

$$\begin{aligned}\text{Credit}_{ijt} = \delta_B\big(\mathbf{1}[\text{Black}_i] \times b_j^B\big) + \delta_H\big(\mathbf{1}[\text{Hispanic}_i] \times b_j^H\big) \\ + \alpha_B(\mathbf{1}[\text{Black}_i]) + \alpha_H(\mathbf{1}[\text{Hispanic}_i]) + \mu_j + \gamma_t + \varepsilon_{ijt}\end{aligned} \tag{3}$$

where b_j^g is the own-race barrier for group g in block j, μ_j is a census block fixed effect, and γ_t is a quarter-year fixed effect. The level coefficients α_B and α_H capture the average credit gap between Black and Hispanic movers and white movers in the same block, evaluated at the zero barrier. Standard errors are clustered at the census block level.

The coefficients of interest are δ_B and δ_H, which capture how the credit gap between Black and Hispanic movers and white movers in the same block varies with own-race discrimination intensity. These δ_g magnitudes are informed by the baseline average credit score differences across racial groups, captured by α_g. A positive δ_g indicates that movers of racial group g are drawn from higher in the credit distribution in blocks where their group faces a larger barrier, consistent with the screening mechanism described in the conceptual framework outlined in the appendix. The block fixed effect μ_j absorbs all unobservable block-level characteristics that shift credit for all renters in the block, such as neighborhood amenities, market tightness, average landlord composition, and the average credit level of movers into the block. Our estimate of δ_g requires only that, across racial groups, any differences in preferences, *within a block*, are orthogonal to the local discrimination environment. This weaker condition is what we assess in the following subsection.

Threats to Interpretation

The barrier measure b_j^g is derived from a randomized correspondence experiment, in which inquirer race was randomly assigned across otherwise identical inquiries, providing experimentally grounded variation in measured discrimination intensity across blocks. By contrast, our credit outcomes are observational. We link the experimentally derived barrier measure

to the credit records of renters who moved into experimental blocks. The estimate δ_g comes from cross-block variation in the barrier, and our interpretation of it as discrimination-driven selection rests on the assumption that, within blocks, the own-race barrier is uncorrelated with unobserved preferences that affect credit score differences across racial groups.

The natural concern is sorting: renters who sort into blocks on credit in a way correlated with the barrier, for reasons unrelated to landlord behavior. Census block fixed effects absorb any block-level characteristics that shift credit for all renters in a block, so a story in which higher-barrier blocks attract higher- or lower-credit renters through amenities, prices, or desirability does not generate δ_g, provided such a channel shifts the credit of all groups in the block equally. In ongoing work, we additionally allow the credit gap to vary with landlords' response rate to white inquiries, which separates the racial gap in treatment from race-neutral screening by landlords who are unresponsive to all applicants.

The largest obstacle to a causal interpretation is a sorting channel that these features cannot rule out: one that operates on Black and Hispanic renters differently than on white renters, within block, and is correlated with the barrier. If, for reasons unrelated to landlord screening, high-credit Black or Hispanic renters are more likely than white renters to move into the blocks that happen to be high-barrier, perhaps because those blocks are whiter and entering them is selective on resources correlated with credit, then the Black/Hispanic-white credit gap would narrow with the barrier even absent any landlord credential screening. Our design cannot separate landlord-side selection from renter-side selective sorting of this specific form, and we therefore interpret δ_g as consistent with discrimination-driven credential screening rather than an estimate of its causal effect.

However, several observations bear on how much of δ_g this channel could plausibly account for. First, the experimental barrier reflects landlord behavior towards randomized inquiries and is constructed independently of the credit of renters who ultimately move in, so the relationship we estimate is not an artifact of the measure sharing inputs with the outcome. Second, for renter-side sorting to account for our estimates, the sorting must be generated from different preferences by race within a census block, and those different preferences must generate a credit score gap that tracks local discrimination intensity. Third, this within-block across-group preference difference must generate concentration in the upper tail of the credit distribution specifically, since the effect does not appear uniformly across credit tiers but is concentrated among renters with Very Good or Exceptional credit. To make the residual concern concrete: the confound would require an unobservable individual-level characteristic that, within the same census block, leads high-credit Black and Hispanic renters who move in to differ systematically from the high-credit white renters who move in, in ways that happen to be stronger where the own-race barrier is larger. Crucially, this

characteristic must manifest in credit-building behavior, concentrated in the upper tail of the credit distribution, without producing corresponding differences in occupation, education, or household composition, since the gradient is stable when we control for those observable attributes. While we cannot rule out an unobservable of this precise form, the conjunction of requirements, differential by race, unabsorbed by the census block fixed effect and individual-level controls, and concentrated in the upper tail of the credit distribution, makes it an increasingly specific and demanding joint condition rather than a generic sorting concern, and we interpret δ_g accordingly. Distinguishing these channels conclusively would require application-level data on who applied to each unit and who was ultimately selected, which remains a priority for future work.

