

Racial Discrimination in Eviction Filing

Alison Lodermeier*

March 1, 2026

Abstract

This paper develops and applies a novel framework to test whether landlords discriminate when filing eviction cases. Discrimination is measured by racial disparities in back rent owed at the time of filing, conditional on landlord and contract rent. Using rich administrative data from Philadelphia, I find that landlords discriminate by tolerating 4.9% less back rent from minorities before filing. To distinguish the sources of discrimination, I conduct a marginal outcome test and find evidence consistent with accurate statistical discrimination: landlords win cases at indistinguishable rates by tenant race, suggesting they accurately predict that successful minority cases require less back rent.

*University of Chicago (email: alodermeier@uchicago.edu). I thank Anna Aizer, Peter Hull, Matthew Pecenco, Neil Thakral, Matthew Turner, John Friedman, Jesse Bruhn, Emily Oster, Carlos Fernando Avenancio-León, Niles Fernando, Tess Scharlemann, Jaclene Begley, and numerous seminar and conference participants for valuable feedback and suggestions. I also thank the Civil Justice Data Commons at Georgetown University and Jonathan Pyle at Philadelphia Legal Assistance for providing access to data. The Orlando Bravo Center for Economic Research at Brown University provided financial support.

1 Introduction

Each year, landlords file 3.6 million eviction cases in the U.S. (Gromis et al. 2022), placing families at risk of homelessness. Minority renters are disproportionately exposed to this risk as they are nearly three times as likely to face an eviction case as white renters (Graetz et al. 2023).¹ This disparity could be explained by racial differences in lease violations, racial differences in matching to strict landlords, or landlords discriminating by tolerating lease violations differently by tenant race. Understanding these potential sources—and in particular whether tenants face discrimination in eviction decisions—has important implications for addressing eviction disparities. This paper develops and applies novel methods to test whether landlords exhibit racial discrimination when evicting tenants and distinguish between potential sources of discrimination.

When measuring discrimination, researchers must account for differences in the underlying qualification of individuals for a treatment. In most settings, qualification is at best partially unobserved, motivating strategies that leverage random assignment to decision-makers (Arnold et al. 2018; Dobbie et al. 2021; Arnold et al. 2022; Baron et al. 2024), structural models of decision-making (Knowles et al. 2001; Anwar and Fang 2006), control groups of non-discriminatory decision-makers (Goncalves and Mello 2021), or the ability to experimentally vary fictitious individuals’ race (Bertrand and Mullainathan 2004; Ewens et al. 2014; Kline et al. 2022). The context of eviction, however, is uniquely well-suited to study discrimination given the plausible assumption that a tenant’s qualification for an eviction filing (caused only by non-payment) is fully summarized by the amount of back rent owed at the time of the filing. Using detailed eviction filing data from Philadelphia, I develop a novel test for discrimination that leverages this comprehensive and observable measure of qualification.

I consider a simple framework in which a landlord tolerates some non-payment of rent before filing an eviction case. This means a landlord has some threshold of back rent for filing an eviction case against a particular tenant, absent any other lease violations. Under certain assumptions, differences in landlords’ filing thresholds by tenant race can be interpreted as evidence of racial discrimination. Legally, this constitutes a Fair Housing Act violation if there exists no legitimate reason for a landlord to selectively file eviction cases against tenants equally behind on rent who have committed no other lease violations. A restated version of this assumption—that a tenant’s qualification for an eviction filing (over non-payment only)

¹While an estimated 50.5% of individuals living in rental housing in the U.S. are white and 49.5% are racial or ethnic minorities, 26.3% of individuals living in households facing eviction filings are white and 73.7% are racial or ethnic minorities (Graetz et al. 2023). This equates to a minority eviction filing rate that is 2.9 times higher than the white eviction filing rate.

is fully summarized by the back rent owed—aligns this disparity with the concept of “total discrimination” by race defined in [Bohren et al. \(2022\)](#).²

To derive my test for racial discrimination in eviction filing, I leverage the assumption that a tenant gradually falls behind on rent until the amount of back rent reaches their landlord’s filing threshold. The average racial disparity in filing thresholds is identified by racial disparities in back rent owed at the time of eviction filings, restricting to cases caused only by non-payment of rent. To allow a landlord’s filing threshold of back rent to be measured in units of dollars owed or months of back rent owed, I condition these disparities on landlord and monthly contract rent.³ Assuming dollars or months of back rent fully captures a tenant’s qualification for a non-payment based eviction, I interpret this conditional racial disparity as total racial discrimination in eviction filing.

I apply this framework to the universe of residential eviction cases filed in Philadelphia between 2006 and 2019. Since race is not reported in eviction cases, I impute tenant race using first names, last names, and census blocks, and verify that the methodology is well-calibrated for individuals living in Philadelphia over this time period.⁴ Comparing back rent owed by white and minority tenants at the time of filing conditional on landlord and contract rent, I find evidence of racial discrimination. On average, landlords tolerate 4.9% (\$131) less back rent from minority tenants than from white tenants before filing an eviction case. This disparity cannot be explained by observable non-race characteristics of tenants, including proxy measures of future rent payment ability, which is consistent with landlords discriminating directly on the basis of tenant race as opposed to other characteristics correlated with race.⁵ These findings are robust to using various other imputed measures of tenant race and bounding the conditional disparities to account for race measurement error following [Elzayn et al. \(2025\)](#). I also find statistically significant evidence of discrimination across narrower definitions of landlords and coarser conditioning on contract rent.

Next, I extend the framework to distinguish between the sources of discrimination. It may be driven by accurate statistical discrimination arising from actual racial differences in

²[Bohren et al. \(2022\)](#) define total racial discrimination as racial differences in treatment (e.g. eviction filing) conditional on qualification (e.g. back rent).

³I consider months of back rent owed as the ratio of back rent to monthly contract rent. Given that distressed renters often make partial rent payments ([Manville et al. 2023](#)), a tenant may, for example, accrue two months of back rent over a longer period of time than two months.

⁴I compare observed and predicted race for borrowers on all mortgages originated in Philadelphia between 2007 and 2019 using CoreLogic deeds data linked to the Home Mortgage Disclosure Act (HMDA) data.

⁵If back rent realized at the time of filing fully captures eviction qualification, any differences arising from these non-race characteristics would amount to systemic (or indirect) discrimination in [Bohren et al. \(2022\)](#). Alternatively, if predictions about a tenant’s future ability to pay (coming from non-race factors other than current back rent) is considered another legitimate or qualifying reason for eviction, then controlling for payment ability is necessary to interpret racial disparities in filing thresholds as racial discrimination.

the landlord’s payoff from filing an eviction case (Phelps 1972; Arrow 1974; Aigner and Cain 1977) or landlords having racially biased beliefs or preferences (Becker 1957; Bordalo et al. 2016; Bohren et al. 2025).⁶ Landlords may exhibit accurate statistical discrimination if they observe that winning repossession of the property in court from minority tenants tends to require less back rent than cases against white tenants. The data indeed reflect this potential motivation for statistical discrimination: landlords are more likely to win repossession from minority tenants than white tenants with comparable back rent.⁷ If landlords exhibit accurate statistical discrimination, then they should be equally likely to win repossession in marginal cases against white and minority tenants. Any deviation from equal repossession rates reflects racial bias in eviction filing, where landlords are biased against the group for whom they tolerate a lower repossession rate.

I formalize this logic with a marginal outcome test for racial bias in eviction filing (Canay et al. 2024; Hull 2021), which follows recent evaluations of bias in other settings such as vehicle searches, capital sentencing, parole, bail, lending, and foster care (Knowles et al. 2001; Persico and Todd 2006; Alesina and La Ferrara 2014; Anwar and Fang 2015; Arnold et al. 2018; Marx 2021; Dobbie et al. 2021; Baron et al. 2024). I conduct my test by comparing a landlord’s rate of winning repossession in white and minority cases with the same contract rent. The test is valid when landlords file eviction cases at their threshold of back rent (marginal outcomes) and a landlord’s true payoff from winning repossession of the property does not vary with tenant race conditional on control variables (no omitted payoff bias).

I find no evidence that landlords exhibit racial bias against minority tenants. Landlords achieve statistically indistinguishable rates of winning repossession in court against white and minority tenants, and if anything landlords win repossession slightly *more* often from minority tenants which allows me to reject small degrees of anti-minority bias. This finding is consistent with landlords exhibiting accurate statistical discrimination and accounting for racial disparities at later stages of the eviction process when deciding to file a case. An implication of this result is that policies to eliminate racial disparities in eviction court could reduce earlier discrimination by landlords filing eviction cases. Further upstream, however, it is possible that landlords would adjust how they screen potential tenants for rental units which could disproportionately harm minority renters.

While an extensive literature on housing market disparities provides evidence of racial discrimination during the housing search process, this paper contributes evidence in a new

⁶Racial bias includes taste-based discrimination/animus based on biased preferences and inaccurate statistical discrimination based on incorrect beliefs.

⁷This pattern could reflect racial differences in whether tenants appear in court, unequal access to legal representation, racial disparities in bargaining power during landlord-tenant negotiations, discrimination by judges, or some other factor.

and high-stakes context: housing loss. In the housing search literature, correspondence studies that send fictitious rental applications to landlords generally find lower callback rates for applicant names with a high likelihood of association with a minority group (Carpusor and Loges 2006; Ahmed and Hammarstedt 2008; Hanson and Hawley 2011; Ewens et al. 2014; Christensen et al. 2021).⁸ Other research finds that minority homebuyers face discrimination when seeking mortgage loans (Munnell et al. 1996; Ladd 1998; Turner and Skidmore 1999; Ross et al. 2008; Hanson et al. 2016). Despite the known adverse effects of eviction and foreclosure (Diamond et al. 2020; Collinson et al. 2024, 2025), much less is known about the role of race in the context of housing loss.⁹ I contribute new evidence that landlords tolerate non-payment differently by tenant race, indicating that minority renters face higher exposure to housing loss than white renters.

Methodologically, this paper contributes to a recent literature developing new tools to measure discrimination. While correspondence studies estimate discrimination by experimentally varying fictitious applicant names that are distinctly associated with different groups (Bertrand and Mullainathan 2004; Jacquemet and Yannelis 2012; Ewens et al. 2014; Edelman et al. 2017; Agan and Starr 2018; Kline et al. 2022), measuring discrimination is generally more challenging in observational data on high-stakes decisions such as eviction. Researchers must account for racial differences in the underlying qualification of individuals for a treatment, but qualification is at best partially unobserved in most settings. Strategies to overcome this problem leverage random assignment to decision-makers (Arnold et al. 2018; Dobbie et al. 2021; Arnold et al. 2022; Baron et al. 2024), structural models of decision-making (Knowles et al. 2001; Anwar and Fang 2006), control groups of non-discriminatory decision-makers (Goncalves and Mello 2021), or observable outcomes from both random and algorithmic decision-making (Elzayn et al. 2025). Across this literature, these methods have overwhelmingly been developed for static decisions such as traffic stops (Knowles et al. 2001; Anwar and Fang 2006; Goncalves and Mello 2021), bail (Arnold et al. 2018, 2022), hiring (Bertrand and Mullainathan 2004), and lending (Dobbie et al. 2021), where the decision-maker first encounters the individual subject to the decision at the time treatment is assigned.

Eviction filings represent a dynamic decision by landlords who observe tenants over time and decide if and when to file a case. This paper contributes a novel approach to measure discrimination in dynamic settings by leveraging the plausible assumption that back rent

⁸A related literature uses in-person audits to test for discrimination against prospective homebuyers (Yinger 1986; Page 1995).

⁹In survey data from Milwaukee, Greenberg et al. (2016) find that among renters habitually late on rent, Hispanic renters are significantly more likely to be evicted. However, without quantifying the severity of non-payment or other lease violations among these renters, it is difficult to distinguish landlord discrimination from different levels of eviction-warranting behavior across tenants.

fully captures a tenant’s qualification for an eviction filing (over non-payment only). My approach relies on two key features of eviction filing decisions. First, the measure of an individual’s qualification (back rent) for a given treatment (eviction filing) accrues gradually over time. Second, qualification is observable at the time of treatment. Since treatment is assigned the moment an individual’s qualification crosses the decision-maker’s threshold, discrimination equates to racial differences in qualification at the time of treatment. This approach requires only data on treated units (i.e. filed eviction cases) and does not require random assignment to decision-makers, allowing it to be applied to a wide range of other important markets.¹⁰ In settings where the outcome of the treatment decision (e.g. winning court-ordered repossession) is also observable, it is possible to distinguish the sources of discrimination using a marginal outcome test. This implementation of the marginal outcome test avoids the common “infra-marginality” problem of outcome tests (Ayres 2002) since the gradual accrual of an individual’s qualification up to the decision-maker’s threshold means that all individuals are treated at the margin.¹¹

The remainder of this paper is organized as follows. Section 2 details the legal framework governing housing discrimination and the Philadelphia eviction process. Section 3 describes the Philadelphia eviction court records used in the analysis. Section 4 models a landlord’s eviction filing decision and formalizes discrimination concepts in this setting. Section 5 outlines the empirical approach. Section 6 reports estimates of racial discrimination in eviction filing. Section 7 distinguishes between the sources of racial discrimination. Section 8 concludes.

2 Institutional Context

2.1 Fair Housing Act

The foundation of anti-discrimination law in the U.S. housing market is the Fair Housing Act. It prohibits discrimination by housing providers, such as landlords, whose discriminatory practices make housing unavailable on the basis of race, religion, sex, national origin,

¹⁰Other high-stakes settings where these new tools may be applied include worsening symptoms before a physician begins medical treatment, school absenteeism before an educator imposes disciplinary action, debt accrual before a creditor sells to a collection agency, or workplace misconduct leading up to an employee being fired.

¹¹The “infra-marginality” problem refers to the fact that racial differences in average outcomes cannot inform differences in marginal outcomes if the racial groups have different underlying risk distributions. The intuition of my marginal outcome test is most similar to Anwar and Fang (2015) which tests for racial bias in parole decisions under the assumption that all inmates released by a Parole Board are released at the moment their potential recidivism rate is lowered enough via time spent incarcerated to hit the Parole Board’s threshold for release.

family status, or disability. When considering evictions in particular, the Department of Housing and Urban Development (HUD) states that “Selective evictions because of protected class (such as race or national origin) violate the Fair Housing Act even when the eviction might otherwise have been lawful” (HUD 2021a). More explicitly, HUD provides guidance to tenants that “Even if you are behind on your rent and subject to eviction, a landlord may not pick and choose which tenants to evict based on any protected characteristic” (HUD 2021b). They offer the following example of discrimination: “A landlord evicts a Black tenant for unpaid rent but does not evict a White tenant who has unpaid rent” (HUD 2021b).

These statements suggest a landlord may be in violation of the Fair Housing Act if they tolerate lease violations differently by tenant race before filing an eviction case.¹² When non-payment of rent is the only lease violation that has occurred, this is equivalent to a landlord using different filing thresholds of back rent for white and minority tenants. Since landlords may reasonably measure back rent in dollars or months owed, I focus on racial disparities in filing thresholds of back rent among tenants with the same landlord and monthly contract rent.¹³

To understand the conditions under which this conditional disparity can be interpreted as unlawful discrimination under the Fair Housing Act, it is important to understand the three potential sources of the disparity. First, landlords may incorporate non-race characteristics correlated with tenant race into filing decisions, which can generate different filing thresholds by race. Second, accurate statistical discrimination can lead to racial disparities in thresholds if landlords account for actual racial differences in their payoff from filing (Phelps 1972; Arrow 1974; Aigner and Cain 1977). Third, racial bias can generate racial disparities in thresholds stemming from landlords’ biased preferences or beliefs (Becker 1957; Bohren et al. 2025; Bordalo et al. 2016). I formalize these sources and demonstrate how they can generate racial disparities in eviction filing thresholds in Section 4.

The Fair Housing Act protects against selective evictions driven by racially discriminatory intent and against selective evictions that have a disparate impact on racial minorities if they were not necessary to achieve a legitimate interest for the landlord (HUD 2023). While selective evictions driven by discriminatory intent reflect only direct discrimination on the basis of race (accurate statistical discrimination or racial bias), the broader concept

¹²HUD has issued discrimination charges consistent with this interpretation. In 2012, HUD charged a landlord with discrimination in violation of the Fair Housing Act for attempting to evict a Black tenant in response to a disorderly conduct incident, alleging that the landlord did not attempt to evict white tenants with equally or more serious lease violations. See *United States v. Altoona Housing Authority* (W.D. Pa.) for details.

¹³In Section 6.3, I allow landlords to consider non-payment of rent in an alternative unit of measure: back rent relative to the tenant’s tenure in the rental property.

of disparate impact also includes differences arising from non-race characteristics that do not contribute to achieving some legitimate interest. As such, the conditional racial disparity in eviction filing thresholds that I estimate in this paper aligns with unlawful disparate impact if there exists no legitimate interest for a landlord to selectively file non-payment based eviction cases against tenants equally behind on rent. A restated version of this assumption—that tenants equally behind on rent are equally qualified for an eviction filing, absent other lease violations—also aligns the conditional disparity in thresholds with the concept of total discrimination in [Bohren et al. \(2022\)](#)—racial disparities in treatment among equally qualified individuals.

This assumption may not hold if landlords selectively evict tenants that they predict to be less likely to pay going forward than tenants with higher predicted payment ability. To the extent that observable (non-race) predictors of future payment ability are correlated with race, this would lead to landlords evicting minority tenants over less back rent than white tenants. One could argue that this behavior does not meet the legal standard to be considered unlawful discrimination if a landlord’s prediction about future rental earnings coming solely from observed non-race characteristics of tenants is considered a legitimate or qualifying reason to tolerate less back rent from minority tenants.¹⁴ However, this behavior would be unlawful if tolerating less back rent from minority tenants was even in part based on tenant race or if landlords use tenant race as a signal for payment ability. In practice, I find little evidence that racial disparities in eviction thresholds can be explained by differences in proxies for future payment ability in Section 6.2, supporting the interpretation that my estimates capture unlawful discrimination.

2.2 Eviction Filing Process in Philadelphia

Eviction filings are the first publicly documented step in the eviction process, although tenants may receive a short period of notice beforehand.¹⁵ When filing an eviction case in Philadelphia, landlords indicate all causes of the filing, selecting from the following choices: non-payment of rent, breach of a condition of the lease, and/or termination of the lease term.¹⁶ Note that they may indicate multiple causes. In the filing, the landlord also reports

¹⁴It is important to note that landlords operating many rental units—which are the landlords for which I can identify racial disparities in filing thresholds—may be unlikely to observe such non-race characteristics that predict future payment, beyond the information collected when initially screening the tenant.

¹⁵Philadelphia landlords are mandated to provide tenants at least ten days notice prior to filing eviction cases over non-payment, but leases often waive this notice requirement in practice. [Hoffman and Strezhnev \(2022\)](#) find that 67% of unsubsidized lease agreements associated with Philadelphia eviction cases waive or reduce the right to notice.

