Enforcement Risk and Screening in Rental Markets: Evidence from Eviction Moratoria *

Alina Arefeva
University of North Dakota &
University of Wisconsin – Madison

Qihui Hu
Carnegie Mellon University

Christopher Timmins
University of Wisconsin – Madison

Abstract

We study how contract enforcement shapes screening by intermediaries in the U.S. rental market, where eviction enforces leases and tenant selection is delegated to property managers. We combine a hand-collected record of state eviction moratoria with a nationwide field experiment of more than 25,000 inquiries. Exploiting staggered moratorium expirations, we show that property managers discriminated more while eviction was unenforceable, redirecting responses from African American toward white applicants without changing their overall responsiveness. The reallocation altered who ultimately moved in and is consistent with statistical discrimination: a benefit–cost calculation shows that the screening was privately profitable despite being illegal.

Keywords: contract enforcement, screening, financial intermediation, real estate markets, discrimination

JEL Codes: G18, G21, G32, D83, R31, J15

*Arefeva: alina.arefeva@und.edu, Hu: qihuih@andrew.cmu.edu, Timmins: ctimmins@wisc.edu (corresponding coauthor). This paper was previously circulated as "Discrimination During Eviction Moratoria." We thank Peter Christensen, Ignacio Sarmiento-Barbieri, and Christopher Timmins for sharing the experimental data from the correspondence study on housing discrimination, Coulson, Le, Ortego-Marti, and Shen (2025) for sharing their tenant protection index data, and Kay Jowers for sharing her eviction moratoria data. We are also grateful for feedback from Anthony DeFusco and our discussants Nitzan Tzur-Ilan, Carlos Fernando Avenancio-León, Sofie R Waltl, Matthew Freedman, Robert Collinson, Hanchen Jiang, Andrew Spieler, Boaz Abramson, and Cristian Badarinza. The paper has benefited from the feedback of seminar and conference participants at the UW-Madison, UCLA, University of Illinois Urbana-Champaign, Indiana University, University of North Georgia, College of Charleston, University of North Dakota, University of Nevada, Old Dominion University, the Conference on Auctions, Firm Behavior, and Policy at the University of Oklahoma, the NBER Summer Institute Real Estate Workshop, the Real Estate Finance and Investment Symposium at the University of Cambridge, the UEA Conference, the MFA Conference, the AREUEA National Conference, the CREDA Conference, FMA, the New Economic School, the University of Hawaii, REALPAC, USC, UC-Riverside, pre-WFA, and WFA. Ailey Fang and Lucas Marron provided excellent research assistance. The RCT was registered as AEARCTR-0005338 (Sarmiento-Barbieri 2020). The UW-Madison Institutional Review Board determined that this research is exempt from IRB review: submission ID number: 2024-0560. The authors do not have conflicts of interest to disclose.

1 Introduction

Screening counterparties is a core function of intermediaries. In the canonical view, intermediaries manage counterparty risk on two margins: they screen before the contract is signed and they enforce after it is breached (Stiglitz and Weiss, 1981; Diamond, 1984). The margins are substitutes, so when the ex post remedy weakens, the burden of managing default risk shifts to the ex ante screen. Testing this substitution directly has proven difficult for two reasons. Changes in enforceability rarely arrive independently of economic conditions. And the screening decision itself is rarely observable: empirical work sees equilibrium contract terms and ex post performance, such as interest rates and defaults in credit markets, rents and evictions in rental housing, while the choice of whom to transact with is made before any contract exists and is rarely observable in the data.

This paper observes the screening decisions directly and measures how they respond to a large, temporary, and staggered change in contract enforcement. The setting is the U.S. rental market, where the lease is the contract, eviction is its enforcement technology, and screening is delegated to property managers acting on behalf of property owners. In 2020, state eviction moratoria suspended enforcement; their staggered expirations restored it, state by state, over the spring and summer. Because tenant selection operates through responses to rental inquiries, a correspondence experiment that sends fictitious inquiries and records who receives a reply captures screening decisions.

Our empirical approach combines three elements. First, we construct a hand-collected dataset that traces each state moratorium to its operative legal instrument (executive order, court order, or statute) and codes its effective start and end dates, giving precise variation in the credibility of enforcement. Second, we use the largest nationwide correspondence experiment in rental housing, more than 25,000 inquiries sent to listings in the fifty largest U.S. metropolitan areas (Christensen et al., 2021), which lets us observe screening decisions at the inquiry stage. Third, we link the experimental listings to residential address histories, connecting inquiry-stage screening to the tenants who subsequently move in. We estimate treatment effects with the staggered difference-in-differences estimator of Callaway and Sant'Anna (2021), applied to state-by-day discrimination measures constructed from the experiment.

Our central finding is that suspending enforcement intensified discriminatory screening. During the moratoria, African American applicants received responses to 54.5 percent of inquiries, against 59.8 percent for observably identical white applicants. When moratoria expired, the log relative response ratio for African American applicants rose by 0.120, which is a 12.8 percent increase that closes the gap, with the confidence interval bracketing equal

response rates. The effect appears at the legal expiration date, shows no differential pre-trends, and is robust across event windows, control groups, alternative staggered estimators, and other robustness checks.

The moratoria changed the composition of screening, not its level. Decomposing response probabilities into a common level and race-specific components, we find the overall responsiveness of property managers flat through expiration: they answered as many inquiries during the moratoria as after. What changed is who received the answers: responses shifted toward white applicants while eviction was unenforceable and shifted back when it was restored. Weakened enforcement therefore reallocated access to housing rather than rationing it, concentrating the incidence of the policy on the applicants perceived as higher default risks.

Why did an illegal practice persist for the length of the moratoria? Following Becker (1968), a property manager deciding whether to screen on race weighs the expected benefit of the practice against its expected cost. The benefit was risk management: steering responses toward the applicants perceived as safer while eviction was unenforceable was worth roughly $26 per property manager per moratorium. The cost was legal exposure under the Fair Housing Act, and it was small in expectation: race-based fair-housing complaints were filed at 0.0113 percent per renter household-year, an order of magnitude below the 0.121 percent detection probability at which the statutory penalty would deter, putting the expected penalty at $2.43. With the perceived benefit exceeding the cost, property managers shifted responses away from African American applicants for as long as the moratoria lasted, and linked address histories show that the shift carried through to occupancy: non-responses translated into systematically different move-in patterns.

The mechanism is statistical rather than taste-based discrimination. The screening response is, if anything, stronger in high-rent markets, where the option value of a vacant unit is largest; it does not vary with measured racial animus (Stephens-Davidowitz, 2013); and it is strongest in states whose baseline tenant protections (Coulson et al., 2025) had already narrowed the scope for formal eviction, so that moratoria removed the last remaining remedy. The estimates are also unchanged during the federal CARES Act moratorium, which imposed a statutory restriction without an enforcement mechanism in state courts: screening responds to the credibility of enforcement, not to the statute.

The staggered timing we exploit is not random, but the threats it poses are testable, and we test them. Event studies show no differential pre-trends. Expiration timing does not respond to, or coincide with, changes in state economic conditions measured in the Census Household Pulse Survey (employment, income, and rent and mortgage delinquency are flat through expiration) and is not predicted by contemporaneous COVID-19 dynamics or

predetermined socioeconomic characteristics. What does predict it is the political composition of the state legislature, fixed months before the pandemic; predetermined institutional preferences can shift levels, which the design differences out, but not trends at expiration. The results also survive excluding states that ended moratoria at their governor's discretion, restricting identification to expirations fixed in advance by court calendars, statutes, and the stated terms of executive orders.

Our paper contributes to four literatures. The first studies contract enforcement and intermediary incentives, showing how limited enforceability shapes screening and monitoring when agents act on behalf of principals (Arellano et al., 2012; Lerner and Schoar, 2005; Quintin, 2008). We bring this logic to rental housing, where lease enforcement operates through eviction and screening is delegated to property managers. Relative to this literature, our setting provides a shock to enforceability, direct observation of the screening decision at the micro level, and a trace from screening to realized allocation.