Extensions

We extend the main specification in two ways to probe the mechanism underlying the credit gradient.

First, to examine whether the gradient effect is concentrated in the upper tail of the credit distribution, we replace the continuous credit score outcome with indicators for each credit tier and estimate the specification as a linear probability model. For each tier, the outcome equals one if the mover's credit score falls within that tier, using standard categories: Poor (below 580), Fair (580 to 669), Good (670 to 739), Very Good or Exceptional (740 and above). The coefficients δ_B and δ_H then measure how the probability that a Black or Hispanic mover falls in a given tier, relative to white movers in the same block, varies with the own-race barrier. A selection mechanism operating through a threshold predicts that the gradient is concentrated in the upper tiers rather than distributed uniformly across the credit score distribution.

Second, to examine how the gradient varies with neighborhood credit environment, we split the sample by whether a block's average credit score falls above or below the median block-average credit score and estimate the main specification separately within each subsample. This tests whether the credit gradient is stronger in neighborhoods with lower average credit scores, where a given credit score represents a larger deviation from the local applicant pool.

Both extensions retain the block and quarter-year fixed effects and the block-level clustering of the main specification, and both are identified from the same cross-block variation in the barrier measure.

Measurement Error

The barrier measure b_j^g is constructed from a limited number of experimental inquiries per block, covering a small share of the residential structures in a typical block. To the extent that the sampled listings are not perfectly representative of the block's overall discrimination environment, b_j^g is a noisy proxy for the true barrier facing renters in that block. Under classical measurement error, this attenuates δ_g toward zero.

A second source of measurement error arises from the temporal structure of the design. The barrier is measured during the window spanning February to July 2020, while our estimation sample spans multiple quarters before and after that window. If discrimination intensity varies over time, applying the measure from the experimental window to more distant quarters introduces additional attenuation. We find that the estimates for Hispanic movers decline only modestly as the estimation window widens and remain statistically significant throughout, while the precision sharpens for Black movers, indicating that the barrier measure is sufficiently stable over time to be applied across the panel and that our results are not an artifact of a narrow window around the experiment. For more information, see Table A1 in the appendix. Under classical measurement error, both sources attenuate $\hat{\delta}_g$ toward zero. To the extent that the inquiry-stage measure understates discrimination that occurs later in the application process, and that this understatement is correlated with landlord responsiveness, the error is not classical and the direction of bias is not guaranteed. We therefore treat the attenuation argument as suggestive rather than as a strict lower bound.

Results

Main Results

Table 2 presents our primary specification, equation (3). The coefficients of interest are the interactions between Black and Hispanic race indicators and the own-race barrier, $\hat{\delta}_B$ and $\hat{\delta}_H$, which measure how the credit gap between each group and white movers in the same block varies with the intensity of local discrimination.

The Hispanic interaction $\hat{\delta}_H$ is positive and statistically significant: the credit gap between Hispanic and white movers in the same census block is smaller where the own-race barrier is larger, consistent with prediction P1 (credit score gap narrowing where the barrier is larger). The Black interaction $\hat{\delta}_B$ is similar to the Hispanic interaction result, and is statistically significant. While $\hat{\delta}_B$ is about half the magnitude of $\hat{\delta}_H$, we cannot reject that the two estimates are equal. Columns 1 through 5 vary the estimation along several dimensions,

and the interaction estimate is stable throughout. Columns 1 to 3 widen the time window around the experimental window; the estimate attenuates modestly as the window widens, from about 10 points in the narrow window to about 7 points in the wide window, but remains significant, indicating that the result is not an artifact of a narrow window around the experiment and that the barrier measure is sufficiently stable over time to apply across the panel. Column 4 estimates the specification on all existing residents rather than just those who moved in within the last quarter, and the coefficient is nearly identical, differing by less than one credit point in the wide movers sample. We read this as evidence that the finding is not an artifact of how we define the mover sample: the same within-block relationship between credit and the barrier appears whether we examine the current flow of movers or the broader stock of residents. Column 5 adds individual-level controls, several of which may themselves be affected by the landlord selection process we study. We include them to show the estimate is stable to their inclusion.