¹⁶During the study period (2006-2019), Philadelphia landlords could terminate a lease for any reason except for in retaliation for a tenant reporting a complaint, joining a tenant organization, or being a victim

the amount of back rent, fees, and physical damages owed by the tenant, information on any other lease violations, and details of the lease agreement. After submitting this information, the landlord and tenant are notified of the scheduled hearing date.

2.3 Eviction Court Process in Philadelphia

After a landlord files an eviction case, the eviction case could proceed in a number of ways. A landlord may withdraw the case prior to the scheduled hearing date, typically in response to payment of the past due rent, the tenant voluntarily moving out, or a private settlement agreement. Barring a withdrawal, if the tenant does not appear in housing court for the scheduled hearing, the case results in a default judgment in favor of the landlord. In the relatively rare occurrence that the landlord does not appear in court for the scheduled hearing, the case results in a default judgment in favor of the tenant.

If both the landlord and tenant appear in court, policy in Philadelphia requires the parties to attempt to reach an agreement before seeing a judge. This negotiation occurs between the landlord and tenant and any attorneys present.¹⁷ If the parties reach an agreement, a staff member from the Municipal Court Mediation Unit reviews the agreement and confirms that both parties understand the agreement. Both parties then sign the agreement, and it is entered into the court record as a judgment by agreement (JBA) without any involvement of a judge ([Reinvestment Fund 2020](#)).¹⁸ If the parties cannot reach an agreement, they contest the case before a judge who decides the case. Conditional on both parties appearing in court, 87% of cases result in a JBA, meaning that judges preside over much fewer eviction cases in Philadelphia (less than 5% of all filings) compared to other jurisdictions.¹⁹

Any judgment that is entered via default, a JBA, or decided by a judge may include a judgment for possession of the property and/or a judgment for money. After any judgment for possession in favor of the landlord is entered, the landlord may file for a writ of possession which will then be served to the tenant before the Landlord and Tenant Office or the Sheriff's Office executes the eviction.

of domestic violence. However, most filings (94%) involved non-payment of rent. A 2019 Philadelphia law introduced the requirement of "good cause" for filing an eviction case with allowable causes restricted to non-payment of rent, breach of lease, property damage, and other violations. For details, see Unfair Rental Practices, Philadelphia Code § 9-804.

¹⁷Over 80% of landlords have legal representation, but less than 8% of tenants are represented. If neither party has an attorney, a volunteer mediator is provided.

¹⁸The Philadelphia JBA system came about as a way to process high volumes of cases by allowing cases to proceed without any judge involvement ([Reinvestment Fund 2020](#)).

¹⁹Unlike the JBA process in Philadelphia, judges in Cook County, Illinois and New York City approve settlement agreements, thus influencing a larger share of eviction filings ([Collinson et al. 2024](#)).

3 Data

3.1 Eviction Court Records

I analyze the universe of residential eviction cases filed in Philadelphia between 2006 and 2019. These court records contain information on the filing date, defendant (tenant) names, plaintiff names and addresses, property address, monthly contract rent, cause(s) of the case filing, itemized claim amount (including back rent, late fees, physical damages, etc.), judgments, and whether the landlord filed for a writ of possession.

I identify landlords using plaintiff names and plaintiff addresses listed on eviction filing records. The plaintiff information may correspond to property owners in some cases and property managers in others, depending on which party is responsible for filing the eviction case. I define a landlord as the relevant decision-maker responsible for filing eviction cases, which should be reflected in the plaintiff information. I construct landlord identifiers by linking eviction filings that share plaintiff names or plaintiff addresses. This linkage approach is meant to account for inconsistent spellings and abbreviations, and the fact that many landlords owning or managing multiple buildings list property-specific companies as plaintiff names on eviction filings. See Appendix B.1 for details on the construction of landlord identifiers.

The itemized amounts claimed are a unique feature of these data that are rarely observable in eviction case records from other jurisdictions. This is essential for testing for discrimination as the amounts provide a continuous measure of a landlord’s grievances against a tenant at the time of an eviction filing. Only certain lease violations related to, for example, noise complaints or unapproved guests may not be reflected in the claimed amounts. For this reason, I focus on cases caused only by non-payment of rent for which the amount of back rent claimed should be a comprehensive measure of a tenant’s qualification for eviction.

3.2 Measuring Tenant Race

I impute the race of tenants using their names and property addresses listed on the eviction filing. To do so, I utilize the fully Bayesian Improved Surname Geocoding (fBISG) algorithm developed by Imai et al. (2022). The fBISG approach estimates the conditional posterior probability that tenant i belongs to racial group r given the i ’s first name, last name, and census block. By accounting for measurement error in observed census counts, fBISG has been shown to substantially improve imputation accuracy compared to standard BISG imputations, especially for minority groups (Imai et al. 2022). Figure A1 validates that fBISG performs well (out-of-sample) in Philadelphia over the same time period by

comparing fBISG predictions of borrower race among Philadelphia mortgage originations to true borrower race reported in Home Mortgage Disclosure Act (HMDA) data.²⁰ See Appendix B.2 for more details on the fBISG methodology.

I conduct all analysis at the case filing level and compare landlords' treatment of white and minority tenants, where a minority tenant is defined as Black, Hispanic, Asian, or other. Since an eviction filing may name multiple tenants who live in the same rental unit, I classify filings based on the filing-level probability any tenant listed belongs to a racial minority group. I classify a filing as white if $Pr(\text{any } i \text{ is minority}) \leq 20\%$, and I classify a filing as minority if $Pr(\text{any } i \text{ is a minority}) \geq 80\%$.²¹ Because the underlying distribution of filing-level minority probabilities is bimodal with many cases concentrated close to 0 or 1, the analysis sample only excludes 14.62% of filings that cannot be classified as white or minority.²² Furthermore, the accuracy rate implied by aggregating the fBISG probabilities of correct classification within this sample is 97.25%. See Table A1 for a comparison of the percentage of cases dropped and the implied accuracy rates across alternative race classification rules.

Using these imputation-based tenant race classifications, Table 1 reports the racial composition of Philadelphia eviction filings. For comparison, I also report the racial composition of the population of renters in Philadelphia from the 2010 decennial census. The vast majority of eviction cases are filed against non-white tenants. Black tenants appear especially over-represented while white tenants appear especially under-represented in eviction filings.

3.3 Descriptive Statistics

Philadelphia eviction records offer unique insight into the circumstances under which landlords file eviction cases, revealing that many tenants face eviction cases over strikingly little unpaid rent. Among cases filed only due to non-payment of rent, tenants owe less than one month of contract rent in 10.6% of cases (1,433 cases per year) and owe less than \$600 in 17.7% of cases (2,389 cases per year). Moreover, minority renters appear particularly vulnerable to eviction cases. Not only are they more likely to face an eviction filing (Table 1), but the average claim amount tenants owe at the time of the filing is over \$600 lower for

²⁰I predict borrower race using borrower names and census blocks of Philadelphia deeds reported in data from CoreLogic that could be linked to HMDA data. Access to CoreLogic data was provided by the Becker-DeMunnich Institute for Economics at the University of Chicago and the Fama-Miller Center for Research in Finance.

²¹I refer to this cutoff approach as the 80% rule since it imposes that white filing classifications are at least 80% certain to have all white tenants and minority filing classifications are at least 80% certain to have at least one minority tenant.

²²In Section 6.3.2, I demonstrate that the results are robust to including these cases with less certain race predictions to estimate alternative imputation-based estimators.

Table 1: Racial Composition of Eviction Filings

	Eviction Filings	Renter Population
Asian	1.40%	5.46%
Black	76.97%	44.36%
Hispanic	7.69%	10.92%
White	9.06%	36.83%
Other/Mixed	4.88%	2.43%
N	261,738	

This table reports the racial composition of eviction filings and renters in Philadelphia. The eviction filings sample includes filings that could be classified as white or minority with at least 80% certainty. A minority filing is further classified as detailed race r if the most likely racial group of each individual tenant named in the filing is r . A minority filing is classified as mixed if the most likely racial group of each individual tenant named in the filing differ. The racial composition of renters in Philadelphia is sourced from the 2010 decennial census.

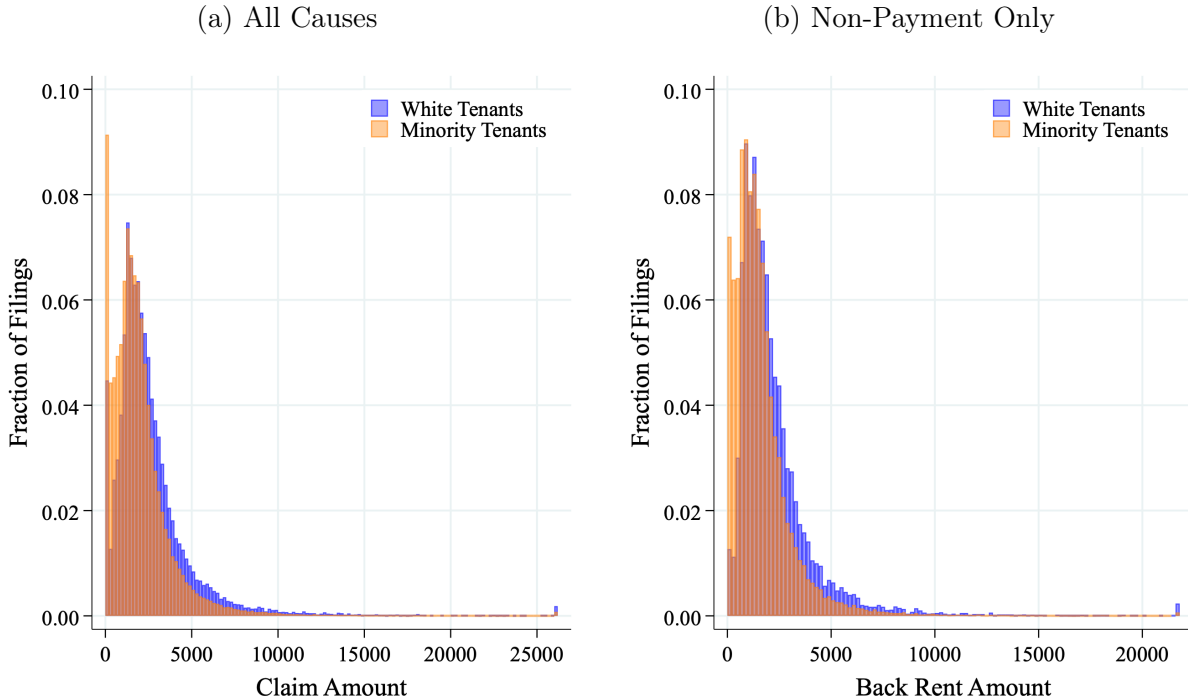
minority tenants than white tenants (Table A2).²³ Considering only the portions of the claim for which tenants are directly responsible—back rent, physical damages, and late fees,—it remains that minority tenants owe less on average than white tenants (Table A2).

Figure 1a plots the full distribution of claim amounts by tenant race and illustrates more clearly that minority tenants owe less than white tenants when landlords file eviction cases. Among cases caused only by non-payment of rent, the minority distribution of back rent owed at filing remains visibly lower than the white distribution of back rent (Figure 1b). Since minority tenants with eviction filings also tend to have lower monthly contract rents than white tenants (Table A2), minority tenants could theoretically owe equal months of back rent at the time of the filing as white tenants. However, this does not appear to be the case in Figure A2a, which shows that minority tenants also appear to owe fewer months of back rent at the time of filing. This pattern is even clearer among tenants with below median contract rents (Figure A2b).

The outcomes of eviction filings also appear to be related to tenant race. Figure 2 plots the relationship between the amount of back rent owed at filing and whether the landlord won repossession of the property—measured as filing for a writ of possession or the tenant agreeing to move out during the JBA process—separately for white and minority cases. It is first important to note that back rent appears positively correlated with the likelihood

²³The claim amount captures the total amount the landlord claims the tenant owed at the time the landlord filed the eviction case. The amount claimed on eviction filings tends to be mostly composed of back rent, but landlords also claim physical damages, late fees, attorney fees, and other fees. Total claim amounts and back rent values are separately winsorized at the 99.9th percentile.

Figure 1: Distribution of Amounts Owed at Filing



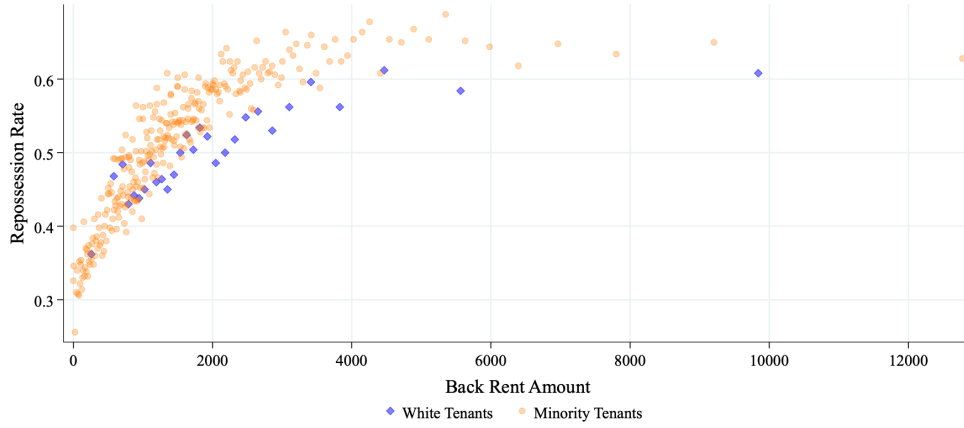
Subfigure (a) plots the claim amount distribution separately for Philadelphia eviction filings against white and minority tenants. Subfigure (b) plots the back rent distribution of filings separately for filings against white and minority tenants that were caused only by non-payment of rent. Both subfigures exclude filings that could not be classified as white or minority with at least 80% certainty based on tenant race imputations. Claim amounts and back rent values are winsorized at the 99.9th percentile.

of the landlord winning repossession. This is consistent with landlords more easily winning repossession from tenants with greater lease violations. The second key takeaway is that landlords tend to win repossession from minority tenants at a higher rate than in cases against white tenants owing similar back rent. This pattern is visible across the entire distribution of back rent and suggests that it requires less back rent for landlords to win repossession from minority tenants compared to white tenants.

4 Conceptual Framework

To understand how racial disparities in the amount owed at filing could be explained by racial discrimination, I provide a framework to conceptualize how discrimination may enter a landlord’s eviction filing decision. In this section, I first provide a model of a landlord’s threshold rule for filing eviction cases. I then use this framework to demonstrate how racial disparities in thresholds can arise from three potential sources—differences in non-race

Figure 2: Racial Differences in Repossession



This figure plots a bin scatter of the relationship between back rent and landlords winning repossession. Repossession is defined as a landlord filing for a writ of possession or the tenant agreeing to vacate the unit during the JBA process. Each point reflects 500 eviction cases against tenants belonging to a given racial group. The sample includes Philadelphia eviction filings caused only by non-payment of rent and excludes filings that could not be classified as white or minority with at least 80% certainty based on tenant race imputations. The x-axis is the mean back rent among cases in the cell. The y-axis is the rate of landlords winning repossession among cases in the cell.

characteristics, accurate statistical discrimination, or racial bias—and discuss the testable implications of each source.

4.1 Landlord Filing Decision

Consider all tenants of a single landlord. Let R_i denote tenant race and let Z_i capture all characteristics of tenant i observable to the landlord besides race. Z_i may include individual characteristics such as income and also characteristics of the rental property including neighborhood factors. Each month, tenant i agrees to pay the landlord the contract rent amount Y_i^0 . If the tenant falls behind on rent, they accrue some amount of back rent, $X_i(t)$, owed to the landlord. For simplicity, assume that the tenant commits no other lease violations besides non-payment of rent. At any time t such that $X_i(t) > 0$, the landlord can legally file an eviction case over non-payment. Let $D_i(t) = 1$ denote this eviction filing action taken by the landlord. Filing an eviction case is costly, however. The landlord pays a fixed cost c to file an eviction case, which reflects any unrecoverable court filing costs and legal fees.²⁴

²⁴I allow c , as well as the \tilde{p} and \tilde{v} functions defined below, to vary across landlords when I bring the model to data. Here, however, I simplify the model's notation by considering the behavior of a single landlord.

If the landlord chooses to file an eviction case against i , the case could either result in the landlord winning repossession of the property to rent to a new tenant ($Y_i^* = 1$) or the current tenant remaining in the property ($Y_i^* = 0$). The true expected probability that the landlord wins repossession of the property is:

$$\tilde{p}(x, y, z, r) = \mathbb{E}[Y_i^* = 1 | X_i(t) = x, Y_i^0 = y, Z_i = z, R_i = r] \quad (1)$$

I assume $\tilde{p}(x, y, z, r)$ is weakly increasing in back rent, x , which is consistent with the positive relationship between repossession rates and back rent in Figure 2. The likelihood of winning repossession also depends on Y_i^0 since the ongoing contract rent amount informs the severity of the amount of back rent. Z_i enters \tilde{p} given that tenants with lower incomes or greater childcare responsibilities, for example, may have different likelihoods of appearing in eviction court as scheduled, affecting the likelihood of the landlord winning repossession. R_i enters \tilde{p} to allow for the possibility that minority tenants face unequal access to court resources or the eviction court process systematically favors white tenants, as suggested by Figure 2.

The landlord may have biased perceptions about the probability of winning repossession of the property. I define the landlord's perceived probability of winning repossession as:

$$p(x, y, z, r) = \tilde{p}(x, y, z, r) + b^p(r) \quad (2)$$

The bias term $b^p(r)$ allows for landlords to have racially biased beliefs about the likelihood of winning repossession from a given tenant (Bohren et al. 2025; Bordalo et al. 2016).

The landlord's true payoff from winning repossession of the property from tenant i is $\tilde{v}(y, z)$ for $Y_i^0 = y$ and $Z_i = z$. Intuitively, \tilde{v} captures the difference in expected future rental income from replacing tenant i with some unknown new tenant. Non-race characteristics, Z_i , enter \tilde{v} since i 's factors such as income may predict the likelihood that i pays rent in the future, affecting the landlord's payoff from replacing them. Contract rent, Y_i^0 , enters \tilde{v} since it captures the rental value of the property that the landlord could earn from a paying tenant. I assume that race cannot impact the true payoff \tilde{v} , holding fixed Z_i and Y_i^0 .

I model the landlord's perceived payoff from winning repossession as:

$$v(y, z, r) = \tilde{v}(y, z) + b^v(r) \quad (3)$$

The perceived payoff v additionally varies with tenant race r to allow for the possibility that the landlord exhibits racial bias, $b^v(r)$. The bias term, $b^v(r)$, may reflect taste-based discrimination (Becker 1957) if the landlord values removing minority tenants from the property differently from removing white tenants, conditional on \tilde{v} . Alternatively, $b^v(r)$ may reflect

racially biased beliefs or stereotypes (Bohren et al. 2025; Bordalo et al. 2016) about the likelihood of the current tenant paying rent in the future.