Second, we contribute to work in mortgage and housing finance on how enforcement institutions and intermediary exposure shape ex ante screening. One strand exploits differences between judicial and non-judicial foreclosure regimes to show that weaker enforcement affects credit supply and outcomes (Pence, 2006; Dagher and Sun, 2016). Another shows that exposure to downside risk disciplines origination: securitization that removes skin in the game weakens screening (Keys et al., 2010), and the ability to sell or delay the disposition of originated loans shapes it (Adelino et al., 2025). We study the same force, the intermediary's exposure to losses from a bad match, but on a different margin. Existing work varies exposure through contract design at origination or risk transfer after origination, and infers screening from loan terms and subsequent performance; we vary the credibility of enforcement at the moment the screening decision is made and observe the decision itself.

Third, we contribute to the literature on discrimination in housing and credit markets (Yinger, 1995; Galster and Godfrey, 2005; Hanson and Hawley, 2011; Ghent et al., 2014; Christensen et al., 2021; Ambrose et al., 2020; Bartlett et al., 2022; Frame et al., 2023; Fuster et al., 2022). This work documents unequal treatment; it provides limited evidence on how the legal environment shapes its intensity. We show that discrimination intensifies when contract enforcement weakens, trace it from inquiry responses to occupancy, and price the private return to the illegal screening. This connects fair-housing compliance to the economics of crime.

Finally, we contribute to the literature on eviction policy and landlord behavior. Existing work documents the prevalence and consequences of eviction for tenants (Desmond, 2016; Collinson et al., 2023), the equilibrium effects of eviction restrictions on rents, defaults, and welfare (Abramson, 2025; Corbae et al., 2024; Coulson et al., 2025), and the landlord's ex

post eviction decision (Humphries et al., 2025). This literature evaluates eviction policy through rents and ex post outcomes because the screening stage has not been observable. We observe it, and we find that the incidence of moratoria arrives before any tenancy begins: while eviction was unenforceable, access to vacant units was reallocated across groups even though total access was unchanged, which is a margin invisible in rents. Welfare evaluations of eviction policy must therefore incorporate delegated screening and market access alongside the protection of incumbent tenants.

2 Theoretical Framework

We develop a dynamic screening model in which a property manager chooses whether to respond to rental inquiries when tenant default is possible and eviction enforcement may be temporarily suspended. The model allows for both group-specific beliefs about default risk (statistical discrimination) and group-specific utility costs (taste-based discrimination), and yields a simple testable prediction: by suspending enforcement, moratoria strengthen the influence of discriminatory priors on tenant selection decisions.

2.1 Setup

Consider a forward-looking property manager who discounts future payoffs at rate $\beta < 1$ and decides whether to respond to rental inquiries from two types, $i \in \{M, W\}$, of applicants: a minority applicant ($i = M$) and a white applicant ($i = W$). Applicants accept any lease offer, after which they occupy the unit as tenants. Each period, a tenant of type i pays rent $R > 0$ with probability π_i and defaults with probability $1 - \pi_i$. The parameter π_i captures the landlord's subjective belief about tenant reliability and may differ across groups. We interpret $\pi_M < \pi_W$ as evidence of *statistical discrimination*.

In addition to pecuniary payoffs, the property manager may incur a non-pecuniary flow cost κ_i from leasing to type-i applicants, reflecting taste-based preferences. We normalize $\kappa_W = 0$ and let $\kappa_M \equiv \kappa \geq 0$. A positive value of κ captures *taste-based discrimination*.

Tenants who default are evicted unless a moratorium is in place. If the property becomes vacant, the manager can list it for lease and start responding to inquiries about renting this unit. Let ψ denote an idiosyncratic shock to the relative payoff from responding to a minority rather than a white applicant. We assume $\mathbb{E}[\psi] = 0$, with cumulative distribution function $F(\cdot)$ and density $f(\cdot)$. We impose the technical condition $\lim_{\psi \to \psi_{\min}} \psi F(\psi) = 0$ to ensure well-behaved limits. Importantly, shock ψ captures idiosyncratic variation in search costs and is not itself a source of discrimination.

2.2 Property Manager's Problem

Let u_i denote the expected continuation value of leasing to an applicant of type i when eviction is enforceable

$$u_i = \pi_i(R + \beta u_i) + (1 - \pi_i)\beta V - \kappa_i, \tag{1}$$

where V denotes the option value of keeping the unit vacant and continuing to search.

Upon observing inquiries from applicants, the manager compares the payoff from responding to a minority applicant with the payoff from responding to a white applicant. The manager responds to the minority applicant whenever $u_M - \psi > u_W$, so the probability of response is:

$$P_M^{\text{Response}} = \mathbb{P}(u_M - \psi > u_W) = F(u_M - u_W). \tag{2}$$

Defining $\Delta u \equiv u_M - u_W$, the option value of a vacant unit satisfies:

$$V = \mathbb{E}[\max\{u_M - \psi, u_W\}] = u_W + \int_{\psi_{\min}}^{\Delta u} F(\psi)\, d\psi. \tag{3}$$

An equilibrium is a fixed point $\{u_M^*, u_W^*, V^*\}$ that solves the above equations.

2.3 Eviction Moratorium

We model an eviction moratorium as a temporary, one-period suspension of eviction following tenant default. During the moratorium period, the manager continues to incur the utility cost κ_i associated with tenant i and cannot re-list the unit. The continuation value under a moratorium becomes:

$$u_i = \pi_i(R + \beta u_i) + (1 - \pi_i)(-\beta\kappa_i + \beta^2 V) - \kappa_i. \tag{4}$$

Relative to the previous expression, the moratorium increases the expected cost of default by introducing both a time delay and an additional utility penalty. These effects interact directly with the manager's beliefs (π_i) and preferences (κ_i) regarding different applicant types.

2.4 Discrimination Effects of the Moratorium

Let $\Delta u^* \equiv u_M^* - u_W^*$ denote the cross-type value difference before the moratorium, and let $\Delta\bar{u} \equiv \bar{u}_M - \bar{u}_W$ denote the same under the moratorium. The difference in this difference is:

$$\Delta^2 u = \Delta\bar{u} - \Delta u^*. \tag{5}$$

A negative value of $\Delta^2 u$ implies that the moratorium exacerbates pre-existing disparities in applicant response rates.

Taste-Based Discrimination. Suppose $\pi_M = \pi_W = \pi$ and $\kappa > 0$. Then:

$$\Delta^2 u = -\frac{\beta(1-\pi)}{1-\beta\pi}\kappa < 0. \tag{6}$$

The moratorium amplifies the effect of taste-based preferences by prolonging the period over which the utility loss from leasing to minority tenants is incurred.

Statistical Discrimination. Suppose $\pi_W > \pi_M$ and $\kappa = 0$. Then:

$$\Delta^2 u = -\frac{\beta(1-\beta)(\pi_W - \pi_M)}{1-\beta\pi_M - \beta(\pi_W - \pi_M)F(\Delta u^*)}\bar{V} < 0. \tag{7}$$

Here, the moratorium increases expected losses from applicants perceived as higher risk, with the effect magnified in high-rent markets where the option value of the unit $\bar{V}$ is larger.

2.5 Empirical Implications

The model predicts that eviction moratoria intensify discriminatory selection, decreasing the likelihood that minority applicants receive responses. The mechanisms differ by discrimination type. Under taste-based discrimination, response gaps widen uniformly across markets. Under statistical discrimination, the impact is more pronounced in high-rent markets, where the opportunity cost of delayed re-leasing is higher. These heterogeneous implications motivate our empirical test on the source of discrimination in Section 8.2.

3 Data

Correspondence experiment. Our main dataset comes from the large-scale randomized correspondence study conducted by Christensen et al. (2021), which examines racial dis-

crimination in online rental markets in the United States. The experiment was implemented between February 6 and July 31, 2020 using an automated software bot that sent rental inquiries to newly posted listings across the fifty largest U.S. metropolitan areas. Appendix Tables A3 and A4 list the metropolitan areas and states in the experiment. Each day the bot targeted six listings per metro area: three located in downtown zip codes (defined as the five percent closest to the city center) and three in suburban zip codes. One day after a listing appeared, the bot submitted inquiries under fictitious applicant identities in randomized order, ensuring that no property manager received more than one inquiry per day.