Column 6 replaces the outcome with the mover's credit percentile within the block credit distribution, a relative rather than absolute measure of standing. The interaction is positive and significant for both groups, at the 1 percent level for Hispanic movers and the 10 percent level for Black movers, indicating that Black and Hispanic renters moving into higher-barrier blocks occupy higher relative positions in the local credit distribution. This relative measure corroborates the interpretation that what matters is a mover's standing relative to the local applicant pool.

Taken together, the main results are consistent with a partial credential offset in the rental housing market, paralleling labor market evidence that stronger credentials can offset part of a racial penalty in hiring (Gaddis, 2015). A key distinction is that we observe the offset through the actual financial characteristics of the renters who clear the barrier, rather than inferring it from hiring outcomes or callback rates. The credit profiles of successful movers are a more direct measure of credential screening than downstream labor market indicators, though they come at the cost of not observing the screening process itself.

Threshold Screening and Credit Tier Composition

If the partial credential offset operates through a screening threshold, as the labor market analog suggests, its effect should not distribute uniformly across the credit score distribution. Only renters with sufficiently strong credentials can clear a higher bar, so the offset should concentrate in the upper tail, among renters with Very Good or Exceptional credit, while renters with Poor credit should be underrepresented in high-barrier blocks relative to white renters. Table 3 estimates a linear probability model replacing the continuous credit outcome

with indicators for each credit tier: Poor (below 580), Fair (580-669), Good (670-739), and Very Good or Exceptional (740 and above). The results confirm this pattern, consistent with P2 (the offset concentrating in the upper tail).

For Hispanic movers, the interaction is positive and significant in the Very Good or Exceptional tier and negative and significant in the Poor tier: Hispanic movers become more likely to hold high credit and less likely to hold poor credit as the barrier rises, relative to white movers in the same block. The concentration in the tails rather than the middle of the distribution is precisely what a threshold mechanism predicts. The offset operates where credit is strong enough to clear a higher bar, and not where it sits far below it.

For Black movers, the upper-tier interaction is positive and significant, while the Poor-tier interaction is negative, though estimated with less precision. As the barrier rises, the Black–white gap in the probability of holding Very Good or Exceptional credit narrows, while the gap for Poor credit widens.

This distributional evidence adds something the labor market literature has rarely been able to show. Studies of employer discrimination typically observe callbacks or hiring outcomes but not the qualifications of those who clear the screen, making it difficult to see where in the distribution of credentials an offset operates. Observing the credit of successful movers lets us show that the offset is concentrated among the strongest credentials rather than spread across the distribution.

Neighborhood Credit Environment

Table 4 estimates the main specification separately for blocks above and below the median block-average credit score. The credential gradient is stronger in below-median blocks, consistent with P3 (the gradient being stronger in lower-credit neighborhoods), and the contrast is informative about why the offset may operate at all.

If credit functions as a signal to landlords not in absolute terms but relative to the local applicant pool, then when the surrounding population has lower average credit, a Very Good or Exceptional credit score stands further above the local mean. This provides the high-credit individual with a more powerful distinguishing signal in the landlord's screening decision. The barrier-countering benefit of a higher credit score is therefore larger precisely where the local credit environment makes it more informative.

This result reinforces the central finding of the paper. Black and Hispanic renters appear to partially offset racial barriers through credit, and that offset is most powerful where the signal carries the most weight. In lower-credit neighborhoods, a high credit score is not merely above average, it is exceptional relative to the local pool, and that exceptionalism is

what makes it effective against a higher screening threshold. The finding suggests that the partial credential offset is not uniform across the housing market but varies predictably with the information content of credit in a given neighborhood.

This heterogeneity connects to a broader point about when economic credentials are most effective as a partial offset against discrimination. Credit is not necessarily valuable in the abstract, but it is valuable relative to the pool against which a renter is evaluated. Where that pool has lower average credit scores, the signal is more powerful and offset is larger. This implies that the credential offset, to the extent it operates, is most accessible to high-credit Black and Hispanic renters in lower-credit neighborhoods, a pattern with implications for understanding where in the housing market the partial offset can and cannot reach.