The landlord files an eviction case against tenant i if the perceived payoff weighted by the perceived probability of winning repossession exceeds the fixed cost of filing:²⁵

$$D_i(t) = \mathbb{1}\{p(X_i(t), Y_i^0, Z_i, R_i)v(Y_i^0, Z_i, R_i) \geq c\} \quad (4)$$

By monotonicity of $p(x, y, z, r)$ over x , this is equivalent to a threshold rule for filing an eviction case:

$$D_i(t) = \mathbb{1}\{X_i(t) \geq \tau_i\} \quad (5)$$

where

$$\tau_i = p^{-1}\left(\frac{c}{v(Y_i^0, Z_i, R_i)}, Y_i^0, Z_i, R_i\right) = \tilde{p}^{-1}\left(\frac{c}{\tilde{v}(Y_i^0, Z_i) + b^v(R_i)} - b^p(R_i), Y_i^0, Z_i, R_i\right) \quad (6)$$

From this threshold rule, it is apparent that the landlord's threshold τ_i varies with the tenant's non-race characteristics, contract rent, and race.²⁶

4.2 Discrimination Definition and Sources

As discussed in Section 2.1, I define a landlord as exhibiting total racial discrimination if they use different filing thresholds, τ_i , by tenant race conditional on contract rent:

$$\mathbb{E}[\tau_i | Y_i^0 = y, R_i = \textit{minority}] \neq \mathbb{E}[\tau_i | Y_i^0 = y, R_i = \textit{white}] \quad (7)$$

Total discrimination against minority (white) tenants is the specific case of lower thresholds for minority (white) tenants. Given that τ_i may vary by a tenant's race and non-race characteristics (Equation 6), total racial discrimination can arise from three potential sources: differences in non-race characteristics, accurate statistical discrimination, or racial bias.

First, components of Z_i that are correlated with race can generate racial differences in τ_i even if neither p nor v depend on R_i . If this alone explains total racial discrimination, then thresholds should be equal across race holding fixed Z_i :

$$\mathbb{E}[\tau_i | Y_i^0 = y, Z_i = z, R_i = \textit{minority}] = \mathbb{E}[\tau_i | Y_i^0 = y, Z_i = z, R_i = \textit{white}] \quad (8)$$

²⁵This follows from a landlord maximizing utility $U_0 + D_i(t)(p(X_i(t), Y_i^0, Z_i, R_i)v(Y_i^0, Z_i, R_i) - c)$ for some U_0 .

²⁶I model τ_i as static here but allow for time-varying thresholds in Section 6.3.5. Thresholds may vary over time due to macroeconomic conditions changing the cost c of filing an eviction case.

If instead the landlord exhibits direct discrimination on tenant race (accurate statistical discrimination or racial bias), then thresholds would vary by tenant race even after holding fixed Z_i .

Second, a landlord exhibits accurate statistical discrimination if $b^p(\textit{minority}) = b^p(\textit{white})$ and $b^v(\textit{minority}) = b^v(\textit{white})$. Accurate statistical discrimination can generate racial disparities in τ_i due to the dependence of \tilde{p} on R_i . In essence, a landlord’s accurate predictions about any racial differences in the likelihood of winning repossession in court results in their filing threshold varying by tenant race. To test this source, I can compare the marginal outcomes of filings across tenant race:

Proposition 1. If a landlord exhibits accurate statistical discrimination, then $\mathbb{E}[Y_i^*|X_i(t) = \tau_i, Y_i^0 = y, Z_i = z, R_i = \textit{minority}] = \mathbb{E}[Y_i^*|X_i(t) = \tau_i, Y_i^0 = y, Z_i = z, R_i = \textit{white}]$.

The intuition of Proposition 1 is that a landlord accurately accounting for any racial differences in the likelihood of winning repossession in court should achieve equal rates of winning repossession across tenant race at the margin, holding fixed non-race characteristics and contract rent. For the proof of Proposition 1, see Appendix C.1.

Third, racial bias arises from the dependence of τ_i on $b^p(R_i)$ and $b^v(R_i)$. A landlord exhibits racial bias if $b^p(\textit{minority}) \neq b^p(\textit{white})$ or $b^v(\textit{minority}) \neq b^v(\textit{white})$. This may reflect taste-based discrimination based on racially biased preferences (Becker 1957) or inaccurate statistical discrimination based on racially biased beliefs (Bohren et al. 2025). Holding fixed Z_i , racially biased beliefs or preferences about the probability of repossession or the payoff from repossession can generate racial differences in filing thresholds even if \tilde{p} does not actually depend on R_i . A testable implication of racial bias is that landlords deviate from equal rates of winning repossession across race at the margin, holding fixed non-race characteristics and contract rent:

Proposition 2. If $\mathbb{E}[Y_i^*|X_i(t) = \tau_i, Y_i^0 = y, Z_i = z, R_i = \textit{minority}] \neq \mathbb{E}[Y_i^*|X_i(t) = \tau_i, Y_i^0 = y, Z_i = z, R_i = \textit{white}]$, then a landlord exhibits racial bias.

The intuition of Proposition 2 is that a landlord tolerating a lower expected repossession rate from one racial group must reflect racially biased preferences or beliefs. Racial bias against minority (white) tenants is the specific case where the expected repossession rate is lower for minority (white) tenants. For the proof of Proposition 2, see Appendix C.2.

Propositions 1 and 2 constitute a marginal outcome test for racial bias in eviction filing which is consistent with Canay et al. (2024) and follow the same intuition of outcome-based tests for racial bias conducted in other settings (Knowles et al. 2001; Persico and Todd 2006; Alesina and La Ferrara 2014; Anwar and Fang 2015; Arnold et al. 2018; Marx 2021; Dobbie

et al. 2021; Baron et al. 2024).²⁷ I note that the null hypothesis of the marginal outcome test is that landlords exhibit accurate statistical discrimination with respect to court outcomes of eviction filings. If I relax the assumption discussed above that \tilde{v} —the landlord’s payoff from winning repossession of the property from the current tenant—does not depend on R_i , then it would also be possible that landlords exhibit accurate statistical discrimination with respect to future rental earnings (e.g. filing minority cases over less back rent because they are less likely to catch up on rent in the future). In this case, rejecting accurate statistical discrimination with respect to court outcomes would imply some combination of racial bias and accurate statistical discrimination with respect to future rental earnings. By conditioning the marginal outcome test on non-race characteristics that predict future payment ability (Z_i) and assuming race does not predict future payment ability conditional on Z_i , the test is meant to compare the court outcomes of marginal white and marginal minority filings that are likely to generate the same future rental earnings, therefore isolating racial bias from accurate statistical discrimination with respect to future payments.

5 Empirical Approach

I next describe my empirical test for average racial disparities in eviction filing thresholds which I interpret as total racial discrimination. I return to the question of the sources of any disparity in Section 7. The total racial discrimination parameter of interest for a single landlord is:

$$\mathbb{E}[\tau_i|Y_i^0 = y, R_i = \textit{minority}] - \mathbb{E}[\tau_i|Y_i^0 = y, R_i = \textit{white}] \quad (9)$$

To estimate Equation 9, I leverage the notion that tenants gradually accrue back rent until back rent reaches the landlord’s eviction filing threshold. I observe the amount of back rent tenants owe upon crossing this threshold, which is the time the landlord files an eviction case. I denote this observable value $X_i(t^*)$, where $t^* = \min(t)$ such that $D_i(t) = 1$. My empirical approach utilizes observable racial differences in $X_i(t^*)$ to estimate racial differences in τ_i , conditional on landlord and contract rent.

To build intuition for this approach, suppose that back rent evolves continuously over time. Continuous $X_i(t)$ implies that back rent owed at the time of an eviction filing is exactly the landlord’s threshold: $X_i(t^*) = \tau_i$. Therefore, Equation 9 is equivalent to:

$$\mathbb{E}[X_i(t^*)|Y_i^0 = y, R_i = \textit{minority}] - \mathbb{E}[X_i(t^*)|Y_i^0 = y, R_i = \textit{white}] \quad (10)$$

²⁷Propositions 1 and 2 are a stricter version of the marginal outcome test in Hull (2021) which allows bias to arise via non-race characteristics.

In reality, back rent may evolve discretely over time $t \in \{0, 1, 2, 3, \dots\}$. Discrete changes in back rent, however, are bounded due to the monthly payment schedule of rent. Tenant i cannot owe more than one additional month of contract rent than in the previous month:

$$X_i(t) \leq X_i(t-1) + Y_i^0 \quad (11)$$

In practice, distressed renters often make partial rent payments (Manville et al. 2023), suggesting that discrete changes in back rent may be smaller than one month of contract rent in many cases. Discrete changes in back rent may place back rent owed at time t^* above the landlord’s threshold, but the distance above is bounded by the value of contract rent:

$$X_i(t^*) \in \left[\tau_i, \tau_i + Y_i^0 \right) \quad (12)$$

I assume that back rent owed at filing exceeds the threshold by the same amount on average for white and minority tenants who have the same landlord and contract rent. Formally, this assumption is as follows for tenants of a single landlord with contract rent y :

$$\mathbb{E}[X_i(t^*) - \tau_i | Y_i^0 = y, R_i = \textit{minority}] = \mathbb{E}[X_i(t^*) - \tau_i | Y_i^0 = y, R_i = \textit{white}] \quad (13)$$

This imposes that back rent owed at filing by white and minority tenants is the same amount above their respective τ_i on average. From this assumption, it follows that any racial differences in back rent equate to racial differences in thresholds conditional on landlord and contract rent, meaning that estimating Equation 9 in the discrete case is identical to the continuous case (Equation 10). The assumption would only be violated if the size of the last discrete change in back rent before the filing is statistically different by race among tenants with the same landlord and contract rent. See Appendix D for a discussion of bias if this assumption is violated.

To implement this test for total racial discrimination, I estimate the following specification for eviction filings caused only by non-payment of rent against tenant(s) i by landlord j on filing date t :

$$\textit{Ln}(\textit{Back Rent})_{ijt} = \beta \textit{Minority}_i + \alpha_{j,y(ijt)} + \varepsilon_{ijt} \quad (14)$$

Back Rent is the dollars of back rent owed at the time of the eviction filing. *Minority* is an indicator equal to one if the filing is at least 80% likely to be against a minority tenant based on imputed tenant race. I include landlord-by-contract rent fixed effects, $\alpha_{j,y(ijt)}$, to condition on landlord j and contract rent y . I cluster standard errors at the landlord level.

For landlords that observe tenant race, β is an unbiased estimate of the average racial disparity in eviction filing thresholds under the assumption that race is conditionally indepen-

dent of the amount of back rent owed at filing given the imputation inputs—an individual’s first name, last name, and census block. Landlords that own or manage a large number of properties may not directly observe or recall a tenant’s race when deciding to file an eviction case. In this case, imputation-based race classifications should accurately capture a landlord’s belief about the tenant’s race based on their name and location.

The importance of the independence assumption depends on the share of landlords observing tenant race in the analysis sample of eviction filings. Since β is only identified for all j, y such that $Filings_{j,y,white} > 0$ and $Filings_{j,y,minority} > 0$, β reflects average racial disparities among this selected sample of likely large landlords who may be less likely to directly observe tenant race. I discuss remaining concerns about imputed tenant race and potential violations of the independence assumption in Section 6.3.

6 Discrimination Results

6.1 Total Racial Discrimination

In Table 2, I present the results of estimating Equation 14. Column 1 reports the raw racial disparity in back rent owed at the time of eviction filings, and column 2 estimates the conditional disparity that I interpret as total racial discrimination. From column 2, minority tenants owe 4.9% less back rent at the time of a filing than white tenants with the same landlord and contract rent. This is evidence that landlords discriminate by using 4.9% lower filing thresholds for minority tenants.

Comparing columns 1 and 2, the conditional disparity explains about 11% of the raw disparity in back rent owed at filing. The difference between these coefficients reflects tenants sorting to landlords and contract rents differently by race. Some of this sorting may reflect earlier discrimination that occurs when landlords consider rental applications (Hanson and Hawley 2011; Ewens et al. 2014; Christensen et al. 2021). This means that the coefficient reported in column 2 isolates total discrimination in eviction filing from any earlier housing market discrimination, providing evidence that discrimination persists beyond entry into rental housing.

The magnitude of total discrimination is economically meaningful. Landlords using 4.9% lower filing thresholds for minority tenants equates to tolerating \$131 fewer dollars of back rent (column 3) or 0.14 fewer months of back rent (column 4) from minority tenants. In other settings, even small amounts of credit can reduce eviction filings (Lodermeier 2024), suggesting that informal credit from landlords raising minority thresholds by 4.9% may allow a substantial share of minority tenants to avoid a filing entirely.

Table 2: Total Racial Discrimination

	Back Rent			
	Log Dollars		Dollars	Months
	(1)	(2)	(3)	(4)
Minority Tenant	-0.4493** (0.2273)	-0.0489*** (0.0168)	-130.5750*** (38.6363)	-0.1389*** (0.0422)
Observations	159,799	120,328	120,328	119,766
R ²	0.0176	0.6447	0.4799	0.6000
Mean White Outcome	7.4644	7.3683	1,991.68	2.2575
Landlord \times Contract Rent FE		✓	✓	✓

*p < 0.1, **p < 0.05, ***p < 0.01. This table reports the results from estimating Equation 14. The sample includes Philadelphia eviction filings caused only by non-payment of rent that could be classified as white or minority with at least 80% certainty based on tenant race imputations. Columns 2 through 4 exclude singleton landlord-by-contract rent observations. Contract rent fixed effects refer to the exact dollar value of contract rent. The dependent variable in column 4 is the dollar value of back rent divided by contract rent. Standard errors are clustered at the landlord level.

6.2 Non-Race Characteristics

Assuming that back rent fully captures qualification an eviction filing, conditional on landlord and contract rent, then the estimates reported in Table 2 reflect total racial discrimination. These estimates reflect some combination of differences arising from indirect or systemic factors—differences in non-race characteristics across racial groups—and direct discrimination on the basis of race alone. To understand the relative importance of indirect and direct discrimination, I estimate these racial disparities with additional controls for tenant covariates. If differences in non-race characteristics fully drive total racial discrimination, then filing thresholds should be equal across race holding fixed non-race characteristics Z_i among tenants with the same landlord and contract rent (Equation 8). Rejecting equal thresholds conditional on Z_i is evidence of direct discrimination on the basis of tenant race.

I do not observe the full vector of non-race tenant characteristics Z_i that may be observable to the landlord but test for direct discrimination by controlling for the subset of Z_i in the data. These characteristics include an imputed measure of gender, the number of tenants named in the filing, the number of previous cases filed against the tenant at the same address, two proxy measures of rent burden, and census tract fixed effects. I construct rent burden as individual-level contract rent listed in the filing as a share of median household income among similar households in the American Community Survey (ACS). The first measure of rent burden uses median household income among renters in the census tract and year of the filing. The second measure uses median household income in the census tract and year

Table 3: Controls for Non-Race Characteristics

	Log Back Rent			
	(1)	(2)	(3)	(4)
Minority Tenant	-0.0489*** (0.0168)	-0.0540*** (0.0155)	-0.0628*** (0.0159)	-0.0474*** (0.0162)
Observations	120,328	120,328	120,328	120,328
R ²	0.6447	0.6460	0.6465	0.6576
Mean White Back Rent	1,584.90	1,584.90	1,584.90	1,584.90
Landlord × Contract Rent FE	✓	✓	✓	✓
Tenant Controls		✓	✓	✓
Rent Burden FE (Tract × Renter Inc.)		✓	✓	✓
Rent Burden FE (Tract × Race Inc.)			✓	✓
Census Tract FE				✓
Bias-adj coef.				-0.0433

*p < 0.1, **p < 0.05, ***p < 0.01. This table reports the results from estimating Equation 14 with controls. The sample includes Philadelphia eviction filings caused only by non-payment of rent that could be classified as white or minority with at least 80% certainty based on tenant race imputations. Column 1 restates the estimate of total discrimination from Table 2. Columns 2 through 4 introduce control variables. Tenant controls include the imputed gender of first tenant named, fixed effects for the number of tenants listed on the filing, and fixed effects for the number of previous cases filed against the tenant at the same property. Rent burden fixed effects are 2-percentage point bins of the ratio of individual-level contract rent to median household income among similar individuals in the ACS. The first set of rent burden fixed effects uses median household income among renters in the census tract around the year of the filing. The second set of rent burden fixed effects uses median household income by household head race and census tract around the year of the filing. Household head race in the eviction filing is defined based on the imputed race of the first tenant listed if the imputation is at least 80% likely. Median household income measures are sourced from 5-year ACS estimates, where the midpoint of the 5-year ACS period is matched to the case filing year. All control variables include a category for missing values. All specifications exclude singleton observations for the fixed effects used. Contract rent fixed effects refer to the exact dollar value of contract rent. Standard errors are clustered at the landlord level. The bias-adjusted coefficient is estimated following Oster (2019) and sets $R_{max}^2 = R_{Col.4}^2 + (R_{Col.4}^2 - R_{Col.1}^2)$.

of the filing among householders with the same detailed race as the imputed race of the first tenant listed on the filing.²⁸ For this test to be informative, these control variables must vary with tenant race after conditioning on landlord and monthly contract rent in order for there to be scope for these characteristics to explain total racial discrimination. Table A3 confirms tenant race is a highly significant predictor of each of the observable non-race characteristics of tenants, conditional on landlord and contract rent.

Table 3 provides evidence that differences in non-race characteristics do not explain the racial disparity in eviction filing thresholds. Column 1 re-states the previous estimate of

²⁸I match median household income by tract, year, and race if the most likely race—Asian, Black, Hispanic, white, or other—of the first tenant listed on the filing is imputed with at least 80% certainty.

total racial discrimination for reference, and columns 2 through 4 estimate discrimination coefficients controlled for gender, the number of tenants, previous filings, rent burden, and neighborhood characteristics. Controlling for these characteristics does not substantially attenuate the racial disparity and if anything, the disparity widens in columns 2 and 3. This is consistent with total racial discrimination entirely reflecting direct discrimination on the basis of tenant race.

To the extent that landlords selectively file eviction cases according to other tenant characteristics not captured here, the estimates in Table 3 could include omitted variable bias. In particular, I do not observe individual-level income which landlords may observe and incorporate into filing decisions if income predicts a tenant’s ability to pay the contract rent going forward and thus the landlord’s payoff from eviction (\tilde{v}). It is important to note, however, that landlords likely observe at best a lagged measure of tenant income from when they initially screened the tenant for the unit.

This potential omitted variable does not appear to be a substantial issue for a number of reasons. First, the two proxy measures of rent burden and the inclusion of census tract fixed effects should account for a considerable portion of any racial disparities in future rent payment ability that landlords observe. In particular, the second rent burden proxy used in columns 3 and 4 directly incorporates racial disparities in income within census tract and year. If unobserved variation in tenant income was biasing the direct discrimination estimate in column 4, I would expect the inclusion of these controls to attenuate the discrimination coefficient relative to column 1. Instead, the magnitude increases in columns 2 and 3 and is barely attenuated in column 4. Second, I follow Oster (2019) to estimate bias-adjusted coefficients. Assuming that controlling for any remaining unobservables would add as much to the R^2 as the observables, I estimate a bias-adjusted coefficient of -0.0433 which falls within the 95% confidence interval of both the total discrimination estimate in column 1 and the controlled estimate in column 4.