Table 1: First and Last Names of Identities Used in the Correspondence Study

African American	**Hispanic**	**White**
Nia Harris	Isabella Lopez	Aubrey Murphy
Jalen Jackson	Jorge Rodriguez	Caleb Peterson
Ebony James	Mariana Morales	Erica Cox
Lamar Williams	Pedro Sanchez	Charlie Myers
Shanice Thomas	Jimena Ramirez	Leslie Wood
DaQuan Robinson	Luis Torres	Ronnie Miller

Applicant identities were constructed using eighteen first-last name pairs designed to signal African American, Hispanic, or white racial identities (see Table 1). Identities were additionally stratified by gender and maternal education background to reduce potential correlations with unobservable socioeconomic characteristics (Guryan and Charles, 2013; Fryer Jr and Levitt, 2004). Each fictitious applicant was assigned a unique email address and a New York City-area phone number, and all inquiries used the same standardized platform-generated message requesting to schedule a viewing. Responses were coded as one if the property manager confirmed availability within seven days and zero otherwise; approximately 80% of responses occurred within the first 24 hours.

The unit of observation is an individual inquiry sent to a property listing. The final dataset contains 25,428 inquiry-listing interactions. For each inquiry, the data report the response indicator, the race and gender of the fictitious applicant, the maternal education category, the order of the inquiry within the listing, the county and state FIPS codes, the CBSA, and the date the inquiry was sent. This dataset has two important limitations. First, information about the property manager, including identity and race, is not available because the vast majority of listings are posted by limited liability companies (LLCs). Second, the dataset also does not report listing characteristics such as rent levels or property type (e.g., single-family versus multi-family). These limitations restrict our ability to conduct important

heterogeneity analyses by the race and type of the property manager and the property type. However, we study whether eviction moratoria change the composition of rental listings using CoreLogic data, and do not find any effects.

In the empirical analysis, the response indicator serves as the main outcome variable, while applicant race and gender indicators (and their interactions) are the key explanatory variables. We construct time variables based on the inquiry date and drop observations from Virginia after June 8, 2020, when the state introduced a second eviction moratorium.

Eviction moratoria. We construct a hand-collected dataset that records the start and end dates of state-level eviction moratoria during the spring and summer of 2020. Each policy action is linked to its legal or administrative source, including executive orders, court rulings, legislative acts, and regulatory directives. Building on the database compiled by Benfer et al. (2023) and Kay Jowers, we independently verify all entries using primary sources and reconcile discrepancies in policy timing or classification.

Our analysis focuses on states that implemented a moratorium during the sample period. Appendix Figures A1a and A1b report the geographic distribution of moratoria and the timing of enactment. Section 7.2 shows that our results are robust to alternative control groups that include these never-treated states.

Census Household Pulse Survey. We obtain high-frequency data on household economic conditions from the U.S. Census Bureau's Household Pulse Survey, available at `https://www.census.gov/programs-surveys/household-pulse-survey.html`, an online survey launched in April 2020 to measure the economic and social effects of the COVID-19 pandemic in near real time. We use the public-use microdata files for the first forty collection waves, covering late April 2020 through December 2021; the survey was conducted weekly through July 2020 and in two-week waves thereafter. We aggregate the microdata to a state-by-week panel and construct four measures: the employment share, defined as the person-weighted share of prime-age (25–54) adults who worked in the seven days preceding the interview; household income, the household-weighted mean of the survey's eight household-income bands, converted to dollars at band midpoints; the rent delinquency rate, the household-weighted share of renter households not caught up on rent payments; and the mortgage delinquency rate, the household-weighted share of mortgage-holding homeowners not current on their mortgage. We use these series to test whether the timing of moratorium expiration followed, or coincided with, changes in time-varying economic conditions at the state level in Section 5.1.

COVID-19 cases and deaths. We obtain daily state-level COVID-19 case and death counts from the *New York Times* COVID-19 dataset, available at `https://github.com/nytimes/covid-19-data`. The data provide cumulative counts of confirmed cases and deaths by state and day. We downloaded the dataset on March 3, 2024, and set case and death counts to zero for state–day cells prior to the first reported case. We convert case counts to COVID-19 cases per 100,000 residents using state population totals.

We use COVID-19 cases per 100,000 population for two purposes. First, we use this measure to assess potential endogeneity in the timing of eviction moratoria. Second, we test whether the expiration of eviction moratoria affects COVID-19 case rates, verifying that treatment is not associated with changes in pandemic dynamics.

Rental listings data. We use rental listing data from the CoreLogic Multiple Listing Service (MLS). The raw dataset contains all MLS listings from January 2018 through August 2020. We restrict the sample to rental listings and apply a series of cleaning procedures to remove non-residential properties, foreclosures, short sales, duplicate listings, and observations with missing or invalid dates; the full sample construction is reported in Appendix Table A6. The final cleaned dataset contains approximately 1.06 million rental listings. The data report detailed unit characteristics, including square footage, number of bedrooms and bathrooms, and property age, as well as listing and lease timing. We use these data to test whether the expiration of eviction moratoria affected the number of rental listings or the observable characteristics of the units.

Racial animus index. We measure racial animus using the Google search-based index developed by Stephens-Davidowitz (2013), which captures the relative frequency of racially charged search terms across U.S. Designated Market Areas (DMAs) between 2004 and 2007. The DMA-level index is mapped to Census block groups and merged with the experimental addresses to examine heterogeneity in discriminatory responses across locations with higher versus lower levels of racial animus; the precise geographic crosswalk used in this merge cannot be disclosed due to Institutional Review Board confidentiality requirements.

Tenant protection index. Coulson, Le, Ortego-Marti, and Shen (2025) compile a state-year dataset of legal provisions governing landlord-tenant relationships in the United States for the period 1997-2016. Using these data, they construct a Tenant Rights Index (TRI) by extracting the first principal component from twelve provisions capturing the strength of tenant protections. We use the most recent available value of the index (2016) to measure cross-state differences in tenant protection in our heterogeneity analysis in Section 8.4.

Political Affiliation. We measure state-level partisan composition using 2019 data from the National Conference of State Legislatures (NCSL).[1] For each state, we compute the Republican seat share in the lower chamber, defined as the ratio of Republican seats to total seats. Legislative composition is measured in 2019, prior to our sample period, and is used to construct heterogeneity splits based on the median Republican share.

Zillow Observed Rent Index (ZORI). We obtain rental market data from the Zillow Observed Rent Index (ZORI), provided by Zillow Research. The ZORI series reports smoothed, seasonally adjusted rent levels for all homes plus multifamily units and is available at the CBSA-month level. We measure local rental market conditions using the latest ZORI value prior to the start of our sample period, February 2020, and use this rent level to classify markets into above- and below-median rent areas for heterogeneity analysis. To merge the data with the experimental sample, we harmonize CBSA names between Zillow's naming convention and the Office of Management and Budget definitions using a crosswalk.

Rental Vacancy Rates. We obtain state-level rental vacancy rates from the U.S. Census Bureau's Housing Vacancy Survey (HVS). The HVS reports the percentage of rental housing units that are vacant in each state and quarter. We measure local rental market slack using the rental vacancy rate in 2019:Q4, the last observation prior to the start of our sample period. We use this measure in heterogeneity analyses.

American Community Survey (ACS). We obtain state-level socioeconomic characteristics from the 2015–2019 American Community Survey (ACS) five-year estimates accessed through Social Explorer. The data include total population, population density, median household income, the share of residents aged 65 and older, racial and ethnic composition, the share of renters, education shares, the share of residents living in group quarters, the share of essential workers, the uninsured share, average rent, and commuting mode shares, including public transportation, carpooling, driving alone, motorcycle, bicycle, walking, and working from home. We construct all shares as ratios of the relevant subgroup count to its corresponding ACS group total. For regression use, we transform median household income using the natural logarithm and rescale population density by dividing by 1,000. These variables are used to test for potential endogeneity in the timing of moratorium expiration and to conduct heterogeneity analyses across states with different demographic and socioeconomic characteristics.