Results by Group

Estimates for Black movers mirror those for Hispanic movers, but are about half as large and estimated with somewhat less precision: the main interaction is about 3.6 points and statistically significant at the 5 percent level. Since we cannot reject that the two estimates are equal, we do not interpret the difference in magnitudes as evidence that discrimination operates differently for the two groups. The smaller Black sample in California also limits precision, and we view the cross-group comparison as an open question rather than a central finding.

Conclusion

This paper provides evidence consistent with a partial credential offset in the rental housing market. The credit score gap between Black and Hispanic movers and white movers in the same census block is smaller where racial barriers are larger, the offset concentrates in the upper tail of the credit distribution among renters with Very Good or Exceptional credit, and it is strongest in neighborhoods where average credit is lower and a high score provides a more powerful relative signal. The pattern is consistent with a screening mechanism in which discrimination raises the effective bar for Black and Hispanic applicants, and strong credit can partially offset that bar when it is sufficiently strong.

These findings extend the study of credential offsets from labor markets to housing. In hiring, stronger credentials can offset part of a racial penalty (Gaddis, 2015); we find an analogous pattern among renters, with credit score functioning as the observable proxy for the broader financial qualifications that landlords screen on. The evidence here is more direct than what the labor market literature has typically been able to produce. Rather

than inferring credential selection from hiring outcomes or callback rates, we observe the actual credit profiles of the renters who secure housing. The cost of this directness is that we do not observe the screening process itself, only its imprint on who moves in.

These findings bear on how discrimination operates. Strong credit can partially offset a higher screening bar, but the offset is limited: it operates only for sufficiently strong credit and is larger where that credit stands out from the local pool. Discrimination therefore still shapes access to housing even where it is partially overcome.

To the extent these patterns reflect landlord screening, they bear on policies governing how landlords use credit. If strong credit is one channel through which Black and Hispanic renters clear a higher bar, restricting its use in tenant screening could close that channel for the renters currently using it, much as removing criminal-record information from job applications appears to have increased statistical discrimination in labor markets (Agan and Starr, 2018; Doleac and Hansen, 2020).[1] These potential costs must be weighed against the disparate impact of credit screening itself, since credit scores carry racial gaps rooted in structural inequality (Bhutta et al., 2025).

Two limitations bound our interpretation. First, our design establishes a within-block relationship between credit and discrimination intensity but cannot fully separate landlord screening from differential renter sorting. The sorting story that survives our design requires a specific and demanding joint condition: unobserved individual characteristics that generate upper-tail credit concentration, operate differentially by race within block, and track the barrier specifically after absorbing block fixed effects and individual controls. We interpret δ_g as consistent with the discrimination-driven credential screening while acknowledging this residual. Establishing the causal role of landlord screening conclusively would require application-level data on who applied and who was selected, which remains a priority for future work. Second, the analysis is restricted to renters in California. The six markets span a meaningful range of the discrimination intensity documented across the fifty cities in Christensen et al. (2021), and we see no reason the mechanism would be California-specific, but we cannot test external validity directly. Whether the partial credential offset operates similarly in markets with different housing stock compositions and rental structures remains open.

Beyond these limitations, this paper opens two directions for future work. First, whether a partial credential offset shapes where Black and Hispanic renters ultimately sort, and whether it reinforces or attenuates residential segregation, remains an open question with

[1]This implication holds whether discrimination is taste-based or statistical: in either case, removing credit eliminates a signal high-credit applicants use to overcome a higher bar. Its magnitude depends on how readily landlords substitute toward other signals, such as income, rental history, or references, and on whether those substitutes are noisier or more open to discretionary bias.

direct implications for neighborhood inequality. Second, the panel structure of the credit data used here enables research on whether clearing a discriminatory barrier has lasting effects on renters' financial trajectories, a dynamic dimension of the costs of housing discrimination that has received little attention. More broadly, the findings suggest that offsetting credentials operate in rental housing markets under discrimination, extending a mechanism long studied in labor markets to a setting where its implications for neighborhood access and residential inequality are direct.

References

Agan, Amanda, and Sonja Starr. 2018. "Ban the Box, Criminal Records, and Racial Discrimination: A Field Experiment." *Quarterly Journal of Economics* 133 (1): 191–235.

Arrow, Kenneth J. 1973. "The Theory of Discrimination." *Discrimination in Labor Markets* 3–33.

Bartlett, Robert, Adair Morse, Richard Stanton, and Nancy Wallace. 2022. "Consumer-lending discrimination in the FinTech Era." *Journal of Financial Economics* 143 (1): 30–56. 10.1016/j.jfineco.2021.05.047.