Overall, the results in Table 3 suggest limited scope for differences in non-race characteristics to explain the racial disparity in eviction thresholds. This suggests racial discrimination in eviction filing mostly reflects landlords discriminating directly on the basis of tenant race rather than correlates of race.

6.3 Robustness Checks

I next address potential concerns about interpreting these estimates as evidence of racial discrimination. I report the results of variation robustness checks that coarsen the landlord-by-contract rent fixed effects, address concerns about race measurement error, use narrower

definitions of landlords, consider the assumption on discrete accrual of back rent, allow for time-varying eviction filing thresholds due to macroeconomic conditions, exclude cases with potential omitted lease violations, use alternative measures of non-payment of rent, and address concerns about selection from different responses to eviction notices.

6.3.1 Coarser Fixed Effects

Measuring total racial discrimination requires landlords to have filed multiple eviction cases against tenants with the same contract rent that belong to different racial groups. The effective sample that identifies total discrimination is 35% ($N = 56,526$) of cases in the analysis sample that are filed by 486 unique landlords.²⁹ To broaden the set of landlords reflected in the estimates, I use coarser fixed effects in Table A4. Using additive landlord and contract rent fixed effects yields a nearly identical estimate, as does grouping contract rent into \$100, \$50, and \$25 bins. The samples that identify discrimination with these coarser fixed effects reflect many more landlords (up to 1,274). The stability of coefficients in these larger and more representative samples is consistent with my main estimate reflecting the average amount of eviction filing discrimination in Philadelphia.

6.3.2 Race Imputation

Estimates of total racial discrimination using imputation-based race classifications are unbiased if race is conditionally independent of the amount of back rent owed at filing given an individual’s name and census block. If this assumption is violated, the estimated racial disparities in back rent using cutoff-based race classifications may be biased in either direction (Chen et al. 2019). For example, I may understate disparities (in absolute value) if conditional on the imputed probability a tenant is a racial minority, landlords tolerate less back rent from true minority renters than true white renters. I may overstate disparities if racial disparities in back rent among the unclassified cases that fall below the 80% cutoff and are excluded from the analysis are smaller than disparities among the classified cases.

I address concerns about these potential biases in Table A5. I verify in panel A that the choice of cutoff (80%) for classifying filings by race is not driving the results. I find significant evidence of discrimination using a 70% cutoff, 90% cutoff, or using the continuous probability that the filing is against a minority tenant without excluding any low probability cases. In panel B of Table A5, I follow Elzayn et al. (2025) and report probabilistic and

²⁹This is smaller than the 120,328 cases reported in column 2 of Table 2 because the sample of 120,328 cases includes landlord-by-contract rent cells with multiple filings, but some of these cells only include a single race. I include cases in these single-race cells in the analysis sample to identify the parameters for non-race characteristics in the Section 6.2 analysis.

linear estimators of conditional disparities that bound the true racial disparity under certain conditions. The probabilistic estimator uses the race probabilities as weights to estimate the minority-white difference in probability-weighted average log back rent within each landlord-by-contract rent cell, and reports the weighted average of the probabilistic disparity across cells.³⁰ The linear estimator is based on separate regressions of log back rent on the minority probability for each landlord-contract rent cell, and I report the weighted average of the minority coefficient across cells. I note that the magnitude of the linear estimator is somewhat sensitive to the inclusion of small landlord by contract rent cells with small ranges of minority probabilities. For this reason, I report estimates using \$25 bins of contract rent and compare estimates using various restrictions on cells with small minority probability ranges.

From [Elzayn et al. \(2025\)](#), the probabilistic estimator is an upper bound on the true conditional disparity when

$$\mathbb{E}[\text{Cov}(X_i, M_i | m_i, J_i, Y_i^0) | J_i = j, Y_i^0 = y] \leq 0 \quad (15)$$

for the average landlord-by-contract rent cell, where M is an indicator for true minority status, m is the continuous probability of minority status, J is the landlord, Y^0 is the contract rent, and X is back rent. The linear estimator is a lower bound when

$$\mathbb{E}[\text{Cov}(X_i, m_i | M_i, J_i, Y_i^0) | J_i = j, Y_i^0 = y] \leq 0 \quad (16)$$

for the average landlord-by-contract rent cell. While I cannot verify whether these bounding conditions outlined are satisfied in my setting, I highlight that my main estimate from [Table 2](#) tends to fall in the range between the linear and probabilistic estimators. In particular, the probabilistic estimator—which is less sensitive to small cells—is significantly different from 0, which allows me to reject equal filing thresholds by true race if [Equation 15](#) holds.

6.3.3 Landlord Identifiers

The landlord identifiers based on linked networks of shared plaintiff names and plaintiff addresses are intended to group inconsistent spellings and property-specific companies that may be managed by a single owner or property manager. When constructing these landlord identifiers, the linkage algorithm finds eleven landlords that account for 43.6% of all residential eviction filings ([Table B1](#)). While this may reflect true eviction filing volumes by large landlords, it is also possible that rental properties that switch management while maintaining the same ownership could generate false links between otherwise separate property

³⁰For this probabilistic estimator and the linear estimator discussed below, I report the average conditional disparity across landlords and contract rents, weighting by the filing share of each landlord and contract rent.

managers.

To account for this potential concern, I report estimates using much more conservative landlord definitions in Table A6. Compared to the main estimates reported in column 1, column 2 defines plaintiff names that are associated with different plaintiff addresses in different time periods (potentially reflecting a change in the property manager) as separate landlords. Columns 3 and 4 use simply the plaintiff address or the plaintiff name to define landlords. These narrower definitions of landlords likely fail to fully account for shared ownership or management across multiple buildings, but I still find significant evidence of discrimination. Compared to the main results, the uncontrolled coefficients are smaller in magnitude when restricting to comparisons within more narrowly defined landlords, but the controlled coefficients are quite similar across landlord definitions.

6.3.4 Discrete Back Rent Assumption

Measuring total discrimination also requires the assumption that the amount by which back rent exceeds the landlord’s threshold at the time of filing is race-invariant conditional on landlord and contract rent (Equation 13). If instead the last discrete change in back rent prior to the filing is statistically smaller for minority tenants, then white tenants’ back rent would be further above the landlord’s threshold at the time of the filing than that owed by minority tenants. This would negatively bias my estimate of total racial discrimination. In Appendix D, I demonstrate that while minority tenants appear to have slightly smaller discrete changes in back rent in the data, the magnitude could only explain up to 8% of my total discrimination estimate.

6.3.5 Time-Varying Filing Thresholds

I model a landlord’s eviction filing threshold, τ_i , for tenant i as static in Section 4. However, τ_i could be time-varying if, for example, macroeconomic conditions influence the cost c of filing an eviction case or payoffs from repossessing the property to rent to a new tenant.³¹ If the minority share of renters is correlated with such conditions, then β could be biased. To address this issue, I control flexibly for time trends in filing thresholds in Table A7. Column 1 controls for a cubic function of the filing date, column 2 includes filing year fixed effects, and column 3 includes filing month-year fixed effects. To account for any regional shocks within the city of Philadelphia, I include zip code-by-year fixed effects in

³¹For example, periods with high eviction filing volumes may have slower case progression, leading to universally lower filing thresholds. Alternatively, periods of high unemployment equate to a lower likelihood of the landlord finding a better-paying tenant, and may lead to universally higher filing thresholds.

column 4. The results are robust to the inclusion of these controls, suggesting minimal scope for bias from time-varying filing thresholds.

6.3.6 Omitted Lease Violations

By focusing on eviction cases caused only by non-payment of rent, I assume back rent owed at the time of filing fully captures how qualified a tenant is for an eviction filing. However, some landlords indicate the cause of the filing is only non-payment but also claim physical damages or other fees in the filing (Table A2). This may indicate some unobserved lease violations in the sample of cases appearing to be caused only by non-payment of rent. To ensure unobserved lease violations are not biasing my estimate of total discrimination, I further restrict the sample to cases caused only by non-payment of rent with \$0 of physical damages or other fees claimed in the filing in Table A8. The main results are robust to using this further restricted sample (column 1). I also find similar evidence when comparing the sum of back rent and late fees owed at filing in this sample (column 2).

Conversely, some eviction filings indicate that some breach of the lease occurred in addition to non-payment of rent, but the description of the additional breach of the lease lists only non-payment. It is reasonable to assume these cases have no omitted lease violations after accounting for back rent. I demonstrate in column 3 of Table A8 that my results are also robust to including these additional cases in the analysis sample.

6.3.7 Measuring Non-Payment of Rent

My total discrimination test allows for equal thresholds across race measured in either dollars or months of back rent to be considered non-discriminatory behavior. However, an additional unit of measure to consider for thresholds of non-payment is the ratio of back rent to a tenant's tenure in the rental property. For example, a landlord may make different filing decisions for a new tenant who failed to pay their first rental payment and a longstanding tenant who gradually accrued an equivalent amount of back rent. One could argue that this form of selective eviction filing behavior is non-discriminatory if it is reasonable for a landlord to measure non-payment of rent in terms of back rent relative to tenure. I allow for landlords to utilize this additional measure of non-payment by including landlord-by-contract rent-by-tenure fixed effects in the total discrimination test in Table A9.³² I find similar and if anything, larger magnitudes when also conditioning on tenure. This suggests that total racial discrimination remains even if landlords using equal filing thresholds of back rent divided by tenure is considered non-discriminatory behavior.

³²Due to power concerns from the high-dimensionality of these fixed effects and since tenure is only observable for cases filed after February 2011, I use \$25 bins of contract rent for this test.

6.3.8 Notice Period

Philadelphia law mandates that landlords provide tenants with ten days notice prior to filing eviction cases over non-payment of rent. In practice, however, many eviction cases are filed with substantially less than ten days notice or no notice at all due to notice waivers in lease agreements (Hoffman and Strezhnev 2022). Among tenants that do receive an eviction notice, some may quickly pay back rent or vacate the property to avoid an eviction filing. If such responses to a notice vary by race or if landlords provide different amounts of notice by race, my estimates of discrimination could include selection bias. In particular, if filings against white tenants owing low amounts of back rent are disproportionately censored from the eviction filing data due to those tenants’ responses to notices, then I may incorrectly detect lower thresholds for minority tenants.

To understand the extent to which notice responses could explain my main results, I re-estimate total discrimination among sub-samples of eviction filings with varying amounts of notice provided. Among tenants receiving little notice prior to an eviction filing, there is much less scope for responses to notices to generate selection bias. In Figure A3, I find that even among these cases with short notice periods, filing thresholds appear to be lower for minority tenants.³³ This suggests that any selection bias from notice responses cannot fully explain the main results.

6.4 Landlord-Specific Discrimination

The estimates thus far reflect the average amount of total racial discrimination across landlords in my sample. This could be driven by a small number of highly-discriminatory landlords or reflect widespread discrimination among Philadelphia landlords. To better understand the prevalence of discrimination in eviction filing, I estimate landlord-specific coefficients. To do so, I estimate an alternative version of Equation 14 that allows for the total discrimination parameter to vary by landlord:

$$\text{Ln}(\text{Back Rent})_{ijt} = \beta_j \text{Minority}_i + \alpha_{j,y(ijt)} + \varepsilon_{ijt} \quad (17)$$

Estimates of β_j are most informative for landlords with a sufficient number of filings against white and minority tenants with the same contract rent, but many landlords file few eviction cases. Therefore, I adopt three different approaches to estimate β_j , varying the filing volume of landlords required for inclusion.

³³Due to power concerns from the small samples of filings with very short or non-existent notice periods, I use \$25 bins of contract rent instead of the dollar value of contract rent. I also report estimates controlled for the amount of notice provided.

First, I restrict to cases that are found in landlord-by-contract rent cells with at least two non-payment of rent filings against white tenants and minority tenants. Due to the small size of many landlord-by-contract rent cells, this method can only identify landlord-specific discrimination coefficients for 127 unique landlords, but these landlords are responsible for 51% of Philadelphia eviction cases. As a second approach, I estimate β_j parameters from a regression with landlord and contract rent fixed effects included separately instead of interacted, and I again restrict to landlords that filed at least two non-payment-based filings against each racial group. This second approach allows me to estimate β_j coefficients for 520 unique landlords who are responsible for 59% of Philadelphia eviction cases. Finally, as a third approach, I again estimate β_j parameters using additive landlord and contract rent fixed effects but remove the restriction to landlords that filed at least two non-payment cases against each racial group. This third approach allows me to estimate β_j coefficients for 1,274 unique landlords who are responsible for 64% of all Philadelphia eviction cases.

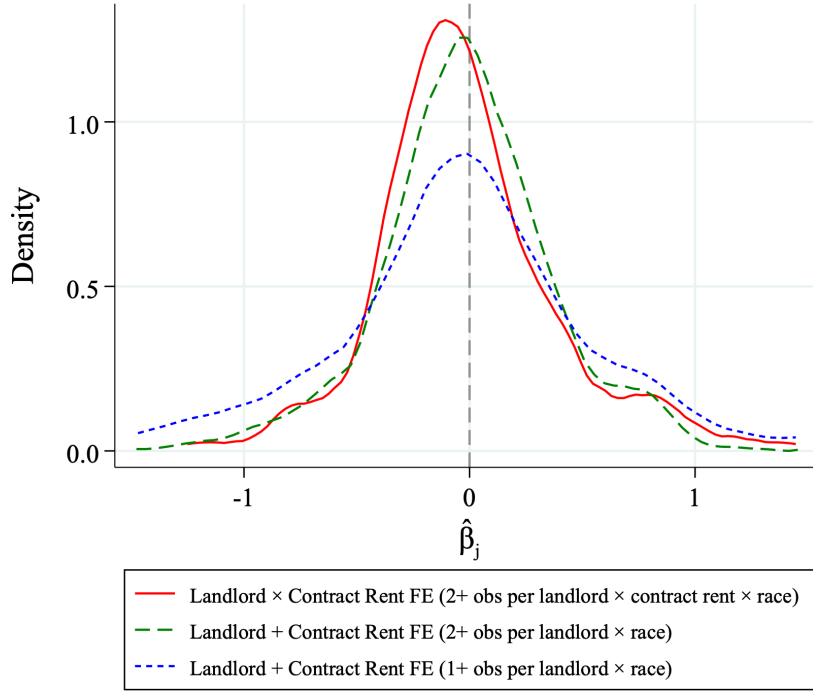
I plot the distributions of these landlord-specific coefficients in Figure 3. The median $\hat{\beta}_j$ from the three methods are -0.0683, -0.0233, and -0.0281, respectively, suggesting that the median landlord exhibits total racial discrimination in eviction filing. From Figure 3, however, it remains unclear what share of landlords exhibit total discrimination. The standard errors of each individual β_j estimated by Equation 17 are quite imprecise due to the often small number of eviction cases filed by a single landlord.³⁴

To account for noise in the $\hat{\beta}_j$ s and estimate the share of landlords exhibiting total racial discrimination, I follow [Goncalves and Mello \(2021\)](#) by modeling the estimation error in $\hat{\beta}_j$ as normal with variances taken from the estimated standard errors. I then estimate the true distribution of landlord-specific parameters by maximum likelihood and report the shares of discriminatory landlords ($\beta_j < 0$) and reverse-discriminatory landlords ($\beta_j > 0$) in Table A10. I estimate that at least 49% of landlords in the sample exhibit any discrimination against minority tenants, and at least 40% of landlords exhibit substantial discrimination ($\beta_j < -0.02$). Across all approaches, I find that a higher share of landlords exhibit total discrimination against minority tenants than reverse-discriminate against white tenants. It is important to note that these estimated discriminatory shares reflect only the selected sample of landlords where $\hat{\beta}_j$ is identified and are estimated with considerably more noise than is achieved by [Goncalves and Mello \(2021\)](#) given that many landlords file a small number of eviction cases.³⁵

³⁴The means of the standard errors of the $\hat{\beta}_j$ s plotted in Figure 3 are 0.24, 0.25, and 0.19, respectively.

³⁵[Goncalves and Mello \(2021\)](#) develop this method for estimating the share of discriminatory agents for the setting of traffic stops. Given that police officers stop a much higher number of drivers than landlords file eviction cases, their results are more precise than the results in this paper.

Figure 3: Landlord-Specific Estimates of Total Discrimination



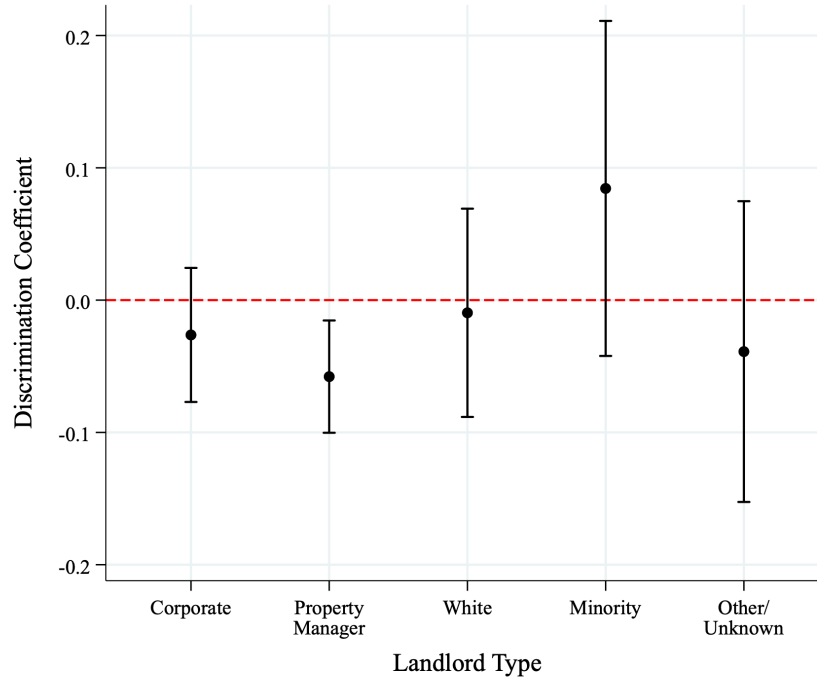
This figure plots the distribution of landlord-specific discrimination coefficients from estimating Equation 17. The sample restricts to eviction filings caused only by non-payment of rent that could be classified as white or minority with at least 80% certainty based on tenant race imputations. The estimates plotted in the solid red line are controlled for landlord-by-contract rent fixed effects and estimated from a sample of cases found in landlord-by-contract rent cells with at least two cases against both racial groups. The estimates plotted in the green dashed line are controlled for landlord and contract rent fixed effects separately and restrict to landlords with at least two non-payment cases filed against both racial groups. The estimates plotted in the blue dashed line are controlled for landlord and contract rent fixed effects separately and restrict to landlords with at least one non-payment cases filed against each racial group.