[1] National Conference of State Legislatures, State Partisan Composition, available at `https://www.ncsl.org/about-state-legislatures/state-partisan-composition`.

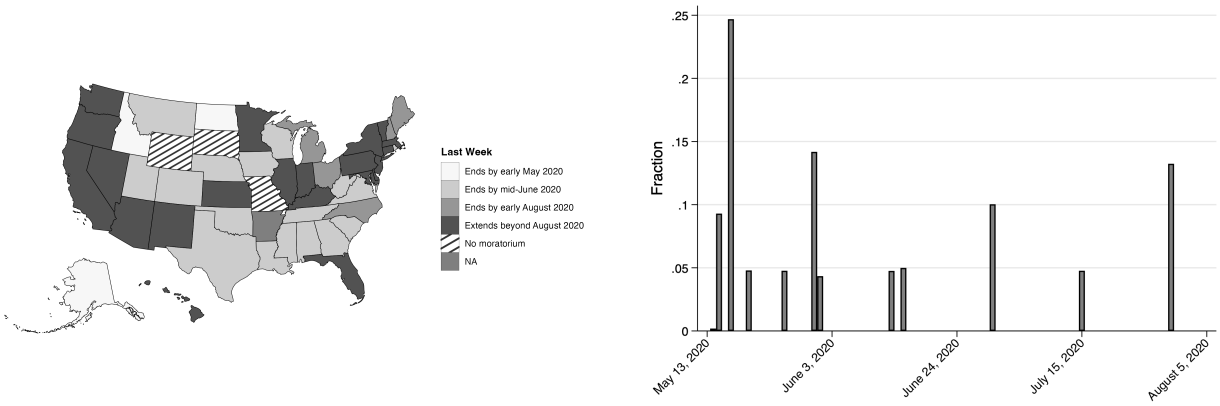
Air quality (fine particulate matter $PM_{2.5}$) data. To measure local air pollution exposure, we use satellite-derived estimates of fine particulate matter ($PM_{2.5}$) developed by van Donkelaar et al. (2021). These data combine satellite observations of aerosol optical depth with chemical transport models and ground monitoring stations in a geophysical–statistical framework to produce spatially continuous estimates of surface-level $PM_{2.5}$ concentrations. The dataset provides high-resolution measures of ambient air pollution that are widely used in environmental and health research, particularly in locations where ground monitoring networks are sparse. In the heterogeneity analysis, we classify areas into above- and below-median $PM_{2.5}$ exposure to test whether discrimination patterns differ between more and less polluted environments.

InfoUSA (Data Axle) residential data. To study whether discriminatory screening affects realized housing outcomes, we link the experimental rental listings to the residential address history database compiled by InfoUSA (now Data Axle). The database is a commercial marketing dataset that covers nearly the entire U.S. population and reports household-level information, including address histories, owner-renter status, ethnicity codes, and the length of residence at the current address.

We merge the 2020–2022 vintages of the database with the experimental addresses using several matching procedures described in Appendix D. Using the observation date and the reported length of residence, we infer the tenant's move-in date and retain households whose move-in occurs after the experimental inquiry date. Ethnicity codes are mapped into four race categories (White, Black, Hispanic, and Asian/Other) using vendor and manually constructed crosswalks. The resulting dataset provides a mapping between experimental rental listings and the race of the tenant who subsequently occupied the unit and forms the basis for our outcome test of discrimination in Section 9.1.

4 Results

4.1 Defining Treatment



(a) The Last Week of the Eviction Moratorium

(b) The Last Date of Eviction Moratorium

Figure 1: Eviction Moratoria Expiration across U.S. States. Panel (a) shows the last week of the eviction moratorium across the United States. Panel (b) presents the distribution of moratorium expiration dates across states.

Because eviction moratoria were introduced within a narrow window at the onset of the COVID-19 pandemic, policy onset provides limited identifying variation across states. Moreover, the timing of enactment is unlikely to be exogenous because states tended to adopt moratoria in response to worsening local pandemic conditions. Expirations, in contrast, occurred in a staggered fashion and therefore provide substantially richer variation in the timing of restored eviction enforcement. Figure 1 depicts this variation in the timing of expiration of state-level moratoria across the United States.

The expiration timing also does not appear to be systematically related to daily COVID-19 rates. Table 2 presents estimates from Callaway and Sant'Anna (2021)'s staggered Difference-in-Differences of the effect of moratorium expiration on daily COVID-19 cases per 100,000 residents. Within windows of 30, 45, and 60 days around the expiration dates, repeal has no statistically significant effect on case rates. The expiration of eviction moratoria therefore does not predict short-run changes in COVID-19 dynamics, indicating that public health conditions do not mechanically determine treatment timing.

Table 2: Staggered DiD Estimates of the Effect of the End of Moratorium on the Daily Number of the COVID-19 Cases in a State Per 100,000 People

	Days around treatment		
	30	45	60
ATT	33.221	-4.768	-8.274
	(-35.506, 101.948)	(-59.742, 50.207)	(-60.693, 44.145)
Number of Observations	868	1373	1889

Notes: 1) ATT stands for the Average Treatment Effect on the Treated. 2) 95% confidence intervals in parentheses. 4) ***, **, and * denote significance at the 1%, 5%, and 10% levels.

A natural alternative would be to condition on contemporaneous COVID-19 cases directly, for example by matching treated and control states on case rates. We deliberately do not do so. Identification under propensity score matching with staggered adoption requires a strong overlap condition: the conditional probability that a state s ends a moratorium on day g, given the number of COVID-19 cases X_s, must be bounded away from zero and one, $\varepsilon < \Pr(G_s = g \mid X_s) < 1 - \varepsilon$, for all repeal dates g. This requirement is difficult to guarantee in settings with staggered policy adoption (Baker et al., 2025), and it fails decisively in ours because treatment cohorts are thin. The twenty-four treated states end their moratoria on seventeen distinct dates, and thirteen of these dates correspond to a single state. With so many singleton cohorts, contemporaneous case rates can mechanically separate the expiring state from the rest of its risk set: estimating the probability of expiration on each date yields fitted probabilities outside $[0.05, 0.95]$ for 77 percent of state–cohort observations. We therefore rely on the absence of differential pre-trends rather than on conditioning on pandemic conditions.

These reasons make expiration dates a more credible source of quasi-experimental variation than enactment dates. We thus define treatment based on the *expiration* of a moratorium, in line with prior work that likewise studies the effects of moratoria ending rather than their initial adoption (Benfer et al., 2021; Hatamyar and Parmeter, 2023).

During our sample period, eviction restrictions also arose from the federal CARES Act, which imposed a temporary moratorium on evictions for properties with federally backed mortgages or participating in federal housing programs. In Section 7.3, we examine whether our results could be driven by this federal policy and show that the main findings are not explained by the CARES Act coverage. In addition, the federal eviction moratorium issued by the Centers for Disease Control and Prevention (CDC) began on September 4, 2020. Because our experimental sample ends on July 31, 2020, the analysis exploits only cross-

state variation in the expiration of state-level moratoria that occurred prior to the CDC policy.

The correspondence experiment spans February 6 through July 31, 2020. To avoid conflating baseline discrimination with periods prior to the policy intervention, we restrict the sample to observations that occur after a state enacted its eviction moratorium. The resulting analysis sample begins on March 13, 2020, the earliest date on which a moratorium was active in any treated state. Figure 1b shows the distribution of moratorium expiration dates in the sample, which range from May 7 to July 30, 2020.

We define the treatment indicator $Treatment_{jt}$ as equal to one if inquiry j was sent after the expiration of the eviction moratorium in the corresponding state and zero otherwise. This design exploits cross-state variation in the timing of the restoration of eviction enforcement to identify its effect on property managers' screening behavior.