Becker, Gary S. 1957. *The Economics of Discrimination*. Chicago: University of Chicago Press.

Berkovec, James A., Glenn B. Canner, Stuart A. Gabriel, and Timothy H. Hannan. 1994. "Race, Redlining, and Residential Mortgage Loan Performance." *The Journal of Real Estate Finance and Economics* 9 (3): 263–294. 10.1007/BF01099279.

Bertrand, Marianne, and Sendhil Mullainathan. 2004. "Are Emily and Greg More Employable Than Lakisha and Jamal? A Field Experiment on Labor Market Discrimination." *American Economic Review* 94 (4): 991–1013. 10.1257/0002828042002561.

Bhutta, Neil, and Aurel Hizmo. 2021. "Do Minorities Pay More for Mortgages?" *The Review of Financial Studies* 34 (2): 763–789. 10.1093/rfs/hhaa047.

Bhutta, Neil, Aurel Hizmo, and Daniel Ringo. 2025. "How Much Does Racial Bias Affect Mortgage Lending? Evidence from Human and Algorithmic Credit Decisions." *The Journal of Finance* 80 (3): 1463–1496. 10.1111/jofi.13444.

Chetty, Raj, Nathaniel Hendren, and Lawrence F. Katz. 2016. "The Effects of Exposure to Better Neighborhoods on Children: New Evidence from the Moving to Opportunity Experiment." *American Economic Review* 106 (4): 855–902. 10.1257/aer.20150572.

Christensen, Peter, Ignacio Sarmiento-Barbieri, and Christopher Timmins. 2021. "Racial Discrimination and Housing Outcomes in the United States Rental Market." Working Paper29516, National Bureau of Economic Research. 10.3386/w29516.

Christensen, Peter, and Christopher Timmins. 2022. "Sorting or Steering: The Effects of Housing Discrimination on Neighborhood Choice." *Journal of Political Economy* 130 (8): 2110–2163. 10.1086/720140.

Christensen, Peter, and Christopher Timmins. 2023. "The Damages and Distortions from Discrimination in the Rental Housing Market." *The Quarterly Journal of Economics* 138 (4): 2505–2557. 10.1093/qje/qjad029.

Dobbie, Will, Andres Liberman, Daniel Paravisini, and Vikram Pathania. 2021. "Measuring Bias in Consumer Lending." *Review of Economic Studies* 88 (6): 2799–2832.

Doleac, Jennifer L., and Benjamin Hansen. 2020. "The Unintended Consequences of "Ban the Box": Statistical Discrimination and Employment Outcomes When Criminal Histories Are Hidden." *Journal of Labor Economics* 38 (2): 321–374.

Elliott, Marc N., Peter A. Morrison, Allen Fremont, Daniel F. McCaffrey, Philip Pantoja, and Nicole Lurie. 2009. "Using the Census Bureau's Surname List to Improve Estimates of Race/Ethnicity and Associated Disparities." *Health Services and Outcomes Research Methodology* 9 (2): 69–83.

Ewens, Michael, Bryan Tomlin, and Liang Choon Wang. 2014. "Statistical Discrimination or Prejudice? A Large Sample Field Experiment." *Review of Economics and Statistics* 96 (1): 119–134.

Gaddis, S. Michael. 2015. "Discrimination in the Credential Society: An Audit Study of Race and College Selectivity in the Labor Market." *Social Forces* 93 (4): 1451–1479.

Hanson, Andrew, and Zackary Hawley. 2011. "Do Landlords Discriminate in the Rental Housing Market? Evidence from an Internet Field Experiment in US Cities." *Journal of Urban Economics* 70 (2–3): 99–114.

Holzer, Harry J., and Keith R. Ihlanfeldt. 1998. "Customer Discrimination and Employment Outcomes for Minority Workers." *The Quarterly Journal of Economics* 113 (3): 835–867.

Kline, Patrick, Evan K Rose, and Christopher R Walters. 2022. "Systemic Discrimination Among Large U.S. Employers." *The Quarterly Journal of Economics* 137 (4): 1963–2036. 10.1093/qje/qjac024.

Phelps, Edmund. 1972. "The Statistical Theory of Racism and Sexism." *American Economic Review* 62 659–61.

Yinger, John. 1995. *Closed Doors, Opportunities Lost: The Continuing Costs of Housing Discrimination*. New York: Russell Sage Foundation.