6.5 Heterogeneity

I also explore heterogeneity by landlord, tenant, rent, and neighborhood characteristics to understand which tenants are most vulnerable to discrimination in eviction filing. Figure 4 plots estimates separately for various types of landlords, which I classify based on plaintiff names and addresses.³⁶ Discrimination appears to be driven by non-minority landlords, while

³⁶I classify landlords as large property managers if no single plaintiff name is used in at least 50% of their filings and the landlord filed at least 25 cases. For non-property managers, I classify landlords as corporate if the most common plaintiff name contains “LLC”, “LP”, “LTD”, “INC”, or similar strings. For the remaining landlords, I impute their race using the fBISG methodology given the first name, last name, and ZCTA of the most common plaintiff name and address associated with the landlord. Based on these classifications, large property managers file 36% of cases, corporate landlords file 18% of cases, white individual landlords file 8% of cases, minority individual landlords file 16% of cases, and unclassified landlords file 22% of cases.

Figure 4: Heterogeneity by Landlord Type



This figure plots the discrimination coefficients by landlord type. The outcome variable is log back rent, and the coefficients are controlled for landlord-by-contract rent fixed effects. The sample is restricted to Philadelphia eviction filings caused only by non-payment of rent that could be classified as white or minority with at least 80% certainty based on tenant race imputations. Standard errors are clustered at the landlord level.

minority landlords if anything may exhibit reverse-discrimination against white tenants.³⁷ In Section 7, I discuss the likely drivers of this pattern of heterogeneity. Examining individual tenant, rent, and neighborhood characteristics in Figure A4, I find no clear patterns of heterogeneous discrimination by gender, tenure in the rental property, contract rent, rent affordability, tract racial composition, or tract poverty rates.

7 Sources of Discrimination

Having found evidence of racial discrimination in eviction filing with no evidence that differences in tenants' non-race characteristics drive the results, I next turn to examine the possible sources of discrimination: accurate statistical discrimination or racial bias. From Section 4.2, accurate statistical discrimination may arise if landlords account for any true de-

³⁷I can reject that discrimination by individual minority landlords is equal to discrimination by all other landlords (p-value = 0.043).

pendence of the probability of winning repossession on tenant race. Alternatively, landlords may exhibit racial bias if they deviate from accurate statistical discrimination and tolerate different true payoffs from filing eviction cases by tenant race.

In practice, winning repossession appears to depend on tenant race (Figure 2), motivating the possibility of accurate statistical discrimination. The data suggest that the pattern could be driven by landlords withdrawing fewer cases against minority tenants and that minority tenants are less likely to have legal representation (Figure A5).³⁸ However, these potential explanations are suggestive as I am unable to investigate other possible channels such as differences in bargaining power during private landlord-tenant negotiations, access to information, or discrimination by judges in the small number of cases they decide.³⁹ Regardless of the explanation for Figure 2, if landlords account for the observable pattern that it seemingly requires less back rent to win repossession from minority tenants, then accurate statistical discrimination may explain total discrimination in eviction filing.

A testable implication of accurate statistical discrimination is that a landlord should win repossession of the property at the same rate against a white and minority tenant at the margin conditional on non-race characteristics and contract rent (Proposition 1). Any deviation from equal repossession rates reflects racial bias (Proposition 2). In practice, I do not observe all possible non-race characteristics that the landlord may observe and thus conduct the marginal outcome test controlling for the same observable tenant characteristics as in Section 6.2. I conduct this marginal outcome test using the following specification for filings against tenant(s) i by landlord j on filing date t :

$$Repossession_{ijt} = \theta Minority_i + \gamma_{j,y(ijt)} + X'_{ijt}\lambda + v_{ijt} \quad (18)$$

An estimate of $\theta < 0$ is evidence of bias against minority tenants and $\theta > 0$ is evidence of bias against white tenants under two assumptions: marginal outcomes and no omitted payoff bias (Hull 2021; Kleinberg et al. 2018). First, the test requires that all outcomes are marginal. This assumption is satisfied if back rent evolves continuously, ensuring that back rent owed at filing is equal to the landlord’s filing threshold for tenant i . Second, no omitted payoff bias requires that a landlord’s true payoff from winning repossession does not vary with tenant race after including landlord-by-contract rent fixed effects, $\gamma_{j,y(ijt)}$, and control variables, X_{ijt} , which include imputed gender, the number of tenants listed, the number of

³⁸Differential rates of court appearance does not seem to explain Figure 2 since minority tenants tend to appear in court more often than white tenants conditional on no case withdrawal.

³⁹It is important to highlight that compared to other jurisdictions where judges may play a larger role in eviction court outcomes (Collinson et al. 2024), Philadelphia eviction court is characterized by judges deciding less than 5% of eviction filings. This means that it is unlikely that the racial disparities in eviction court anticipated by landlords are driven in large part by judge behavior.

previous filings against the tenant at the same address, two proxy measures of rent burden, and census tract fixed effects.

If the omitted payoff bias assumption is violated and tenant race predicts future payment ability conditional on controls, then a non-zero estimate of θ would reflect some combination of racial bias or accurate statistical discrimination based on future payment ability. In this case, distinguishing racial bias from accurate statistical discrimination based on future payment ability would require observing a landlord’s future rental income to use as the outcome in the marginal outcome test. By including the control variables in Equation 18, the marginal outcome test should account for racial differences in future payment ability (the most likely omitted payoff) and isolate any racial bias.

The validity of the marginal outcome test also depends on accurately measuring the outcome—the landlord winning repossession of the rental property from the tenant during the court process.⁴⁰ One measure of *Repossession* is whether the landlord filed for a writ of possession, which aligns with the definition of an eviction order used by Collinson et al. (2024) in New York City eviction cases. A drawback of this measure is that it may miss informal repossessions. For example, after winning a judgment, a landlord may not file for a writ of possession if the tenant has already vacated the property.⁴¹ If unobserved repossessions vary by race, this represents another form of omitted payoff bias. To address this, I repeat the marginal outcome test using two broader measures of repossession: including JBAs in which the tenant agrees to vacate or already vacated the property, and additionally including any judgment in favor of the landlord regardless of whether a writ of possession was filed. While each of these measures of repossession has drawbacks, consistent results across these measures should support the interpretation of the outcome test.

Table 4 reports the results of the marginal outcome test. Landlords win repossession of the property at statistically indistinguishable rates across tenant race, which is consistent with accurate statistical discrimination based on court outcomes. The results are somewhat imprecise, but I can reject that landlords tolerate more than 1.22 percentage point (2.77%) lower repossession rates from minority tenants (as measured by writs of possession). Evidence that landlords use lower filing thresholds for minority tenants but achieve indistinguishable

⁴⁰It is possible that landlords may have other goals when filing an eviction case, such as filing as a threat to induce tenants to pay back rent (Garboden and Rosen 2019). A marginal outcome test that uses a landlord’s future rental earnings from the rental unit as the outcome variable could account for the various aims of landlords when filing eviction cases. Given that landlords’ rental earnings are unobservable, I simplify the setting and assume that landlords intend to win repossession of the unit.

⁴¹In the data, a substantial number of cases with a judgment in favor of the landlord have no writ of possession filed (Figure A6), suggesting that the writ of possession measure may indeed underestimate repossession. However, some of these cases could reflect a landlord reaching a private agreement with a tenant post-judgment that allows the tenant to remain in the property.

Table 4: Marginal Outcome Test

	Writ of Possession	or Agree to Leave via JBA	or Landlord Won Judgment
	(1)	(2)	(3)
Minority Tenant	0.0143 (0.0135)	0.0108 (0.0131)	0.0062 (0.0088)
Observations	120,328	120,328	120,328
R ²	0.1907	0.1945	0.2211
White Outcome Mean	0.4380	0.4783	0.6470
Landlord \times Contract Rent FE	✓	✓	✓
Tenant Controls	✓	✓	✓
Census Tract FE	✓	✓	✓

* $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$. This table reports the results from estimating Equation 18. The sample includes Philadelphia eviction filings caused only by non-payment of rent that could be classified as white or minority with at least 80% certainty based on tenant race imputations. Contract rent fixed effects refer to the exact dollar value of contract rent. Tenant controls include the imputed gender of first tenant named, fixed effects for the number of tenants listed on the filing, fixed effects for the number of previous cases filed against the tenant at the same property, and fixed effects for two measures of rent burden each divided into 2-percentage point bins. The first measure of rent burden is the ratio of contract rent to median household income among renters in the census tract. The second measure of rent burden is the ratio of contract rent to median household income by race and census tract. Median household income by race and tract is defined based on the imputed race of the first tenant listed if the imputation is at least 80% likely. All control variables include a category for missing values. The outcome variable in column 1 is an indicator for whether a landlord filed for a writ of possession. The outcome variable in column 2 additionally includes cases where the tenant agreed to vacate the property as part of a judgment by agreement (JBA). The outcome variable in column 3 additionally includes any case in which the landlord won a default judgment or won a contested judgment (decided by the judge) regardless of filing for a writ of possession. All specifications exclude singleton observations for the fixed effects used. Standard errors are clustered at the landlord level.

repossession rates across tenant race is consistent with landlords accurately anticipating that winning repossession from minority tenants requires less back rent than winning repossession from white tenants.

Table 5 quantifies this underlying racial disparity in repossession that landlords appear to accurately predict.⁴² By conditioning on landlord and back rent owed at filing, these estimates can be interpreted as counterfactual repossession disparities if landlords were to use race-invariant filing thresholds. The magnitudes of the underlying racial disparities in repossession (Table 5) are consistently larger than the realized repossession disparities after landlords discriminate at filing (Table 4). This pattern further supports the finding that landlords exhibit accurate statistical discrimination: by filing eviction cases at lower

⁴²In essence, the estimates in Table 5 more precisely quantify the disparity visualized in Figure 2.

Table 5: Underlying Racial Disparities in Repossession

	Writ of Possession		or Agree to Leave via JBA		or Landlord Won Judgment	
	(1)	(2)	(3)	(4)	(5)	(6)
Minority Tenant	0.0567*** (0.0218)	0.0483*** (0.0178)	0.0623*** (0.0204)	0.0528*** (0.0158)	0.0527** (0.0212)	0.0421*** (0.0162)
Observations	116,403	127,358	116,403	127,358	116,403	127,358
R ²	0.1291	0.1231	0.1375	0.1310	0.1608	0.1539
White Outcome Mean	0.4182	0.4216	0.4524	0.4593	0.6248	0.6342
Landlord × Back Rent (\$) FE	✓		✓		✓	
Landlord × Back Rent (Mths) FE		✓		✓		✓

*p < 0.1, **p < 0.05, ***p < 0.01. The sample includes Philadelphia eviction filings caused only by non-payment of rent that could be classified as white or minority with at least 80% certainty based on tenant race imputations. Back rent fixed effects in dollars are bins defined at \$50 increments. Back rent fixed effects in months owed are bins of back rent divided by contract rent defined at increments of 0.25 months. The outcome variable in columns 1 and 2 is an indicator for whether a landlord filed for a writ of possession. The outcome variable in columns 3 and 4 additionally includes cases where the tenant agreed to vacate the property as part of a judgment by agreement (JBA). The outcome variable in columns 5 and 6 additionally includes any case in which the landlord won a default judgment or won a contested judgment (decided by the judge) regardless of filing for a writ of possession. All specifications exclude singleton landlord-by-back rent observations. Standard errors are clustered at the landlord level.

thresholds of back rent for minority tenants, landlords appear to equalize the otherwise-unequal repossession rate.

Recall from Figure 4 that minority landlords appear to differ from other landlords by possibly exhibiting reverse-discrimination against white tenants. This pattern of heterogeneity could be rationalized with the finding of accurate statistical discrimination if the underlying racial disparity in repossession rates varies by landlord type. For example, it is possible that minority landlords can achieve more favorable outcomes when negotiating with white tenants, while white landlords can achieve more favorable outcomes when negotiating with minority tenants. Indeed, Table A11 shows that the underlying repossession disparity has the opposite sign for minority landlords. This is consistent with all landlords exhibiting accurate statistical discrimination: most landlords tolerate less back rent from minority tenants because they accurately predict that less back rent is required for them to win against minority tenants, but minority landlords tolerate less back rent from white tenants because they accurately predict that less back rent is required for them to win against white tenants.

Taken together, the evidence suggests that landlords account for racial disparities at later stages of the eviction process when deciding to file a case. This implies that policy responses to eliminate the racial disparities in eviction court that affect the repossession rate should reduce discrimination by landlords at the margin of eviction filing. Further upstream, however, landlords could also change how they select tenants for rental units in

response to reduced disparities in eviction court. With the data analyzed in this paper, I am unable to investigate this potential margin of adjustment, meaning that welfare implications of eliminating racial disparities in eviction court are unclear.

8 Conclusion

In this paper, I develop novel methods for measuring discrimination in dynamic settings and find evidence of racial discrimination when landlords file eviction cases. On average, landlords tolerate 4.9% less back rent from minority tenants compared to white tenants with the same contract rent. At least 49% of landlords in my sample exhibit racial discrimination. This result appears to be fully driven by direct discrimination on the basis of tenant race and not other characteristics correlated with race.

This evidence of racial discrimination at the margin of housing loss reveals that minorities face significant discrimination in the housing market even after entry into housing. Racial discrimination in this context is particularly troubling given the severe consequences of eviction. Regardless of the case outcome, eviction filings generate a public record that is visible on tenant screening reports for many years, and evidence suggests these public records can increase housing instability (Liao et al. 2026). Higher exposure to eviction filings may disproportionately expose minority tenants to the adverse effects of housing displacement (Collinson et al. 2024, 2025). Landlords tolerating less back rent from minority tenants equates to less informal credit provided by landlords, amplifying known racial disparities in formal credit access (Weller 2009; Cohen-Cole 2011).

My results suggest that racial discrimination in eviction filing can be explained by landlords exhibiting accurate statistical discrimination. Landlords observe that winning repossession from minority tenants requires less back rent. By filing eviction cases at lower thresholds against minority tenants, landlords equalize the repossession rate. This finding implies that policy responses to reduce racial disparities in eviction court may also reduce racial discrimination at the earlier stage of eviction case filing. Right-to-counsel laws for tenants in eviction court and eviction diversion programs could theoretically impact the scope for statistical discrimination by reducing racial gaps in the repossession rate. Other policies that directly regulate when landlords can file eviction cases—for example, “good cause” laws or back rent floors—may also reduce discrimination in eviction filing. However, these types of policy interventions that make evicting tenants more difficult could also induce landlords to adopt stricter tenant screening than could negatively impact minority renters (Arefeva et al. 2024). Evaluating the welfare impacts of recent eviction policies that include any discrimination effects remains an important area for future work.

References

- Agan, A. and Starr, S. (2018). Ban the Box, Criminal Records, and Racial Discrimination: A Field Experiment. *Quarterly Journal of Economics*, 133(1):191–235.
- Ahmed, A. M. and Hammarstedt, M. (2008). Discrimination in the Rental Housing Market: A Field Experiment on the Internet. *Journal of Urban Economics*, 64(2):362–372.
- Aigner, D. J. and Cain, G. G. (1977). Statistical Theories of Discrimination in Labor Markets. *Industrial and Labor Relations Review*, 30(2):175–187.
- Alesina, A. and La Ferrara, E. (2014). A Test of Racial Bias in Capital Sentencing. *American Economic Review*, 104(11):3397–3433.
- Anwar, S. and Fang, H. (2006). An Alternative Test of Racial Prejudice in Motor Vehicle Searches: Theory and Evidence. *American Economic Review*, 96(1):127–151.
- Anwar, S. and Fang, H. (2015). Testing for Racial Prejudice in the Parole Board Release Process: Theory and Evidence. *Journal of Legal Studies*, 44(1):1–37.
- Arefeva, A., Jowers, K., Hu, Q., and Timmins, C. (2024). Discrimination During Eviction Moratoria. *National Bureau of Economic Research Working Paper 32289*.
- Arnold, D., Dobbie, W., and Hull, P. (2022). Measuring Racial Discrimination in Bail Decisions. *American Economic Review*, 112(9):2992–3038.
- Arnold, D., Dobbie, W., and Yang, C. S. (2018). Racial Bias in Bail Decisions. *Quarterly Journal of Economics*, 133(4):1885–1932.
- Arrow, K. J. (1974). The Theory of Discrimination. In *Discrimination in Labor Markets*. Princeton University Press, Princeton, NJ.
- Ayres, I. (2002). Outcome Tests of Racial Disparities in Police Practices. *Justice Research and Policy*, 4(1-2):131–142.
- Baron, E. J., Doyle, J. J., Emanuel, N., Hull, P., and Ryan, J. (2024). Discrimination in Multiphase Systems: Evidence from Child Protection. *Quarterly Journal of Economics*, 139(3):1611–1664.
- Becker, G. S. (1957). *The Economics of Discrimination*. University of Chicago Press.

- Bertrand, M. and Mullainathan, S. (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.
- Bohren, J. A., Haggag, K., Imas, A., and Pope, D. G. (2025). Inaccurate Statistical Discrimination: An Identification Problem. *Review of Economics and Statistics*, 107(3):605–620.
- Bohren, J. A., Hull, P., and Imas, A. (2022). Systemic Discrimination: Theory and Measurement. *National Bureau of Economic Research Working Paper 29820*.
- Bordalo, P., Coffman, K., Gennaioli, N., and Shleifer, A. (2016). Stereotypes. *Quarterly Journal of Economics*, 131(4):1753–1794.
- Canay, I. A., Mogstad, M., and Mountjoy, J. (2024). On the Use of Outcome Tests for Detecting Bias in Decision Making. *Review of Economic Studies*, 91(4):2135–2167.
- Carpusor, A. G. and Loges, W. E. (2006). Rental Discrimination and Ethnicity in Names. *Journal of Applied Social Psychology*, 36(4):934–952.
- Chen, J., Kallus, N., Mao, X., Svacha, G., and Udell, M. (2019). Fairness Under Unawareness: Assessing Disparity When Protected Class Is Unobserved. In *Proceedings of the Conference on Fairness, Accountability, and Transparency*, pages 339–348.
- Christensen, P., Sarmiento-Barbieri, I., and Timmins, C. (2021). Racial Discrimination and Housing Outcomes in the United States Rental Market. *National Bureau of Economic Research Working Paper 29516*.
- Cohen-Cole, E. (2011). Credit Card Redlining. *Review of Economics and Statistics*, 93(2):700–713.
- Collinson, R., Dutz, D., Humphries, J. E., Mader, N., Tannenbaum, D., and Van Dijk, W. (2025). The Effects of Eviction on Children. *National Bureau of Economic Research Working Paper 33659*.
- Collinson, R., Humphries, J. E., Mader, N., Reed, D., Tannenbaum, D., and Van Dijk, W. (2024). Eviction and Poverty in American Cities. *Quarterly Journal of Economics*, 139(1):57–120.
- Diamond, R., Guren, A., and Tan, R. (2020). The Effect of Foreclosures on Homeowners, Tenants, and Landlords. *National Bureau of Economic Research Working Paper 27358*.