4.2 Identification Assumptions

Although eviction moratoria were enacted within a narrow time window across states, their expiration occurred at different points during the spring and summer of 2020. This staggered repeal raises a natural concern that states may have lifted moratoria in response to evolving public health conditions or pre-existing socioeconomic characteristics, and that such factors could be correlated with trends in discriminatory screening.

Our baseline design compares states that have already ended their moratoria (the treatment group) with states whose moratoria remain in effect (the control group, consisting of not-yet-treated states). Our outcome measures the rate at which property managers respond to minority applicants relative to white applicants in the same state (defined in Section 4.3). Identification rests on two assumptions. The first is parallel trends: absent moratorium expiration, this measure would have evolved in parallel in treated and not-yet-treated states (Callaway and Sant'Anna, 2021, Assumption 5). The second is limited treatment anticipation: property managers do not adjust their screening behavior in advance of a scheduled expiration (Callaway and Sant'Anna, 2021, Assumption 3). The event-study evidence in Section 4.3 is consistent with these assumptions. Relative response rates exhibit no differential trends across states before moratoria expire, and the change in screening behavior is concentrated at the legal expiration date rather than before it.

The parallel trends assumption does not require that the timing of moratorium expiration be exogenous. Randomly assigned expiration dates would be sufficient, since random assignment implies parallel trends, but they are not necessary (Roth et al., 2023). Expiration timing may be correlated with state characteristics, with political affiliation, and even with

the severity of the pandemic itself, provided these factors do not generate differential trends in the relative outcome at the time of expiration. Because our outcome compares minority and white applicants within the same state, any state-specific shock that affects both groups proportionally (local labor-market deterioration, court backlogs, pandemic severity) differences out. What the assumption rules out is selection on counterfactual trends in the relative outcome: state-specific, race-differential shocks to screening behavior that coincide with moratorium expiration and operate through channels other than eviction enforcement. Section 5 presents direct evidence that such selection is unlikely in our setting.

4.3 Staggered Differences-in-Differences

We examine how racial disparities in housing inquiry responses evolve following the expiration of state-level eviction moratoria. As shown in Appendix B, differential response rates across applicant racial identities are present throughout the sample period, consistent with the framework in Section 2. We therefore focus on whether the removal of eviction moratoria altered the intensity of discriminatory screening.

State eviction moratoria were lifted at different points in time, generating staggered treatment adoption. In such settings, conventional two-way fixed-effects Difference-in-Differences estimators combine comparisons across cohorts with heterogeneous treatment timing, which can lead to attenuation or sign distortions in the presence of dynamic or heterogeneous treatment effects. We therefore rely on an estimator that is robust to staggered adoption and explicitly accommodates treatment effect heterogeneity. For completeness, Appendix C reports conventional DiD estimates, which are directionally consistent with our main findings but statistically attenuated.

Our primary analysis uses the staggered Difference-in-Differences estimator developed by Callaway and Sant'Anna (2021). Applying this estimator directly to the inquiry-level data is not possible: each rental listing appears in our correspondence study only briefly, receives inquiries on a small number of dates, and is rarely observed both before and after its state's moratorium expires. The data therefore do not contain a within-listing pre/post panel, and the listing cannot serve as the unit of analysis. We instead aggregate inquiries to the state level and construct, for each state, a time series of racial discrimination in inquiry responses. We then apply the Callaway–Sant'Anna (CSDiD) procedure to these state-level outcomes to recover dynamic treatment effects associated with moratorium expiration. We implement this in two stages: the first constructs the state-by-day discrimination measures, and the second applies the CSDiD estimator to them.

Stage 1: Estimating Daily Discrimination Coefficients for Each State

In the first stage, we estimate the degree of racial and ethnic discrimination for each state s and day $\tau \in \{1, \ldots, 177\}$, spanning February 6 to July 31, 2020. Specifically, we model the probability of a response to a rental inquiry as a function of race and identity characteristics:

$$Response_{ijst} = \beta_{s\tau}^{AA} \cdot African\ American_j + \beta_{s\tau}^{H} \cdot Hispanic_j + X_j'\theta_{s\tau} + u_{ijst}, \tag{8}$$

where i indexes rental properties, j indexes fictitious identities, s indexes states, and t is the date of inquiry. The outcome variable $Response_{ijst}$ equals one if identity j receives a response from property i in state s on day t. The regressors include indicator variables for race and a vector X_j of a constant and identity-level covariates (gender, maternal education, and inquiry order). Each regression is estimated via logit to ensure predicted probabilities remain bounded between zero and one.

We estimate equation (8) by local likelihood (Tibshirani and Hastie, 1987; Staniswalis, 1989; Cai et al., 2000; Li et al., 2011), a varying-coefficient logit in which the discrimination coefficients $\beta_{s\tau}^{R}$ vary smoothly with the day of inquiry. Using *all* observations from state s, we use all observations from state s and weight each observation by its temporal distance from day τ with a Gaussian kernel:

$$\omega_{ijst}^{\tau} = \frac{1}{h\sqrt{2\pi}} \exp\left(-\frac{1}{2}\left(\frac{t-\tau}{h}\right)^2\right), \tag{9}$$

where h is a smoothing parameter. This approach addresses the concern that sample sizes per state-day may be too small to yield reliable estimates. By pooling all available data for each state and upweighting observations close to τ, we recover smoothly varying state-day estimates of discrimination that reflect local trends without overfitting to daily noise. We use the coefficients only within a window around treatment and show robustness to different smoothing parameters h.

The kernel serves two distinct purposes that are central to our design. First, because each state-day cell contains relatively few inquiries, pooling observations from nearby days raises the effective sample size used to estimate $\beta_{s\tau}^{R}$ and stabilizes the discrimination coefficient. Second, and equally important, smoothing preserves daily resolution in the constructed outcome. This allows the second stage to use the exact calendar day on which each state's moratorium expired as the treatment date, rather than coarsening treatment timing to a weekly or monthly interval.

This estimation procedure gives us predicted discrimination coefficients $\{\beta_{s\tau}^{AA}, \beta_{s\tau}^{H}\}$ for each state s and day τ for African American (AA) and Hispanic (H) identities. Cai et al.

(2000) establish the consistency and asymptotic normality of this estimator, so these coefficients are well-behaved inputs to the second stage. They are nonetheless estimated rather than observed, and the second stage treats them as data. To confirm that our conclusions do not depend on this first-stage sampling error, Section 7.5 reports a bootstrap that propagates the uncertainty in $\{\beta_{s\tau}^{AA}, \beta_{s\tau}^{H}\}$ through to the treatment effects; the results are robust to estimation error in Stage 1.

Stage 2: Estimating Treatment Effects

We use a sequence of discrimination estimates $\{\beta_{s\tau}^{AA}, \beta_{s\tau}^{H}\}$ for each state s and day τ from Stage 1 to compute predicted response probabilities for representative applicants of each racial group — specifically, males with low maternal education who send the first message. Denote the value of these covariates as $\bar{X}$ and the predicted probability of a response for race $R \in \{AA, H, W\}$ as $\rho_{s\tau}^{R}$

$$\rho_{s\tau}^{R} = \Pr\big(\text{Response}_{ijst} = 1 \mid \text{Race} = R,\ X_j = \bar{X}; s, \tau\big) = \Lambda\Big(\hat{\beta}_{s\tau}^{R} + \bar{X}'\hat{\theta}_{s\tau}\Big), \tag{10}$$

where $\Lambda(z) = \exp(z)/(1+\exp(z))$ is the logistic cumulative distribution function. Using these predicted probabilities, we construct the natural logarithm of the relative response ratio:

$$\log(RR_{s\tau}^{R}) \equiv \log\left(\frac{\rho_{s\tau}^{R}}{\rho_{s\tau}^{W}}\right), \tag{11}$$

which serves as the outcome variable in the second-stage CSDiD estimation. Our main results are robust to using the relative response ratio as the outcome variable (Table A5).

Prior to applying the estimator, we restrict the sample in two ways. First, we exclude all observations from before a state enacted its moratorium and all τ for which the state has fewer than 100 inquiries to estimate $\beta_{s\tau}^{R}$ reliably. Second, we define a symmetric event window $|\tau - \tau_s^*| \leq \hat{\tau}$, where τ_s^* is the day the moratorium ends in state s, and $\hat{\tau} \in \{30, 45, 60\}$ determines the window size.