Table 1: Summary Statistics by Racial Group

	White	Hispanic	Black
Panel A: Movers			
N	272,458	193,348	31,166
Credit Score	675.3 (99.3)	631.5 (90.4)	597.2 (93.6)
Own-Race Response Rate	0.691 (0.455)	0.617 (0.482)	0.484 (0.493)
Age	39.4 (15.8)	38.6 (13.8)	40.3 (15.3)
Panel B: Residents			
N	3,100,135	2,786,150	342,368
Credit Score	680.0 (97.1)	638.8 (87.7)	608.2 (91.6)
Own-Race Response Rate	0.659 (0.468)	0.660 (0.470)	0.586 (0.485)
Age	44.2 (17.7)	41.2 (14.6)	43.9 (16.7)
Mobilityj Rate	0.179 (0.146)	0.152 (0.138)	0.174 (0.137)

Notes: Means with standard deviations in parentheses, by racial group. Panel A reports movers into experimental census blocks. Panel B reports existing residents. Own-race reponse rate is the block-level average response rate to inquiries from fictitious renters of that group, from Christensen et al. (2021). Mobility rate is the person-level share of quarters in which a census block change was observed, averaged within racial group.

Table 2: Credit Score Gap and Own-Race Discrimination Intensity

	Credit Score				Credit Percentile
	(1)	(2)	(3)	(4)	(5)
$\mathbf{1}[\text{Black}_i] \times \text{OwnRaceBarrier}_j$	6.308	5.975*	4.762**	3.602**	0.012*
	(4.221)	(3.402)	(1.970)	(1.754)	(0.007)
$\mathbf{1}[\text{Hispanic}_i] \times \text{OwnRaceBarrier}_j$	9.725***	8.311***	6.913***	6.688***	0.025***
	(2.801)	(2.388)	(1.561)	(1.548)	(0.005)
Block FE	Yes	Yes	Yes	Yes	Yes
QY FE	Yes	Yes	Yes	Yes	Yes
Race FE	Yes	Yes	Yes	Yes	Yes
Controls	No	No	No	Yes	No
Observations	41,439	90,924	496,972	496,972	496,972

Notes: Estimates of equation (3). The outcome is credit score, except in column 6, where it is the mover's percentile in the block credit distribution. The sample is movers into experimental blocks, except column 4, which uses all block residents. Columns 1-3 widen the estimation window from narrow to wide. Column 5 adds individual controls (age, marital status, household composition, education, occupation) as a robustness check; several may be endogenous to the selection process, so our preferred specification omits them. All columns include census block, quarter-year, and race fixed effects. Standard errors clustered by census block in parentheses.

$^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

Table 3: Credit Tier Relative Probabilities
Movers, No Controls (2011-2026)

	Poor (<580)	Fair (580–669)	Good (670–739)	VGood or Exc (740+)
$\mathbf{1}[\text{Black}_i] \times \text{OwnRaceBarrier}_j$	-0.019*	0.001	0.002	0.016***
	(0.011)	(0.008)	(0.006)	(0.006)
$\mathbf{1}[\text{Hispanic}_i] \times \text{OwnRaceBarrier}_j$	-0.025***	-0.003	0.008*	0.020***
	(0.006)	(0.004)	(0.004)	(0.005)
Block FE	Yes	Yes	Yes	Yes
QY FE	Yes	Yes	Yes	Yes
Race FE	Yes	Yes	Yes	Yes
Controls	No	No	No	No
Observations	496,972	496,972	496,972	496,972

Notes: Each column reports a linear probability model estimated from equation (3) with an indicator for the named credit tier as the outcome. Tiers are Poor (below 580), Fair (580-669), Good (670-739), and Very Good or Exceptional (740 and above). The sample is movers into experimental census blocks. Census block, quarter-year, and race fixed effects are included. Standard errors clustered by census block in parentheses.
$^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

Table 4: Credit Score Gap and the Local Credit Environment

	Credit Score		Credit Percentile	
	(1)	(2)	(3)	(4)
$\mathbf{1}[\text{Black}_i] \times \text{OwnRaceBarrier}_j$	6.438**	1.551	0.020**	0.003
	(2.667)	(2.607)	(0.009)	(0.008)
$\mathbf{1}[\text{Hispanic}_i] \times \text{OwnRaceBarrier}_j$	8.044***	4.579***	0.029***	0.014***
	(2.053)	(1.568)	(0.007)	(0.005)
Block FE	Yes	Yes	Yes	Yes
QY FE	Yes	Yes	Yes	Yes
Race FE	Yes	Yes	Yes	Yes
Controls	No	No	No	No
Observations	242,676	254,296	242,676	254,296