- Dobbie, W., Liberman, A., Paravisini, D., and Pathania, V. (2021). Measuring Bias in Consumer Lending. *Review of Economic Studies*, 88(6):2799–2832.
- Edelman, B., Luca, M., and Svirsky, D. (2017). Racial Discrimination in the Sharing Economy: Evidence from a Field Experiment. *American Economic Journal: Applied Economics*, 9(2):1–22.
- Elzayn, H., Smith, E., Hertz, T., Guage, C., Ramesh, A., Fisher, R., Ho, D. E., and Goldin, J. (2025). Measuring and Mitigating Racial Disparities in Tax Audits. *Quarterly Journal of Economics*, 140(1):113–163.
- Ewens, M., Tomlin, B., and Wang, L. C. (2014). Statistical Discrimination or Prejudice? A Large Sample Field Experiment. *Review of Economics and Statistics*, 96(1):119–134.
- Garboden, P. M. and Rosen, E. (2019). Serial Filing: How Landlords use the Threat of Eviction. *City & Community*, 18(2):638–661.
- Goncalves, F. and Mello, S. (2021). A Few Bad Apples? Racial Bias in Policing. *American Economic Review*, 111(5):1406–1441.
- Graetz, N., Gershenson, C., Hepburn, P., Porter, S. R., Sandler, D. H., and Desmond, M. (2023). A Comprehensive Demographic Profile of the US Evicted Population. *Proceedings of the National Academy of Sciences*, 120(41).
- Greenberg, D., Gershenson, C., and Desmond, M. (2016). Discrimination in Evictions: Empirical Evidence and Legal Challenges. *Harvard Civil Rights-Civil Liberties Law Review*, 51:115–158.
- Gromis, A., Fellows, I., Hendrickson, J., Edmonds, L., Leung, L., Porton, A., and Desmond, M. (2022). Estimating Eviction Prevalence across the United States. <https://data-downloads.evictionlab.org/#estimating-eviction-prevalenceacross-us/>.
- Hanson, A. and Hawley, Z. (2011). Do Landlords Discriminate in the Rental Housing Market? Evidence from an Internet Field Experiment in US Cities. *Journal of Urban Economics*, 70(2-3):99–114.
- Hanson, A., Hawley, Z., Martin, H., and Liu, B. (2016). Discrimination in Mortgage Lending: Evidence from a Correspondence Experiment. *Journal of Urban Economics*, 92:48–65.
- Hoffman, D. A. and Strezhnev, A. (2022). Leases as Forms. *Journal of Empirical Legal Studies*, 19(1):90–134.

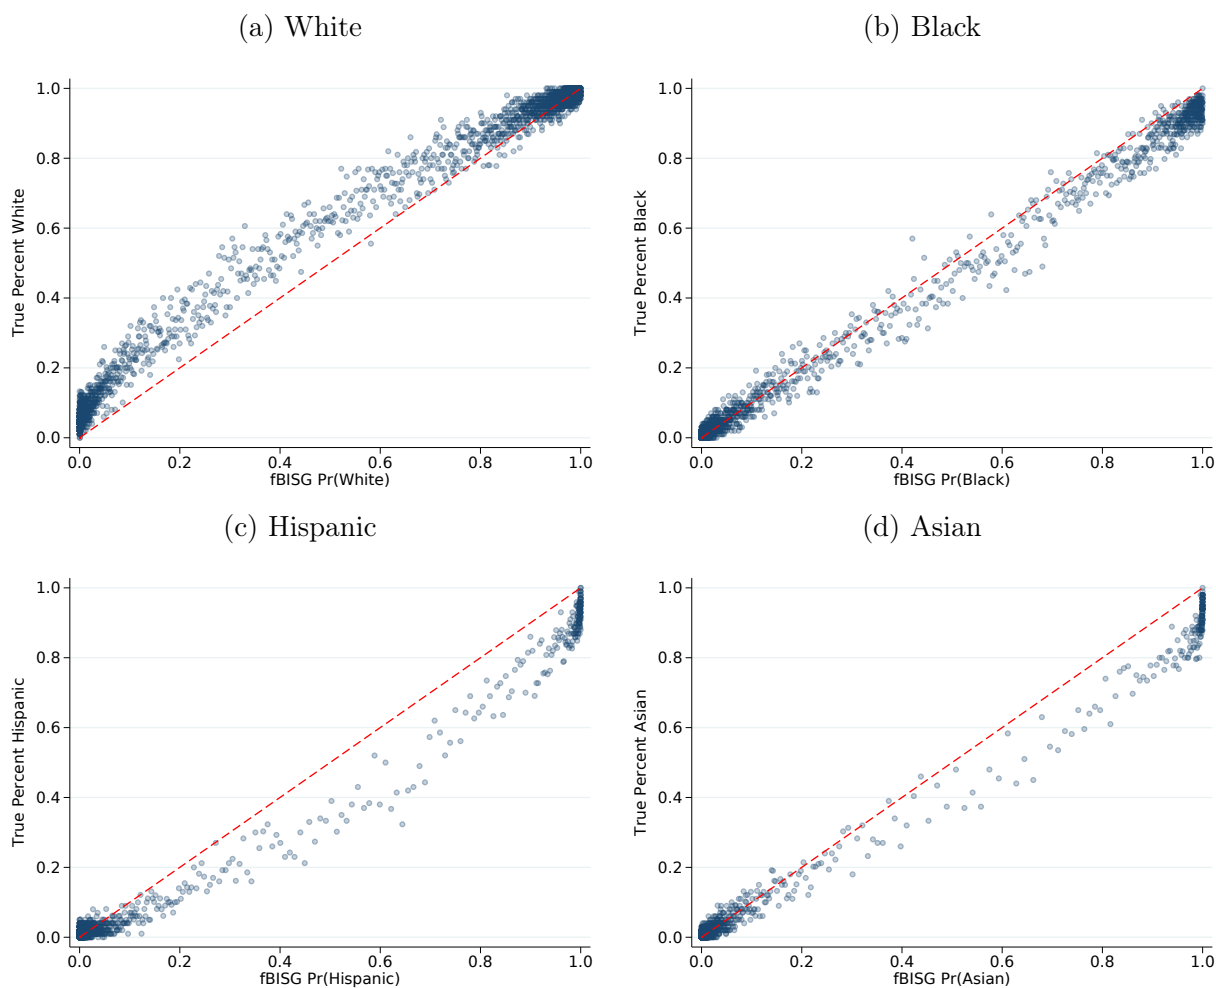
- HUD (2021a). HUD Statement on Fair Housing Obligations, Evictions, and Disparities Faced by Certain Protected Class Groups. https://www.hud.gov/program_offices/fair_housing_equal_opp/Fair_Housing_and_Evictions.
- HUD (2021b). Know Your Fair Housing Rights When You Are Facing Eviction. https://www.hud.gov/program_offices/fair_housing_equal_opp/Know_Your_Fair_Housing_Rights.
- HUD (2023). Reinstatement of HUD’s Discriminatory Effects Standard. Technical report, Office of the Assistant Secretary for Fair Housing and Equal Opportunity, U.S. Department of Housing and Urban Development.
- Hull, P. (2021). What Marginal Outcome Tests Can Tell Us About Racially Biased Decision-Making. *National Bureau of Economic Research Working Paper 28503*.
- Imai, K. and Khanna, K. (2016). Improving Ecological Inference by Predicting Individual Ethnicity from Voter Registration Records. *Political Analysis*, 24(2):263–272.
- Imai, K., Olivella, S., and Rosenman, E. T. R. (2022). Addressing Census Data Problems in Race Imputation via Fully Bayesian Improved Surname Geocoding and Name Supplements. *Science Advances*, 8(49).
- Jacquemet, N. and Yannelis, C. (2012). Indiscriminate Discrimination: A Correspondence Test for Ethnic Homophily in the Chicago Labor Market. *Labour Economics*, 19(6):824–832.
- Kleinberg, J., Lakkaraju, H., Leskovec, J., Ludwig, J., and Mullainathan, S. (2018). Human Decisions and Machine Predictions. *Quarterly Journal of Economics*, 133(1):237–293.
- Kline, P., Rose, E. K., and Walters, C. R. (2022). Systemic Discrimination Among Large U.S. Employers. *Quarterly Journal of Economics*, 137(4):1963–2036.
- Knowles, J., Persico, N., and Todd, P. (2001). Racial Bias in Motor Vehicle Searches: Theory and Evidence. *Journal of Political Economy*, 109(1):203–229.
- Ladd, H. F. (1998). Evidence on Discrimination in Mortgage Lending. *Journal of Economic Perspectives*, 12(2):41–62.
- Liao, H., Lodermeier, A., Ortuzar, G., and Stapleton, S. (2026). Is There a “Scarlet E”? The Effects of Public Eviction Records on Low-Income Households.

- Lodermeier, A. (2024). Credit Access and Housing Insecurity: Evidence from Winter Utility Shutoff Protections. *Journal of Public Economics*, 230.
- Manville, M., Monkkonen, P., Lens, M. C., and Green, R. (2023). Renter Nonpayment and Landlord Response: Evidence From COVID-19. *Housing Policy Debate*, 33(6):1333–1367.
- Marx, P. (2021). An Absolute Test of Racial Prejudice. *Journal of Law, Economics, and Organization*, 38(1):42–91.
- Munnell, A. H., Tootell, G. M. B., Browne, L. E., and McEneaney, J. (1996). Mortgage Lending in Boston: Interpreting HMDA Data. *American Economic Review*, 86(1):25–53.
- Oster, E. (2019). Unobservable Selection and Coefficient Stability: Theory and Evidence. *Journal of Business & Economic Statistics*, 37(2):187–204.
- Page, M. (1995). Racial and Ethnic Discrimination in Urban Housing Markets: Evidence from a Recent Audit Study. *Journal of Urban Economics*, 38(2):183–206.
- Persico, N. and Todd, P. (2006). Generalising the Hit Rates Test for Racial Bias in Law Enforcement, with an Application to Vehicle Searches in Wichita. *Economic Journal*, 116(515):F351–F367.
- Phelps, E. S. (1972). The Statistical Theory of Racism and Sexism. *American Economic Review*, 62(4):659–661.
- Reinvestment Fund (2020). Resolving Landlord-Tenant Disputes: An Analysis of Judgments by Agreement in Philadelphia’s Eviction Process. Technical report, Reinvestment Fund.
- Ross, S. L., Turner, M. A., Godfrey, E., and Smith, R. R. (2008). Mortgage Lending in Chicago and Los Angeles: A Paired Testing Study of the Pre-Application Process. *Journal of Urban Economics*, 63(3):902–919.
- Turner, M. A. and Skidmore, F. (1999). Mortgage Lending Discrimination: A Review of Existing Evidence. Technical report, The Urban Institute, Washington, DC.
- Voicu, I. (2018). Using First Name Information to Improve Race and Ethnicity Classification. *Statistics and Public Policy*, 5(1):1–13.
- Weller, C. E. (2009). Credit Access, the Costs of Credit and Credit Market Discrimination. *Review of Black Political Economy*, 36(1):7–28.
- Yinger, J. (1986). Measuring Racial Discrimination with Fair Housing Audits: Caught in the Act. *American Economic Review*, 76(5):881–893.

Appendix A Supplemental Figures and Tables

A.1 Supplemental Figures

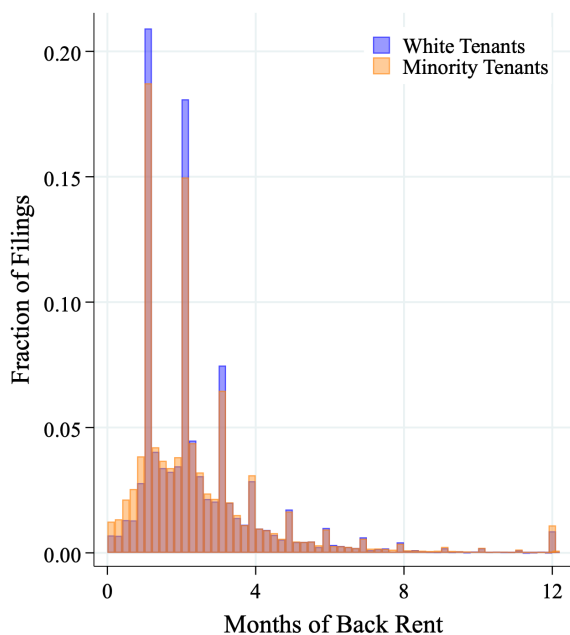
Figure A1: Out-of-Sample fBISG Validation



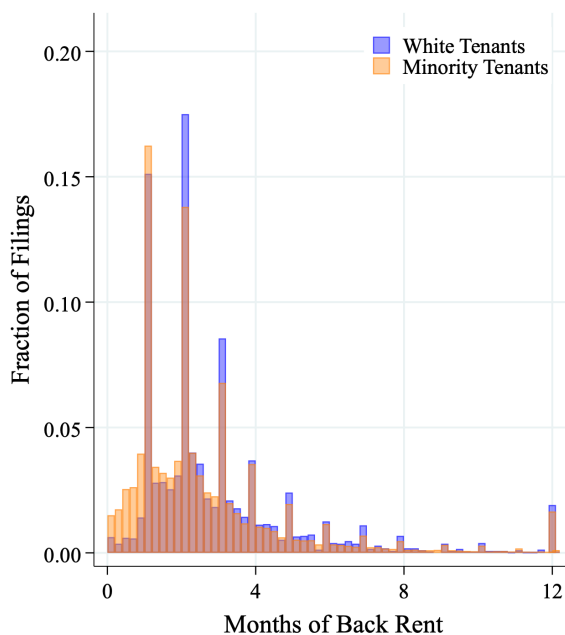
This figure plots a bin scatter of borrower-level fBISG probabilities that an individual belongs to a given racial/ethnic group against the true share of that racial/ethnic group among Philadelphia mortgages. Each point includes 100 borrowers. The red-dashed line is the 45 degree line, representing perfectly-calibrated predictions. The sample includes HMDA owner-occupied mortgages originated between 2007 and 2019 that could be linked to Philadelphia deeds sourced from CoreLogic. The linkage requires an exact match on census tract, year, loan type (conventional, FHA-insured, or VA-guaranteed), and a fuzzy match on the lender. The fBISG predictions are based on a borrower’s first and last name from the mortgage details associated with the CoreLogic record, and the census block of the parcel.

Figure A2: Months of Back Rent Owed at Filing

(a) All Tenants

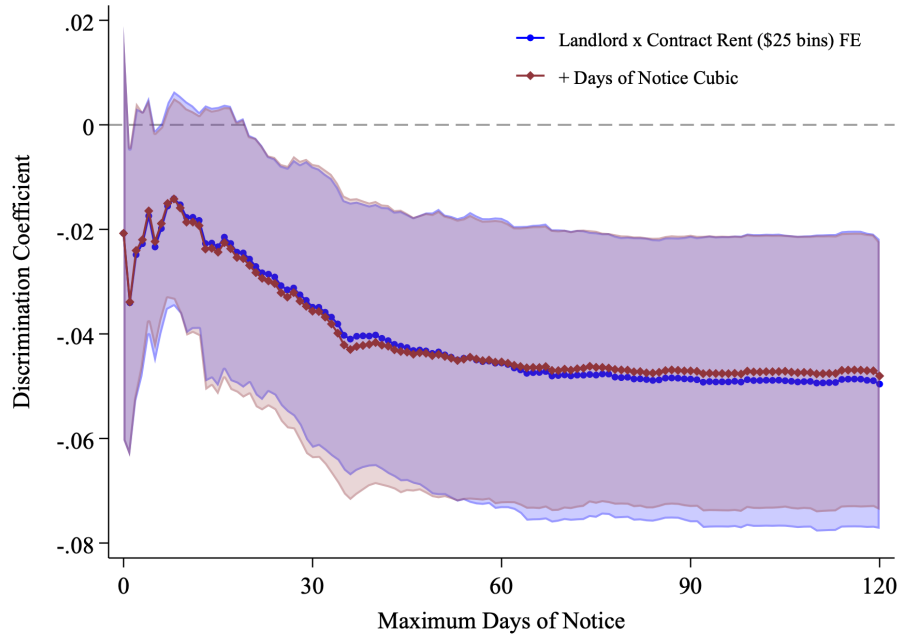


(b) Below Median Contract Rent



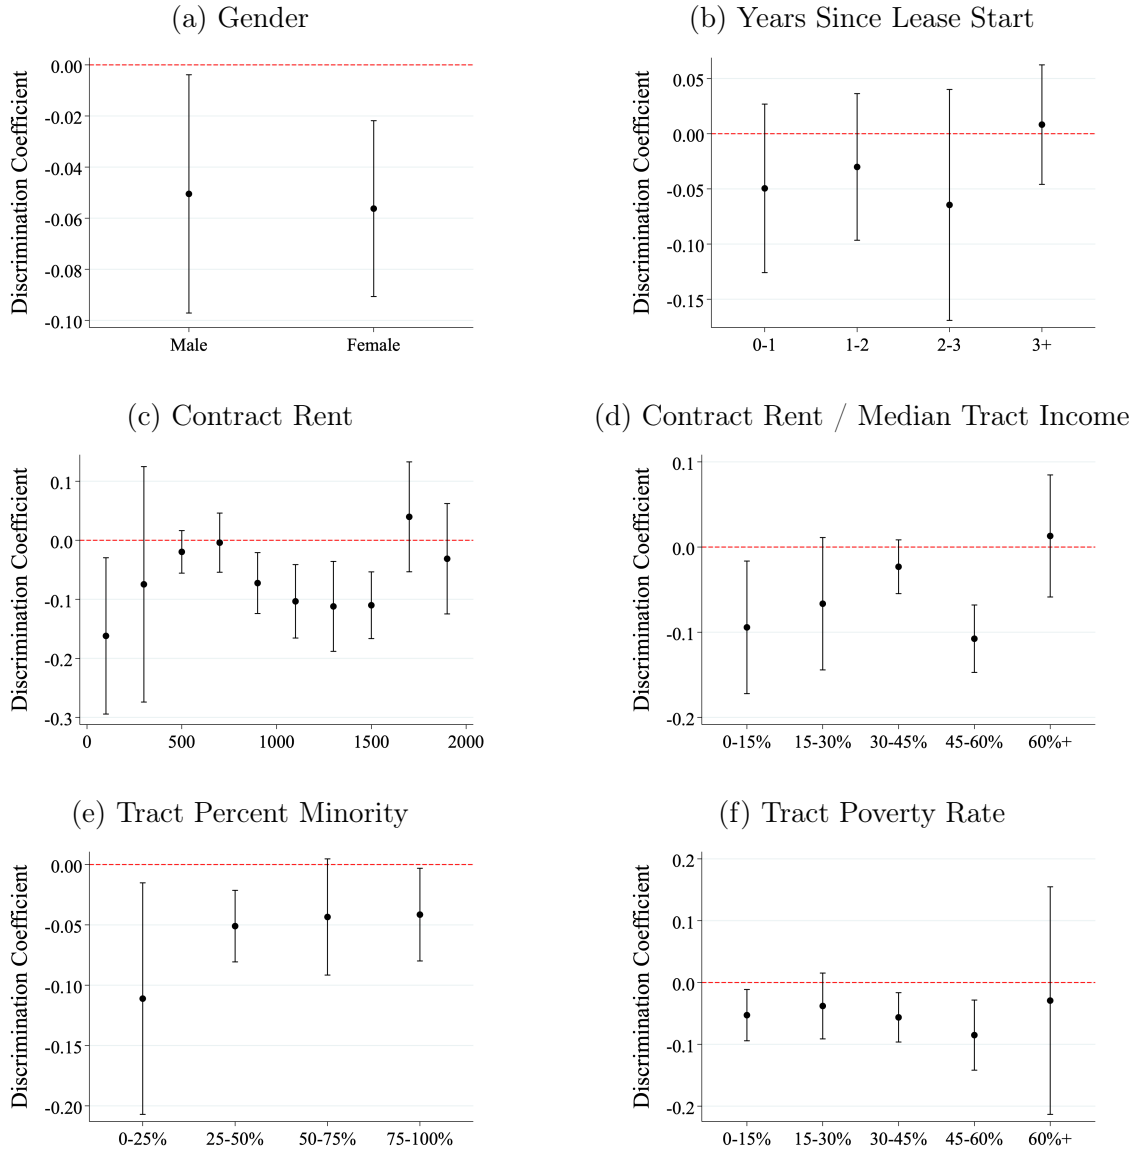
This figure plots the months of back rent distribution of Philadelphia eviction filings separately for filings against white and minority tenants. Months of back rent is equal to dollars of back rent divided by monthly contract rent. The sample in both subfigures restricts to Philadelphia eviction filings caused only by non-payment of rent that could be classified as white or minority with at least 80% certainty based on tenant race imputations. Subfigure (b) further restricts to tenants with contracts rents below the sample median in the year of the eviction filing. When plotting the distributions, months owed are top-coded at 12 months.