We follow Callaway and Sant'Anna (2021)'s methodology and define a state-day observation as treated on day τ if the state ended its moratorium before or on this day, and not treated if the state did not end its moratorium by that time. Our baseline control group is not-yet-treated states, but we show robustness to using not-treated states in the next section. To estimate the average treatment effect on states treated in period g for different periods t, $ATT(g,t)$, we first keep observations for states that had a moratorium on day g or have not yet had the moratorium by time t. We then use these observations to run the

regression

$$\log(RR_{s\tau}) = \alpha_0^g + \alpha_1^g \cdot TREAT_s^g + \alpha_2^g \cdot \mathbb{1}\{\tau = t\} + \alpha_3^{g,t} \cdot TREAT_s^g \times \mathbb{1}\{\tau = t\} + \nu_{s\tau}, \quad (12)$$

where $TREAT_s^g$ indicates whether state s was treated on day g. The interaction coefficient $\alpha_3^{g,t}$ identifies the group-time average treatment effect $ATT(g,t)$, which is then aggregated across groups and periods using the CS weighting scheme.

The estimates from equation (12) identify the causal effect of eviction moratorium expiration on relative response ratios under a conditional parallel trends assumption. Specifically, absent changes in eviction enforcement, the evolution of responses to minority applicants relative to white applicants would have followed similar trends across states with different moratorium expiration dates. This assumption allows for arbitrary cross state differences in levels and for common time-varying shocks, but rules out state-specific shocks coinciding with moratorium expiration that differentially affect property managers' screening behavior by race through channels other than eviction enforcement.

As in all difference-in-differences designs, the parallel trends assumption is not directly testable. Its plausibility is supported by the absence of differential pre-trends in relative response ratios prior to moratorium expiration. The group-time average treatment effects from the CSDiD event study analysis discussed below exhibit flat pre-treatment dynamics.

Results

Figure 2 presents event study estimates for African American applicants using a smoothing parameter of $h = 10$ and an event window of $\hat{\tau} = 45$ days. Relative response ratios increase sharply following moratorium expiration, indicating a reduction in racial discrimination once eviction enforcement is reinstated. There is no evidence of differential pre-trends prior to expiration. The timing of the response aligns closely with the legal expiration date of the moratorium, consistent with a discrete change in enforcement rather than gradual shifts in market conditions.

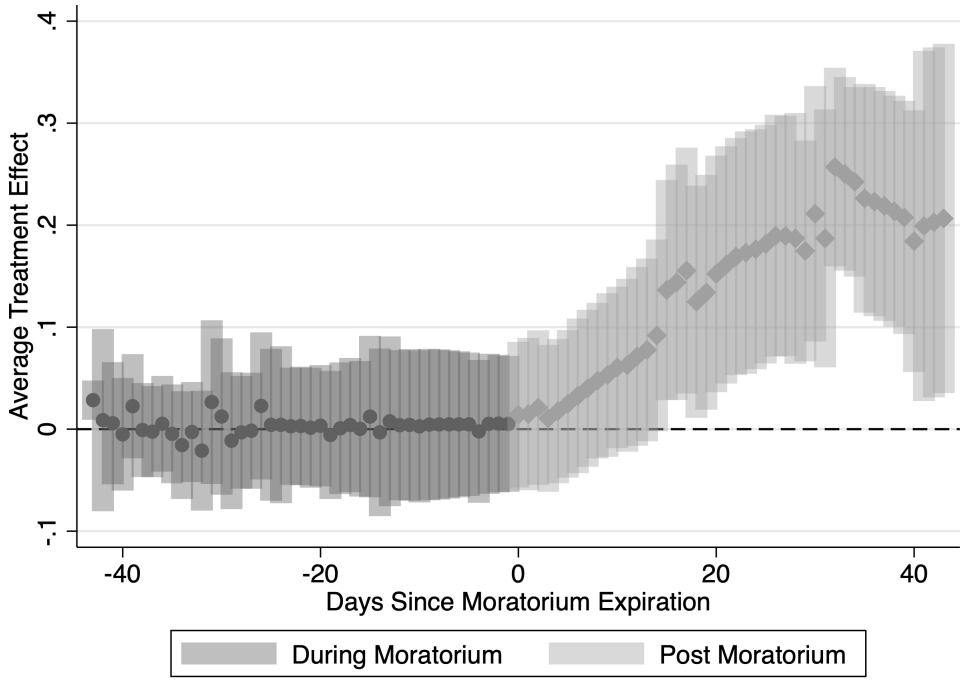


Figure 2: The Event Study Estimates of the Effect of the End of a Moratorium on the Natural Logarithm of the Relative Response Ratios for an African American Identity Relative to a White Identity from Callaway and Sant'Anna (2021)'s Staggered DiD with the Smoothing Parameter $h = 10$ and $\hat{\tau} = 45$ Days around Treatment

Under this design, potential threats to identification must involve state-specific, race-differential shocks that vary at the daily frequency and coincide precisely with moratorium expiration. Appendix Figure A2 shows that these patterns are robust across alternative smoothing parameters h and event-window lengths $\hat{\tau}$.

Table 3: Staggered DiD Estimates of the Effect of the End of Moratorium on the Logarithm of the Relative Response Ratio Relative to White Applicants

	African American			Hispanic		
	$h = 7$	$h = 10$	$h = 15$	$h = 7$	$h = 10$	$h = 15$
Panel A: 30 days around treatment						
ATT	0.111***	0.076***	0.048**	0.029	0.018	0.008
	(0.053, 0.170)	(0.024, 0.128)	(0.003, 0.093)	(-0.029, 0.087)	(-0.031, 0.068)	(-0.031, 0.047)
Number of Observations	724	724	724	724	724	724
Panel B: 45 days around treatment						
ATT	0.164***	0.120***	0.075***	0.063*	0.051**	0.033*
	(0.102, 0.226)	(0.069, 0.172)	(0.033, 0.118)	(-0.001, 0.127)	(0.000, 0.101)	(-0.006, 0.072)
Number of Observations	1139	1139	1139	1139	1139	1139
Panel C: 60 days around treatment						
ATT	0.126***	0.099***	0.065***	0.137***	0.112***	0.081***
	(0.057, 0.194)	(0.045, 0.154)	(0.021, 0.109)	(0.064, 0.209)	(0.055, 0.169)	(0.037, 0.125)
Number of Observations	1561	1561	1561	1561	1561	1561

Notes: 1) ATT stands for the Average Treatment Effect on the Treated. 2) h is the smoothing parameter of the weighted logit, see the text. 3) 95% confidence intervals in parentheses. 4) ***, **, and * denote significance at the 1%, 5%, and 10% levels.

Table 3 reports average treatment effects on the treated across alternative smoothing parameters and event windows. For African American applicants, all nine estimates are positive and significant at the 5% level or better, ranging from 0.048 to 0.164 log points. Point estimates decline in the smoothing parameter h, as expected when a wider kernel averages days on either side of the expiration date, so estimates at larger h are lower bounds. We take $h = 10$ with $\hat{\tau} = 45$, the specification plotted in Figure 2, as our preferred specification, with an ATT of 0.120 and a 95% confidence interval of $[0.069, 0.172]$. Estimates for Hispanic applicants are indistinguishable from zero at $\hat{\tau} = 30$ and reach magnitudes comparable to the African American estimates only at $\hat{\tau} = 60$, which we read as a delayed rather than an absent response. Because the African American estimates are significant in every cell of the table, the remainder of the analysis focuses on African American applicants.