Notes: Estimates of equation (3) on subsamples split by whether the block's average credit score falls below or above the median block-average credit score among experimental blocks. The outcome in columns 1-2 is credit score, while in columns 3-4 the outcome is the percentile ranking of mover i relative to the block's credit distribution. The sample is movers into experimental census blocks. Census block, quarter-year, and race fixed effects are included. Standard errors clustered by census block in parentheses. The median split is intended to capture variation in the informativeness of credit as a signal to landlords: in below-median blocks, where the local credit distribution is left-shifted, a given credit score represents a larger deviation above the local pool.

$^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

Appendix

Conceptual Framework

When a rental unit becomes available, a landlord selects among applicants according to an objective function that depends on applicant characteristics. We model this selection process as:

$$s_j(g_i, z_i) = \lambda_j \cdot \mathbf{1}[g_i = W] + \beta_j \cdot z_i \tag{4}$$

where $g_i \in \{W, H, B\}$ denotes the perceived race of applicant i, $z_i \in \mathbb{R}$ is a scalar index of applicant quality characteristics observable to the landlord (including characteristics such as credit score, rental history, and income), and $\lambda_j \geq 0$ and $\beta_j > 0$ are landlord-specific weights. The landlord allocates the unit to the applicant who maximizes s_j. The weight β_j is strictly positive, reflecting that higher quality applicants are always weakly preferred, all else equal. The weight λ_j captures the landlord's preference for a white tenant above and beyond the quality of the tenant: when $\lambda_j = 0$ the landlord is race-neutral, and when $\lambda_j > 0$ the landlord places additional value on selecting a white applicant.

In a race-neutral neighborhood ($\lambda_j = 0$), credit enters the selection function identically regardless of race and selection is determined entirely by z_i. All applicants compete without a race-specific advantage/disadvantage, and the observed quality characteristics of successful movers reflect the upper tail of each group's quality distribution. Any difference in average observed quality between Black or Hispanic renters and white renters who successfully move in, what we estimate as $\hat{\alpha}_g$ in our empirical specification, reflects baseline differences in the distribution of z across groups rather than discrimination.

When the landlord places positive weight on race ($\lambda_j > 0$), a Black or Hispanic application must overcome a quality premium to win the selection process. Specifically, a Black or Hispanic applicant wins selection over a white applicant when:

$$\beta_j \cdot z_i^M > \lambda_j + \beta_j \cdot z_i^W \tag{5}$$

which rearranges to the threshold condition:

$$z_i^M - z_i^W > \frac{\lambda_j}{\beta_j} \tag{6}$$

The minority applicant must exceed the white applicant's quality index by at least λ_j/β_j to win the unit. This ratio is the quality premium imposed by discrimination: it is increasing in the landlord's racial preference λ_j and decreasing in the weight placed on quality β_j. We use

credit score as the primary observable proxy for z_i, so λ_j/β_j represents the credit premium a minority applicant must present above the competing white applicant's credit score in order to be selected.

The threshold character of equation 6 generates a non-linear prediction about the credit quality among successful minority movers. A minority renter who is only marginally better than the competing white applicant on credit quality cannot clear the threshold λ_j/β_j in discriminatory blocks; marginal improvements in credit below the threshold provide no benefit. A minority renter who is substantially above the competing white applicant on credit quality can reliably clear the threshold. The credential offset is therefore concentrated in the upper tail of the quality distribution, where the threshold can be cleared, rather than distributed evenly across the credit score distribution.

The total quality hurdle a minority renter must overcome reflects two additive components. The first is the baseline quality gap between groups: if minority renters draw from a quality distribution with mean μ^M while white renters draw from a distribution with mean $\mu^W > \mu^M$, the minority renter faces a more difficult selection problem even in the absence of discrimination, because the best white applicant in the pool is drawn from a higher distribution. This baseline component is captured by $\hat{\alpha}_g$ in our empirical specification. The second component is the discrimination-driven quality premium λ_j/β_j, captured by $\hat{\delta}_g$. Expressing each applicant's quality as a deviation from their group mean, the minority applicant wins selection when:

$$(z_i^M - \mu^M) > (z_i^{W,max} - \mu^W) + (\mu^W - \mu^M) + \frac{\lambda_j}{\beta_j} \tag{7}$$

Reading from right to left, a minority renter must overcome three terms: the discrimination premium λ_j/β_j, the baseline quality gap between group means $(\mu^W - \mu^M)$, and the draw of the best white applicant above the white mean $(z_i^{W,max} - \mu^W)$. Since credit score is the primary observable component of z_i, the share of minority renters positioned to clear the threshold is empirically proxied by the share of that group's credit distribution in the upper tail.