Figure A3: Total Discrimination Estimates By Notice Period



This figure plots coefficients from separately estimating Equation 14 using sub-samples of the eviction filing data with varying lengths of notice periods. Each point represents the minority coefficient estimated from the sample of filings with less than the corresponding number of days of notice provided. The number of days of notice for a filing is the difference in days between the date of the filing and the date any notice was provided. Filings with no notice date provided before the filing date have a notice period of 0 days.

Figure A4: Heterogeneity in Total Racial Discrimination




This figure plots the discrimination coefficients estimated from various heterogeneous versions of Equation 14. Subfigure (a) estimates discrimination by the imputed gender of the first tenant named in the filing based on the tenant's first name. Subfigure (b) estimates discrimination by the years the tenant has lived in the rental property prior to the eviction filing. Subfigure (c) estimates discrimination by \$200 contract rent bins. Subfigure (d) estimates discrimination by the ratio of contract rent to median household income among renters in the census tract. Median household income is sourced from 5-year ACS estimates, where the midpoint of the 5-year ACS period is matched to the case filing year. Subfigure (e) estimates discrimination by the minority population share of the census tract. Subfigure (f) estimates discrimination by the poverty rate among renters in the census tract. In all subfigures, the outcome variable is log back rent and the sample is restricted to Philadelphia eviction filings caused only by non-payment of rent that could be classified as white or minority with at least 80% certainty based on tenant race imputations. All coefficients are controlled for landlord-by-contract rent fixed effects interacted with the relevant tenant or neighborhood characteristic. Standard errors are clustered at the landlord level.

Figure A5: Sources of the Underlying Racial Disparity in Repossession



This figure plots bin scatters of the relationship between back rent and various measures of progress through the eviction court process. Each point reflects 500 eviction cases against tenants belonging to a given racial group. The x-axis in each subfigure is the median back rent owed among cases in each cell. Subfigure (a) uses a sample of Philadelphia eviction filings caused only by non-payment of rent that could be classified as white or minority with at least 80% certainty based on tenant race imputations. The y-axis in subfigure (a) is the rate of landlords withdrawing cases prior to the scheduled hearing date. In subfigure (b), the sample is further restricted to non-withdrawn cases, and the y-axis is the rate at which tenants appear in court when scheduled. In subfigure (c), the sample is further restricted to cases in which the tenant appeared in court, and the y-axis is the rate at which tenants had legal representation.

Figure A6: Philadelphia Eviction Case Outcomes



This figure plots the frequency of Philadelphia eviction case outcomes. The sample includes Philadelphia eviction filings caused only by non-payment of rent that could be classified as white or minority with at least 80% certainty based on tenant race imputations. Judgment for tenant includes contested cases where the judge ruled in favor of the tenant and default judgments in favor of the tenant after the landlord did not appear in court. Cases with unclear outcomes have missing outcome variables or multiple occurrences of judgments or withdrawals.

A.2 Supplemental Tables

Table A1: Accuracy and Drops from Race Imputation

	% Dropped	Implied Accuracy Rate (%)
50% Rule	1.27	93.06
60% Rule	5.07	94.58
70% Rule	9.32	95.96
80% Rule	14.62	97.25
90% Rule	22.86	98.50

This table reports the share of eviction cases dropped and the implied accuracy rates for imputed tenant race under various cutoff rules for classifying tenant race. The race classification rules exclude cases that could not be classified as white or minority at the relevant level of certainty. Cases without fBISG predictions—due to missing tenant names or addresses—are dropped under all classification rules. The implied accuracy rate is the share of correct race classifications implied by aggregating the fBISG probabilities within each sample determined by the race classification rule.

Table A2: Descriptive Statistics

	All causes		Non-payment of rent only	
	White	Minority	White	Minority
<i>Claim Details</i>				
Total Claim Amount	2,657.77 (2,481.87)	2,006.25 (1,988.56)	2,791.06 (2,349.58)	2,007.67 (1,830.43)
Back Rent	2,085.19 (2,155.12)	1,575.46 (1,711.62)	2,254.15 (2,071.28)	1,619.43 (1,592.16)
Late Fees	89.27 (242.96)	69.21 (160.33)	97.60 (212.98)	74.62 (149.97)
Physical Damages	3.29 (93.86)	1.97 (99.79)	0.78 (47.40)	0.76 (72.63)
Attorney Fees	277.26 (357.46)	209.56 (267.89)	292.34 (396.13)	202.39 (250.76)
Other Fees	154.88 (757.63)	123.45 (756.14)	145.60 (667.11)	110.67 (785.49)
<i>Other Characteristics</i>				
Contract Rent	887.04 (563.18)	774.99 (28,316.92)	953.49 (573.79)	811.51 (30,635.91)
N	23,704	238,034	14,511	146,655

This table reports the means and standard deviations of variables in the analysis sample of Philadelphia eviction cases. The analysis sample includes eviction filings that could be classified as white or minority with at least 80% certainty based on tenant race imputations and have non-missing claim amounts and back rent. The total claim amount is the sum of back rent, late fees, physical damages, attorney fees, and other fees that the landlord claims the tenant owes at the time of the case filing. Contract rent is the amount of rent owed monthly. Claim amounts and back rent values are winsorized at the 99.9th percentile.

Table A3: Relationship Between Tenant Race and Non-Race Characteristics

	Female	Tenants	Previous Cases	Rent Burden (Tract-Renter Inc)	Rent Burden (Tract-Race Inc)
	(1)	(2)	(3)	(4)	(5)
Minority Tenant	0.0998*** (0.0120)	0.1523*** (0.0155)	0.1784*** (0.0655)	0.0724*** (0.0146)	0.0770*** (0.0198)
Observations	107,270	121,322	121,322	110,856	96,370
R ²	0.3058	0.3686	0.2964	0.6334	0.6271
Mean White Outcome	0.4730	1.2573	0.5376	0.3283	0.2404
Landlord × Contract Rent FE	✓	✓	✓	✓	✓

*p < 0.1, **p < 0.05, ***p < 0.01. This table reports the results from regressing observable non-race characteristics of tenants on tenant race conditional on landlord and monthly contract rent. The sample includes Philadelphia eviction filings caused only by non-payment of rent that could be classified as white or minority with at least 80% certainty based on tenant race imputations. The female outcome used in column 1 is based on the imputed gender of the first tenant named on the filing. The tenants outcome used in column 2 refers to the number of tenants listed on the filing. The previous cases outcome used in column 3 refers to the number of previous cases filed against the tenant at the same property. The rent burden outcome used in column 4 is the ratio of individual-level contract rent to median household income among renters in the census tract around the year of the filing. The rent burden outcome used in column 5 is the ratio of individual-level contract rent to median household by household head race and census tract around the year of the filing. Household head race in the eviction filing is defined based on the imputed race of the first tenant listed if the imputation is at least 80% likely. Median household income measures are sourced from 5-year ACS estimates, where the midpoint of the 5-year ACS period is matched to the case filing year. Both rent burden measures are top-coded at 1.5. All specifications exclude singleton landlord-by-contract rent observations. Contract rent fixed effects refer to the exact dollar value of contract rent. Standard errors are clustered at the landlord level.

Table A4: Total Discrimination Estimates Using Coarser Fixed Effects

	Log Back Rent			
	(1)	(2)	(3)	(4)
<i>A: Total Discrimination (No Controls)</i>				
Minority Tenant	-0.0460*** (0.0134)	-0.0476*** (0.0135)	-0.0471*** (0.0136)	-0.0496*** (0.0141)
<i>B: Direct Discrimination (With Controls)</i>				
Minority Tenant	-0.0528*** (0.0121)	-0.0561*** (0.0145)	-0.0532*** (0.0152)	-0.0543*** (0.0158)
Observations	146,628	138,290	134,140	130,432
Mean White Back Rent	1,651.12	1,614.16	1,601.38	1,595.66
Landlord FE	✓			
Contract Rent FE	✓			
Landlord × Contract Rent (\$100 bins) FE		✓		
Landlord × Contract Rent (\$50 bins) FE			✓	
Landlord × Contract Rent (\$25 bins) FE				✓

*p < 0.1, **p < 0.05, ***p < 0.01. This table reports the results from estimating Equation 14 with coarser fixed effects. The sample includes Philadelphia eviction filings caused only by non-payment of rent that could be classified as white or minority with at least 80% certainty based on tenant race imputations. All specifications exclude singleton observations for the fixed effects used. Contract rent fixed effects used in column 1 refer to the exact dollar value of contract rent. Contract rent fixed effects used in columns 2 through 4 use \$100, \$50, and \$25 dollar bins. Controls in the panel B specifications include the imputed gender of the first tenant named, fixed effects for the number of tenants listed on the filing, fixed effects for the number of previous cases filed against the tenant at the same property, fixed effects for two measures of rent burden each divided into 2-percentage point bins, and census tract fixed effects. The first measure of rent burden is the ratio of contract rent to median household income among renters in the census tract. The second measure of rent burden is the ratio of contract rent to median household income by race and census tract, where median household income by race and tract is defined based on the imputed race of the first tenant listed if the imputation is at least 80% likely. All control variables include a category for missing values. Standard errors are clustered at the landlord level.

Table A5: Total Discrimination Using Alternative Race Imputation Approaches

	Log Back Rent					
	(1)	(2)	(3)	(4)	(5)	(6)
<i>A: Alternative Measures of Race</i>						
Minority Tenant (70% Rule)	-0.0393*** (0.0144)	-0.0364*** (0.0098)				
Minority Tenant (90% Rule)			-0.0574*** (0.0222)	-0.0540*** (0.0171)		
$Pr(\text{Minority Tenant})$					-0.0362** (0.0175)	-0.0344*** (0.0119)
Observations	128,108	128,108	108,605	108,603	140,294	140,293
R ²	0.6425	0.6554	0.6479	0.6611	0.6397	0.6525
Mean White Back Rent	1,589.02	1,589.02	1,603.96	1,603.96	1,103.72	1,101.98
Landlord \times Contract Rent FE	✓	✓	✓	✓	✓	✓
Controls		✓		✓		✓
<i>B: Conditional Disparity Estimates Using Elzayn et al. (2024)</i>						
Minority Tenant	-0.2081 (0.1655)	-0.0135*** (0.0045)	-0.0722 (0.0455)	-0.0135*** (0.0046)	-0.0511** (0.0204)	-0.0140*** (0.0048)
Observations	138,608	138,608	135,410	135,410	129,405	129,405
Landlord \times Contract Rent (\$25 bin) FE	✓	✓	✓	✓	✓	✓
Estimator	Linear	Prob.	Linear	Prob.	Linear	Prob.
Min Range $Pr(\text{Minority})$	0.001	0.001	0.01	0.01	0.05	0.05

* $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$. This table reports the total discrimination estimates under alternative measures of race and using alternative imputation-based estimators of racial disparities. In panel A, columns 1-4 classify filings as white or minority based on 70% and 90% imputation cutoff rules. Columns 5-6 of panel A regress log back rent on the continuous probability that the filing is against at least one non-white tenant. Controls in the panel A specifications include the imputed gender of the first tenant named, fixed effects for the number of tenants listed on the filing, fixed effects for the number of previous cases filed against the tenant at the same property, fixed effects for two measures of rent burden each divided into 2-percentage point bins, and census tract fixed effects. The first measure of rent burden is the ratio of contract rent to median household income among renters in the census tract. The second measure of rent burden is the ratio of contract rent to median household income by race and census tract, where median household income by race and tract is defined based on the imputed race of the first tenant listed if the imputation is at least 80% likely. All control variables include a category for missing values. In panel B, the odd columns report a filing-weighted average of probabilistic racial disparities across landlord-contract rent cells. The even columns of panel B report a filing-weighted average of linear estimators across landlord-contract rent cells. The probabilistic estimator is the minority-white difference in probability-weighted average log back rent. The linear estimator is the minority coefficient from regressing log back rent on the continuous minority probability. The sample includes Philadelphia eviction cases filed only due to non-payment of rent with valid race imputation values. The sample in columns 1-4 of panel A further restrict to cases that could be classified as white or minority with 70% or 90% certainty. Specifications in panel B also require at least 3 observations per landlord-by-contract rent cell and exclude cells with minority probability ranges below the indicated cutoff. All specifications exclude singleton landlord-by-contract rent observations. Contract rent fixed effects refer to the exact dollar value of contract rent in panel A and refer to \$25 bins of contract rent in panel B. Standard errors in panel A are clustered at the landlord level, and standard errors in panel B are derived from the asymptotic distribution of the probabilistic estimator described in [Elzayn et al. \(2025\)](#).

Table A6: Robustness Checks for Alternative Landlord Identifiers

	Log Back Rent			
	(1)	(2)	(3)	(4)
<i>A: Total Discrimination (No Controls)</i>				
Minority Tenant	-0.0489*** (0.0168)	-0.0305** (0.0141)	-0.0325*** (0.0109)	-0.0285* (0.0149)
<i>B: Direct Discrimination (With Controls)</i>				
Minority Tenant	-0.0474*** (0.0162)	-0.0606*** (0.0225)	-0.0512*** (0.0153)	-0.0546*** (0.0205)
Orig. Landlord × Contract Rent FE	✓			
Alt. Landlord × Contract Rent FE		✓		
Plaintiff Address × Contract Rent FE			✓	
Plaintiff Name × Contract Rent FE				✓
Observations	120,328	106,976	96,743	94,012
Mean White Back Rent	1,584.90	1,535.61	1,523.57	1,528.33

*p < 0.1, **p < 0.05, ***p < 0.01. This table reports the results from estimating Equation 14 using alternative landlord identifiers. The sample includes Philadelphia eviction filings caused only by non-payment of rent that could be classified as white or minority with at least 80% certainty based on tenant race imputations. All specifications exclude singleton landlord-by-contract rent observations for the relevant landlord definition. Column 1 re-states the main discrimination estimates using the original landlord definition. Column 2 uses the more conservative landlord identifier that allows plaintiff names to be considered different landlords in different time periods if the plaintiff address listed on the filings has systematically changed. Column 3 uses plaintiff addresses alone to define landlords. For cases in which plaintiff addresses appear to contain attorney addresses, the plaintiff address is replaced with the plaintiff name. Column 4 uses plaintiff names alone to define landlords. Contract rent fixed effects refer to the exact dollar value of contract rent. The controls used in the panel B specifications include the imputed gender of the first tenant named, fixed effects for the number of tenants listed on the filing, fixed effects for the number of previous cases filed against the tenant(s) at the same property, two sets of rent burden fixed effects, and census tract fixed effects. The first measure of rent burden is the ratio of contract rent to median household income among renters in the census tract. The second measure of rent burden is the ratio of contract rent to median household income by race and census tract, where median household income by race and tract is defined based on the imputed race of the first tenant listed if the imputation is at least 80% likely. All control variables include a category for missing values. Standard errors are clustered at the landlord level.

Table A7: Robustness Checks for Time-Varying Filing Thresholds

	Log Back Rent			
	(1)	(2)	(3)	(4)
<i>A: Total Discrimination (No Controls)</i>				
Minority Tenant	-0.0465*** (0.0175)	-0.0468*** (0.0164)	-0.0461*** (0.0159)	-0.0642*** (0.0103)
<i>B: Direct Discrimination (With Controls)</i>				
Minority Tenant	-0.0424*** (0.0125)	-0.0438*** (0.0118)	-0.0445*** (0.0128)	-0.0430*** (0.0110)
Observations	120,328	120,328	120,328	120,316
Mean White Back Rent	1,584.90	1,584.90	1,584.90	1,584.75
Landlord \times Contract Rent FE	✓	✓	✓	✓
Filing Date Cubic	✓			
Filing Year FE		✓		
Filing Month-Year FE			✓	
Zip Code \times Filing Year FE				✓

* $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$. This table reports the results from estimating Equation 14 with controls for time trends in filing thresholds of back rent. The sample includes Philadelphia eviction filings caused only by non-payment of rent that could be classified as white or minority with at least 80% certainty based on tenant race imputations. All specifications exclude singleton observations for the fixed effects used. Contract rent fixed effects refer to the exact dollar value of contract rent. The controls used in the panel B specifications include the imputed gender of the first tenant named, fixed effects for the number of tenants listed on the filing, fixed effects for the number of previous cases filed against the tenant(s) at the same property, two sets of rent burden fixed effects, and census tract fixed effects. The first measure of rent burden is the ratio of contract rent to median household income among renters in the census tract. The second measure of rent burden is the ratio of contract rent to median household income by race and census tract, where median household income by race and tract is defined based on the imputed race of the first tenant listed if the imputation is at least 80% likely. All control variables include a category for missing values. Standard errors are clustered at the landlord level.