To assess the economic significance, note that white applicants in our sample received responses 59.8% of the time during moratoria, while African American applicants received responses at a rate of 54.5%, implying a relative response ratio of 0.911. Our preferred ATT estimate implies that the log of this ratio rises by 0.120 after expiration, an increase in the relative response ratio of $100 \times (e^{0.120} - 1) = 12.8\%$. Applied to the during-moratorium baseline, this moves the relative response ratio from 0.911 to $0.911 \times e^{0.120} = 1.027$, and the 95% confidence interval $[0.069, 0.172]$ corresponds to a post-expiration ratio between

0.976 and 1.082. The point estimate therefore implies that the racial gap observed under the moratorium closes entirely once enforcement is restored, with the confidence interval bracketing parity. We do not interpret a point estimate above one as evidence of a reversal in favor of African American applicants.

The relative response ratio, however, measures only how screening is allocated across applicants; it is silent on whether the overall level of screening changed. We turn to that question next.

4.4 Levels versus Composition of Screening

Suspending eviction enforcement could distort tenant screening on two margins. Property managers could ration access, responding to fewer inquiries from everyone while the cost of a bad match is elevated; or they could reallocate access, redirecting responses across applicants they perceive differently. The distinction matters for welfare and for policy: rationing shrinks the market for all prospective tenants, while reallocation concentrates the burden of weakened enforcement on particular groups. It also matters for the literature, which measures the consequences of eviction protections through rents (Abramson, 2025; Corbae et al., 2024; Coulson et al., 2025) and cannot see either margin. Our data separate the two.

We measure the two margins within a single decomposition. Stage 1 of our procedure delivers, for each state s and each day τ of the experiment, the predicted probability $\rho^R_{s\tau}$ that a property manager responds to an inquiry from an applicant of identity R, where R is white (W), African American (AA), or Hispanic (H) (equation (10)). We define the level of responsiveness in state s on day τ as the equal-weight mean of the log response probabilities,

$$L_{s\tau} = \tfrac{1}{3}\left(\ln \rho^W_{s\tau} + \ln \rho^{AA}_{s\tau} + \ln \rho^H_{s\tau}\right), \tag{13}$$

the log of the geometric-mean response probability across the three identities. Any change in the response probabilities then splits exactly into the two margins,

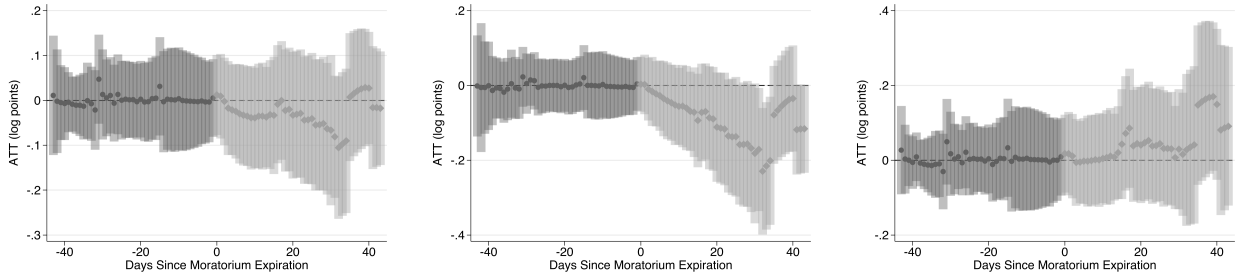
$$\Delta \ln \rho^R_{s\tau} \;=\; \underbrace{\Delta L_{s\tau}}_{\text{level}} \;+\; \underbrace{\Delta\left(\ln \rho^R_{s\tau} - L_{s\tau}\right)}_{\text{composition}}, \qquad R \in \{W, AA, H\}, \tag{14}$$

where the level component is common to all three identities and the composition components sum to zero across identities and are spanned by the two relative response ratios analyzed above. Because L is measured in logs, its treatment effects read as percent changes in responsiveness, on the same scale as the estimates in Table 3.

Table 4: Staggered DiD Estimates of the Effect of the End of Moratorium on the Level and Composition of Responsiveness

	Level L	$\ln \rho^W$	$\ln \rho^{AA}$	$\ln \rho^H$
		Panel A: 30 days around treatment		
ATT	0.015	−0.016	0.059	0.002
	(−0.050, 0.079)	(−0.067, 0.035)	(−0.027, 0.145)	(−0.077, 0.081)
Number of Observations	724	724	724	724
		Panel B: 45 days around treatment		
ATT	−0.014	−0.071***	0.050	−0.020
	(−0.072, 0.044)	(−0.115, −0.026)	(−0.022, 0.121)	(−0.099, 0.059)
Number of Observations	1139	1139	1139	1139
		Panel C: 60 days around treatment		
ATT	−0.012	−0.082***	0.017	0.030
	(−0.072, 0.048)	(−0.128, −0.037)	(−0.057, 0.091)	(−0.054, 0.113)
Number of Observations	1561	1561	1561	1561

Notes: 1) ATT stands for the Average Treatment Effect on the Treated. 2) All outcomes are in log points: $\ln \rho^R_{s\tau}$ is the log predicted probability of a response in state s on experiment day τ to an applicant of identity $R \in \{W, AA, H\}$, where W denotes a white, AA an African American, and H a Hispanic identity. The level is $L_{s\tau} = \frac{1}{3}(\ln \rho^W_{s\tau} + \ln \rho^{AA}_{s\tau} + \ln \rho^H_{s\tau})$, the log of the geometric-mean response probability. 3) Estimates use the smoothing parameter $h = 10$ of the weighted logit defined in Section 4.3; the remaining smoothing parameters, reported in the replication files, display the same pattern. 4) Predicted probabilities are clamped to $[0.001, 0.999]$ before taking logs. 5) 95% confidence intervals in parentheses. 6) ***, **, and * denote significance at the 1%, 5%, and 10% levels.



(a) Level of Responsiveness, L (b) Whites, $\ln \rho^{W}$ (c) African Americans, $\ln \rho^{AA}$

Figure 3: The Event Study Estimates of the Effect of the End of a Moratorium on the Level of Responsiveness and on the Log Response Probabilities by Applicant Identity, from Callaway and Sant'Anna (2021)'s Staggered DiD with the Smoothing Parameter $h = 10$ and $\hat{\tau} = 45$ Days around Treatment

Notes: 1) Panel (a) plots event-study estimates for the level of responsiveness $L_{s\tau}$, the equal-weight mean of the log response probabilities across the white, African American, and Hispanic identities; panels (b) and (c) plot the log response probabilities for the white and African American identities. 2) Bars are 95% confidence intervals: dark bars correspond to days before the state's moratorium expired (pre-treatment) and light bars to days after expiration (post-treatment). 3) All outcomes are in log points.

Moratorium expiration did not change how many responses property managers sent. Table 4 and Figure 3 (a) show that the level of responsiveness is flat through expiration: the estimated effects on L are below 0.015 log points in absolute value in the 45- and 60-day windows, and their confidence intervals exclude level effects as large as the compositional effect documented above. The flat level is not an averaging artifact concealing a common decline. By construction, the effect on L is the equal-weight average of the three group-level effects: the decline for white identities enters with weight one-third and is partly offset by the positive, though imprecise, estimate for African American identities, so the group-level estimates in Panel B imply an aggregate change of -0.014 log points, which is precisely the estimate reported in the first column, and well within its confidence interval.[2]

What expiration changed is who received the responses. Responses to white identities fall by seven to eight percent after expiration (Figure 3 (b)); responses to African American identities do not fall, with positive point estimates throughout (Figure 3 (c)). During the moratoria, that is, property managers did not withdraw from the market. Instead, they tilted their responses toward the applicants they perceived as safer, and the tilt reversed when eviction became enforceable again. The decomposition is also consistent with the headline estimates: the difference between the two group-level estimates in Panel B of Table 4 implies

[2]The arithmetic mean response probability, which under the experiment's randomization of identities equals the expected response rate to a randomly drawn inquiry, tells the same story in levels.

a relative response ratio effect of 0.121, nearly identical to the 0.120 estimated directly from the log relative response ratio in Table 3.

Taken together, the suspension of eviction enforcement reallocated access to rental housing rather than rationing it: it changed *who* received responses, not *how many* were sent. This margin is invisible in rent data, and its burden fell on African American applicants, the group toward whom responses did not flow while eviction was unenforceable.