Finally, if landlords evaluate applicant quality relative to the local applicant pool rather than on an absolute scale, the threshold condition implies that the credential offset will be larger in neighborhoods where the local quality distribution is left-shifted. In neighborhoods where the average quality index is lower, a given absolute credit score represents a larger deviation above the local mean, providing a more powerful relative signal to the landlord. The credential offset is therefore predicted to be stronger in, for example, neighborhoods with lower block-average credit scores.

Together, the selection function framework generates four testable predictions:

MP1. The credit score gap between successful minority and white movers narrows as λ_j increases, reflecting the higher bar minority applicants must clear in more discriminatory blocks.

MP2. The credential offset is concentrated in the upper tail of the credit distribution, where the threshold λ_j/β_j can be cleared; renters with Poor credit face an amplified disadvantage. Renters with Poor credit are further from the threshold and face a correspondingly larger access disadvantage.

MP3. The credential offset is larger in neighborhoods where the population has lower average credit scores, where a given absolute credit score provides a stronger relative signal.

MP4. In high-barrier blocks, credit score is more predictive of minority identity among successful movers, because minority applicants must clear a higher quality threshold to be selected.

Robustness to the Estimation Window

Because the discrimination measure is drawn from the experimental window spanning February to July 2020 while the credit panel spans a longer period, the choice of estimation window trades off two concerns. A narrow window centered on the experiment minimizes any drift between the measured barrier and the discrimination environment renters actually faced, but restricts the sample and reduces precision. A wider window increases the sample but applies the barrier measure to quarters farther from the experiment, introducing additional measurement error if discrimination intensity varies over time.

Table A1 reports the main specifications across estimation windows ranging from five quarters to 14 (asymmetric) years around the experimental period. The interaction estimate $\hat{\delta}_H$ attenuates modestly as the window widens but remains positive and statistically significant throughout. The stability of the estimate indicates that the result is not an artifact of a particular window choice and that the barrier measure is sufficiently stable over time to be applied across the panel. The modest attenuation in wider windows is consistent with the measurement error discussion in the Empirical Strategy section: applying a barrier measured in 2020 to more distant quarters adds noise, which would tend to bias the estimate toward zero if that noise is unrelated to renters' credit.

Table A1: Credit Score Gap by Increasing Quarter-Year Windows

	Credit Score				
	(1)	(2)	(3)	(4)	(5)
Panel A: Movers					
$\mathbf{1}[\text{Black}_i] \times \text{OwnRaceBarrier}_j$	6.308	5.975*	5.158*	5.531**	4.762**
	(4.221)	(3.402)	(2.781)	(2.453)	(1.970)
$\mathbf{1}[\text{Hispanic}_i] \times \text{OwnRaceBarrier}_j$	9.725***	8.311***	7.246***	7.925***	6.913***
	(2.801)	(2.388)	(2.015)	(1.736)	(1.561)
Observations	41,439	90,924	173,593	314,255	496,972
Panel B: Residents					
$\mathbf{1}[\text{Black}_i] \times \text{OwnRaceBarrier}_j$	1.346	1.139	0.130	0.890	1.779
	(3.585)	(3.462)	(2.982)	(2.663)	(2.140)
$\mathbf{1}[\text{Hispanic}_i] \times \text{OwnRaceBarrier}_j$	10.705***	9.581***	8.118***	7.466***	6.619***
	(2.626)	(2.549)	(2.137)	(1.940)	(1.656)
Observations	648,654	1,379,720	2,853,752	4,326,235	6,228,653
CB FE	Yes	Yes	Yes	Yes	Yes
QY FE	Yes	Yes	Yes	Yes	Yes
Race FE	Yes	Yes	Yes	Yes	Yes
Controls	No	No	No	No	No

Notes: Estimates of equation (3) across estimation windows of increasing width around the experimental period. The outcome is credit score. The sample is movers into experimental census blocks in Panel A, and all residents of those experimental blocks in Panel B. Census block, quarter-year, and race fixed effects are included. Standard errors clustered by census block in parentheses.
$^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.