Table A8: Robustness Checks for Omitted Lease Violations

	Log Back Rent		Log (Back Rent + Late Fees)		Log Back Rent	
	(1)	(2)	(3)	(4)	(5)	(6)
Minority Tenant	-0.0490** (0.0231)	-0.0597*** (0.0224)	-0.0463** (0.0220)	-0.0571*** (0.0218)	-0.0478*** (0.0174)	-0.0483*** (0.0153)
Observations	85,810	85,803	85,847	85,840	126,592	126,591
R ²	0.6610	0.6770	0.6643	0.6804	0.6469	0.6593
Mean White Outcome \$ Amount	1,534.00	1,534.00	1,609.37	1,609.37	1,665.00	1,665.00
Landlord × Contract Rent FE	✓	✓	✓	✓	✓	✓
Controls		✓		✓		✓
Nonpay only – cases w/ damages or other fees	✓	✓	✓	✓		
Nonpay only + cases w/ rent-related lease breach					✓	✓

*p < 0.1, **p < 0.05, ***p < 0.01. This table reports the results from estimating Equation 14 using alternative definitions of cases caused only by non-payment of rent and incorporating other components of the claim amount. The sample in all columns includes Philadelphia eviction cases filed only due to non-payment of rent that could be classified as white or minority with at least 80% certainty based on tenant race imputations and had \$0 of physical damages claimed. Columns 1-4 additionally restrict to cases with \$0 of other fees claimed in the filing. Columns 5-6 add cases for which the description of the breach of lease only references non-payment of rent. Controls include the imputed gender of the first tenant named, fixed effects for the number of tenants listed on the filing, fixed effects for the number of previous cases filed against the tenant(s) at the same property, two sets of rent burden fixed effects, and census tract fixed effects. The first measure of rent burden is the ratio of contract rent to median household income among renters in the census tract. The second measure of rent burden is the ratio of contract rent to median household income by race and census tract, where median household income by race and tract is defined based on the imputed race of the first tenant listed if the imputation is at least 80% likely. All control variables include a category for missing values. All specifications exclude singleton landlord-by-contract rent observations. Contract rent fixed effects refer to the exact dollar value of contract rent. Standard errors are clustered at the landlord level.

Table A9: Tenure Robustness Checks

	Log Back Rent					
	(1)	(2)	(3)	(4)	(5)	(6)
Minority Tenant	-0.0544** (0.0259)	-0.0483*** (0.0133)	-0.0592** (0.0287)	-0.0524*** (0.0152)	-0.0590** (0.0290)	-0.0519*** (0.0167)
Observations	78,083	78,077	72,217	72,211	66,332	66,324
R ²	0.6449	0.6660	0.6517	0.6733	0.6582	0.6807
Mean White Back Rent	1,580.12	1,579.64	1,558.82	1,558.26	1,528.74	1,528.08
Landlord × Contract Rent (\$25 bins) × Tenure FE	✓	✓	✓	✓	✓	✓
Tenure Bin Size	12 months	12 months	6 months	6 months	3 months	3 months
Controls		✓		✓		✓

*p < 0.1, **p < 0.05, ***p < 0.01. This table reports the results from estimating Equation 14 with landlord-by-contract rent-by-tenure fixed effects. The sample includes Philadelphia eviction filings caused only by non-payment of rent that could be classified as white or minority with at least 80% certainty based on tenant race imputations. Contract rent fixed effects are \$25 bins of contract rent. Tenure fixed effects are 12 month, 6 month, or 3 month bins of tenure in the rental property, which is defined as the difference between the lease start date and the eviction filing date. Controls include the imputed gender of the first tenant named, fixed effects for the number of tenants listed on the filing, fixed effects for the number of previous cases filed against the tenant(s) at the same property, two sets of rent burden fixed effects, and census tract fixed effects. The first measure of rent burden is the ratio of contract rent to median household income among renters in the census tract. The second measure of rent burden is the ratio of contract rent to median household income by race and census tract, where median household income by race and tract is defined based on the imputed race of the first tenant listed if the imputation is at least 80% likely. All control variables include a category for missing values. All specifications exclude singleton landlord-by-contract rent-by-tenure observations. Standard errors are clustered at the landlord level.

Table A10: Share of Discriminatory Landlords

	Method 1	Method 2	Method 3
Any discrimination			
Share $\beta_j < 0$	0.8864 (0.1074)	0.7316 (0.0881)	0.4910 (0.0378)
Share $\beta_j > 0$	0.1136 (0.0931)	0.2683 (0.0826)	0.3823 (0.0318)
Substantial discrimination			
Share $\beta_j < -0.02$	0.4038 (0.2392)	0.4519 (0.1555)	0.4299 (0.0364)
Share $\beta_j > 0.02$	0.1136 (0.0907)	0.2683 (0.0450)	0.3665 (0.0227)
Median $\hat{\beta}_j$	-0.0683	-0.0233	-0.0281
N	127	520	1274
FE Type	Landlord \times Contract Rent	Landlord + Contract Rent	Landlord + Contract Rent
Sample Restriction	$N_{jyr} \geq 2$	$N_{jr} \geq 2$	$N_{jr} \geq 1$

This table reports the shares of discriminatory landlords ($\beta_j < 0$) and reverse-discriminatory landlords ($\beta_j > 0$). Method 1 uses $\hat{\beta}_j$ s controlled for landlord-by-contract rent fixed effects and estimated from sample of cases in landlord-by-contract rent cells with at least 2 non-payment filings against white tenants and at least 2 non-payment filings against minority tenants. Method 2 uses $\hat{\beta}_j$ s controlled for landlord and contract rent fixed effects separately and estimated from a sample of cases by landlords who file at least 2 non-payment filings against white tenants and at least 2 non-payment filings against minority tenants. Method 3 uses $\hat{\beta}_j$ s controlled for landlord and contract rent fixed effects separately and estimated from a sample of cases by landlords who file at least 1 non-payment filing against a white tenants and at least 1 non-payment filing against a minority tenants. Bootstrapped standard errors are reported in parentheses.

Table A11: Underlying Racial Disparities in Repossession by Landlord Type

	Writ of Possession		or Agree to Leave via JBA		or Landlord Won Judgment	
	(1)	(2)	(3)	(4)	(5)	(6)
Minority Tenant \times Minority Landlord	-0.1633** (0.0794)	-0.1006* (0.0534)	-0.2171*** (0.0812)	-0.0987* (0.0521)	-0.1168* (0.0703)	-0.0984** (0.0448)
Minority Tenant \times Non-Minority Landlord	0.0581*** (0.0217)	0.0503*** (0.0177)	0.0641*** (0.0201)	0.0548*** (0.0156)	0.0537** (0.0210)	0.0440*** (0.0161)
Observations	116,403	127,358	116,403	127,358	116,403	127,358
R ²	0.1291	0.1232	0.1376	0.1311	0.1608	0.1539
White Outcome Mean	0.4182	0.4216	0.4524	0.4593	0.6248	0.6342
Landlord \times Back Rent (\$) FE	✓		✓		✓	
Landlord \times Back Rent (Mths) FE		✓		✓		✓

*p < 0.1, **p < 0.05, ***p < 0.01. The sample includes Philadelphia eviction filings caused only by non-payment of rent that could be classified as white or minority with at least 80% certainty based on tenant race imputations. Back rent fixed effects in dollars are bins defined at \$50 increments. Back rent fixed effects in months owed are bins of back rent divided by contract rent defined at increments of 0.25 months. The outcome variable in columns 1 and 2 is an indicator for whether a landlord filed for a writ of possession. The outcome variable in columns 3 and 4 additionally includes cases where the tenant agreed to vacate the property as part of a judgment by agreement (JBA). The outcome variable in columns 5 and 6 additionally includes any case in which the landlord won a default judgment or won a contested judgment (decided by the judge) regardless of filing for a writ of possession. All specifications exclude singleton landlord-by-back rent observations. Standard errors are clustered at the landlord level.

Appendix B Data Construction

B.1 Landlord Linking Methodology

I construct landlord identifiers using plaintiff names and plaintiff addresses listed in Philadelphia eviction records. This approach is intended to account for inconsistent spellings and abbreviations and the common practice by corporate landlords of using multiple property-specific company names. I implement the following procedure to generate landlord identifiers. For a given plaintiff name, I observe all plaintiff address variations associated with that plaintiff name. For each of those plaintiff address variations, I then observe all plaintiff name variations associated with that plaintiff address. I repeat this process until no additional plaintiff name or address variations can be found. Then, I assign this linked set of landlord aliases a unique landlord identifier. I repeat this process for the entire sample of Philadelphia eviction cases.⁴³

In some cases, the plaintiff address field appears to contain the address of the plaintiff's attorney instead of the landlord's address. Because of this, the linking algorithm may mistakenly group filings with shared attorneys under a single landlord. I prevent this by manually identifying the law firm addresses of all attorneys that represent the plaintiff in at least 50 filings and do not input any plaintiff addresses matching these attorney addresses into the linking algorithm.⁴⁴ All filings with these attorney addresses are assigned landlord identifiers based on linkages derived from the plaintiff names alone. In the end, from 65,438 unique plaintiff names and 52,389 unique plaintiff addresses, I generate 38,347 unique landlord identifiers. Of these, 36,847 unique landlords filed residential eviction cases during the analysis period.

In section 6.3.3, I adopt more conservative landlord definitions. One alternative approach allows a plaintiff name to be considered a different landlord in different time periods, depending on whether the plaintiff address associated with the plaintiff name has systematically changed—possibly due to a change in the property manager's identity. For a given plaintiff name filing cases in a given month, I observe all plaintiff address variations associated with that plaintiff name in that month. For each of these plaintiff address variations, I observe all other plaintiff names associated with that plaintiff address in that same month. This process yields a set of landlord aliases used in a given month. I then construct landlord identifiers by grouping all aliases ever used by a plaintiff address over time and grouping plaintiff addresses that are ever aliases of one another within any filing month. From this more conservative

⁴³I construct landlord identifiers prior to dropping non-residential eviction filings in order to maximize linkages between plaintiff name and address aliases.

⁴⁴In total, I remove plaintiff addresses for less than 8% of filings for this reason.

approach, I generate 52,560 unique landlord identifiers among residential eviction cases.⁴⁵ Other alternative approaches to defining landlords in section 6.3.3 use simply the plaintiff name or plaintiff address as landlord identifiers.

Table B1 reports statistics about landlords as defined under these approaches. Under all landlord definitions, the average landlord files few (5-8) eviction cases during the analysis period, and a small number of landlords account for a large fraction of eviction filings. Since the discrimination estimates in the paper rely on variation within landlord and contract rent, the results reflect the behavior of a small share of landlords, but these landlords are important since they account for a substantial share of eviction cases.

Table B1: Landlord Summary Statistics

	Main Approach	Month Approach	Plaintiff Address	Plaintiff Name
	(1)	(2)	(3)	(4)
N landlords	36,847	52,560	57,397	62,846
% of filings by top 11 landlords	43.62	26.10	14.90	15.27
% of filings by top 50 landlords	52.79	37.46	25.07	21.43
% of filings by top 100 landlords	57.74	43.81	31.06	26.19
Mean filings	8	6	5	5
N landlords w/ required variation	486	594	727	726
% of filings by landlords w/ any req. variation	59.70	47.64	37.68	31.01

This table reports statistics about landlord eviction filing volumes under various landlord definitions. Eviction filings only include residential eviction cases. A landlord has required variation for estimating discrimination if they ever filed eviction cases (caused only by non-payment of rent) against at least one white tenant and at least one minority tenant with the same contract rent. Column 1 reports statistics for landlords defined by the full linkage of plaintiff name and plaintiff address aliases. Column 2 reports statistics for landlords defined by the alternative approach of first linking plaintiff name and address aliases within a filing month and second grouping aliases with any shared plaintiff address across time. Column 3 reports statistics for landlords defined by unique plaintiff addresses. For cases in which the plaintiff address appears to contain an attorney address, the plaintiff address is replaced with the plaintiff name. Column 4 reports statistics for landlords defined by unique plaintiff names.

B.2 Race Imputation Methodology

I impute tenant race based on tenants' first names, last names, and census blocks of rental properties using the fBISG methodology outlined in Imai et al. (2022). Incorporating first names and granular geolocation measures greatly improves the accuracy of race imputations compared using last names alone (Voicu 2018; Imai and Khanna 2016).

⁴⁵This approach likely fails to group filings by landlords with small variations in the plaintiff address listed over time or landlords that changed their home or office address over the analysis period.

The standard BISG approach estimates the conditional probability that an individual i belongs to a racial group given the individual’s name and geolocation. However, this approach can have substantial measurement error problems when using granular geolocation measures. Since the prior distribution $Pr(Race_i|Geolocation_i)$ comes from decennial census counts, some census blocks may have zero counts of individuals in certain racial groups which imposes a 0% probability of belonging to a racial group even if a name is distinctly associated with that racial group. The prior distribution is also decreasing in accuracy in time since the decennial census. Both of these census data problems are more significant for rapidly growing minority groups.

I utilize the improved fBISG method from [Imai et al. \(2022\)](#) which allows for gaining the accuracy from granular geolocation information while addressing these census data problems. This method uses a measurement error model for the observed census counts to ensure that zero counts do not necessarily imply 0% probability of belonging to that racial group. This approach has been shown to substantially improve accuracy compared to BISG imputations, especially for minority groups ([Imai et al. 2022](#)). I use an fBISG model based on the tenant’s first name, last name, and census block to generate the conditional posterior probability that tenant i belongs to each potential racial group (White, Black, Hispanic, Asian, other). These predictions incorporate 2010 decennial census counts by race and census block, the 2010 census surname list which reports the racial distribution of surnames appearing at least 100 times, and the racial distribution of first names from voter registration data compiled by [Imai et al. \(2022\)](#). See [Imai et al. \(2022\)](#) for more detail on the fBISG race imputation methodology.

Appendix C Proof of Marginal Outcome Test

C.1 Proof of Proposition 1

An eviction filing against tenant i at the margin ($X_i(t) = \tau_i$) satisfies:

$$\left(\tilde{p}(\tau_i, Y_i^0, Z_i, R_i) + b^p(R_i)\right) \left(\tilde{v}(Y_i^0, Z_i) + b^v(R_i)\right) = c \quad (19)$$

Since $\tilde{p}(x, y, z, r) = \mathbb{E}[Y_i^* | X_i(t) = x, Y_i^0 = y, Z_i = z, R_i = r]$, this is equivalent to:

$$\mathbb{E}[Y_i^* | X_i(t) = \tau_i, Y_i^0 = y, Z_i = z, R_i = r] = \frac{c}{\tilde{v}(y, z) + b^v(r)} - b^p(r) \quad (20)$$

For $r \in \{white, minority\}$, it thus follows from the law of iterated expectations that:

$$\begin{aligned} & \mathbb{E}[Y_i^* | X_i(t) = \tau_i, Y_i^0 = y, Z_i = z, R_i = r] \\ = & \mathbb{E}\left[Y_i^* \mid \mathbb{E}[Y_i^* | X_i(t) = \tau_i, Y_i^0 = y, Z_i = z, R_i = r]\right] = \frac{c}{\tilde{v}(y, z) + b^v(r)} - b^p(r), Y_i^0 = y, Z_i = z, R_i = r \\ = & \mathbb{E}\left[\mathbb{E}[Y_i^* | X_i(t) = \tau_i, Y_i^0 = y, Z_i = z, R_i = r]\right] \\ & \mathbb{E}[Y_i^* | X_i(t) = \tau_i, Y_i^0 = y, Z_i = z, R_i = r] = \frac{c}{\tilde{v}(y, z) + b^v(r)} - b^p(r), Y_i^0 = y, Z_i = z, R_i = r \\ = & \frac{c}{\tilde{v}(y, z) + b^v(r)} - b^p(r) \end{aligned} \quad (21)$$

Under no racial bias ($b^v(white) = b^v(minority)$ and $b^p(white) = b^p(minority)$):

$$\mathbb{E}[Y_i^* | X_i(t) = \tau_i, Y_i^0 = y, Z_i = z, R_i = minority] = \mathbb{E}[Y_i^* | X_i(t) = \tau_i, Y_i^0 = y, Z_i = z, R_i = white] \quad (22)$$

C.2 Proof of Proposition 2

This proof is by contradiction. Suppose a landlord exhibits no racial bias (accurate statistical discrimination). Then, $b^v(white) = b^v(minority) = b^v$ and $b^p(white) = b^p(minority) = b^p$. It follows from Equation 21 by the law of iterated expectations that:

$$\mathbb{E}[Y_i^* | X_i(t) = \tau_i, Y_i^0 = y, Z_i = z, R_i = r] = \frac{c}{\tilde{v}(y, z) + b^v} - b^p \quad (23)$$

so

$$\begin{aligned} & \mathbb{E}[Y_i^* | X_i(t) = \tau_i, Y_i^0 = y, Z_i = z, R_i = \textit{minority}] - \mathbb{E}[Y_i^* | X_i(t) = \tau_i, Y_i^0 = y, Z_i = z, R_i = \textit{white}] \\ &= \frac{c}{\tilde{v}(y, z) + b^v} - b^p - \frac{c}{\tilde{v}(y, z) + b^v} - b^p \\ &= 0 \end{aligned}$$

(24)

Appendix D Bias from Discrete Back Rent

For my test for total racial discrimination to be valid, I assume that the amount by which back rent may exceed the threshold at the time of filing is the same for white and minority tenants who have the same landlord and contract rent. If this assumption is violated, my estimate includes bias from differential discrete accrual of back rent. Considering all tenants with the same landlord, this bias term for contract rent y is:

$$\mathbb{E}[X_i(t^*) - \tau_i | Y_i^0 = y, R_i = \textit{minority}] - \mathbb{E}[X_i(t^*) - \tau_i | Y_i^0 = y, R_i = \textit{white}] \quad (25)$$

Define the last discrete change in back rent prior to the eviction filing as:

$$\Delta_i = X_i(t^*) - X_i(t^* - 1) \quad (26)$$

Note that $X_i(t^* - 1)$ must be less than the threshold τ_i , implying that $X_i(t^*) - \tau_i \in [0, \Delta_i)$. Assume for simplicity that the tenant distribution of $X_i(t^*) - \tau_i$ has uniform density between 0 and Δ_i .

Suppose that for some share of tenants, back rent accrues at discrete increments equal to contract rent ($\Delta_i = Y_i^0$). For the remaining tenants, assume the most extreme alternative where back rent accrues continuously ($\Delta_i = 0$). Let α^m be the share of minority tenants with $\Delta_i = Y_i^0$ and α^w be the share of white tenants with $\Delta_i = Y_i^0$. Then, the bias term for contract rent y is equal to:

$$(\alpha^m - \alpha^w) \frac{y}{2} \quad (27)$$

I next turn to the data to investigate whether it is a valid concern that any racial differences in the last discrete change in back rent prior to the filing could explain my estimate of total racial discrimination. Given back rent owed at the time of the filing and contract rent, I observe whether tenants owe back rent that is an exact multiple of contract rent at the time of an eviction filing.⁴⁶ It is a reasonable assumption that tenants owing exact multiples of contract rent at the time of the filing have $\Delta_i = Y_i^0$. For the remaining tenants in the data that owe non-integer months of back rent at the time of the case filing, it is likely that $\Delta_i < Y_i^0$ and in the most extreme case, $\Delta_i = 0$.

My estimate of total racial discrimination is $-0.1389y$ (Table 2). For bias from differential discrete accrual of back rent to fully explain this result, it must be the case that $\alpha^m - \alpha^w = -0.2778$. In other words, among tenants with the same landlord and contract rent, the share of minority tenants owing integer months of back rent must be 27.78 percentage points lower

⁴⁶Note that back rent is observed separately from any late fees also accrued.

than the share of white tenants owing integer months of back rent. In the data, minority tenants are only 2.32 percentage points less likely to owe integer months of back rent than white tenants with the same landlord and contract rent. This could only generate bias equal to $-0.0116y$, which is just 8.35% of my true estimate.