5 Assessing the Validity of Identification Assumptions

The estimates in Section 4 carry a causal interpretation only if the timing of moratorium expiration was not driven by forces that independently affect discriminatory screening. This section probes that requirement from two directions. Section 5.1 shows that expiration neither followed nor coincided with changes in time-varying economic conditions. Section 5.2 shows that expiration timing is not predicted by predetermined socioeconomic characteristics, and that the cross-state variation in moratorium duration is instead explained largely by political preferences fixed before the pandemic. Together, these results indicate that the timing variation we exploit reflects predetermined institutional differences across states rather than responses to concurrent shocks.

5.1 Selection on Time-Varying Economic Conditions

The decision to lift an eviction moratorium was made by state officials and need not have been exogenous. Our parallel trends assumption would fail if states timed expiration in response to evolving economic conditions, such as a recovering labor market or improving household finances, and those same conditions independently affected screening behavior. In that case, the estimated effect of restored enforcement would conflate the change in enforcement risk with the economic changes that prompted it. We therefore test directly whether moratorium expiration followed, or coincided with, shifts in time-varying economic conditions at the state level.

We construct a state-by-week panel from the U.S. Census Bureau's Household Pulse Survey, described in Section 3, and measure four outcomes: the employment share of prime-age (25–54) adults, household income, the share of renter households behind on rent, and the share of homeowner households behind on mortgage payments.[3] We then re-estimate our staggered Difference-in-Differences specification on this panel, replacing the discrimination

[3]We focus on prime-age employment because employment at the age margins shifted during the pandemic for reasons unrelated to local economic conditions; estimates for ages 18–67 are similarly indistinguishable from zero.

outcome with each economic indicator: we use the survey week in which a state's moratorium expired as its treatment week, and not-yet-treated states serve as the control group, as in the baseline specification.[4]

Figure 4 presents the event studies. Two patterns stand out. First, there are no differential pre-trends: in the weeks leading up to expiration, treated states do not exhibit rising employment, income, or improving delinquency relative to not-yet-treated states, and joint tests of the pre-expiration coefficients fail to reject the null of no pre-trends for any outcome (p-values range from 0.101 to 0.925). Had governors lifted moratoria once local labor markets recovered or once households caught up on rent, these series would trend upward (or, for delinquency, downward) ahead of expiration; they do not. Second, economic conditions do not change after expiration: for all four outcomes, the average post-expiration effect is small and statistically indistinguishable from zero.

These results address the selection concern from both directions. The absence of pre-expiration trends indicates that expiration timing was not a response to evolving economic conditions, and the absence of post-expiration shifts indicates that the increase in relative response rates documented above cannot be attributed to economic changes coinciding with the restoration of enforcement. The rent delinquency series is particularly informative: it is the economic margin most directly tied to eviction policy, and it is flat through the expiration window.

[4]The Pulse Survey was fielded weekly through July 2020 and in two-week collection waves thereafter, which the Census Bureau continues to label "weeks." We refer to all collection waves as weeks; the horizontal axis of Figure 4 counts these waves.

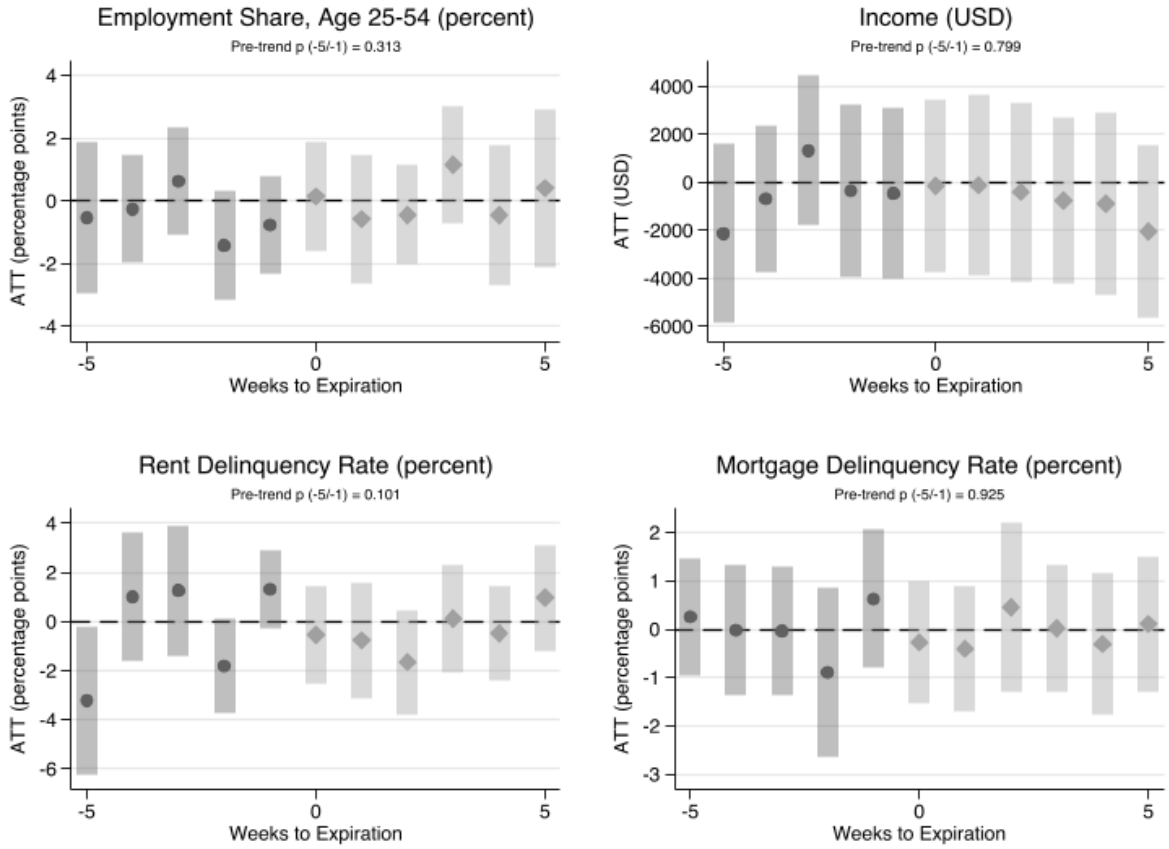


Figure 4: The Event Study Estimates of the Effect of the End of a Moratorium on State-Level Economic Conditions from Callaway and Sant'Anna (2021)'s Staggered DiD, Using the Census Household Pulse Survey

Notes: 1) Each panel plots event-study estimates of the effect of moratorium expiration on a state-level economic outcome, by survey week relative to expiration. 2) Bars are 95% confidence intervals: dark bars correspond to weeks before the state's moratorium expired (pre-treatment) and light bars to weeks after expiration (post-treatment). 3) "Pre-trend p $(-5/-1)$" is the p-value from a joint test that the five pre-expiration coefficients shown (weeks -5 through -1) are all equal to zero; large values indicate no differential trends between treated and not-yet-treated states ahead of expiration. 4) Employment and delinquency outcomes are measured in percent, so estimates are in percentage points; income is measured in U.S. dollars.

5.2 Selection on Time-Invariant State Characteristics

Table 5: Predicting the Timing of Eviction Moratorium Expiration

	(1)	
	Coefficient	95% Confidence Interval
First day of moratorium	-2.61	(-21.14, 15.92)
Total population in 100k	-0.09	(-1.55, 1.36)
Population density	-0.09	(-0.27, 0.09)
Log median income	99.20	(-1507.98, 1706.39)
Percent of people who are over 65 years old	20.97	(-44.94, 86.89)
Percent of African Americans	1.09	(-14.80, 16.98)
Percent of Asians	-19.08	(-59.90, 21.73)
Percent of American Indian	-7.11	(-63.96, 49.74)
Percent of Hispanics	-3.70	(-24.82, 17.42)
Percent of renters	27.19	(-6.92, 61.30)
Percent of people without high school degrees and below	-41.46	(-134.57, 51.64)
Percent of people with a college degree and above	-1246.39	(-5684.00, 3191.22)
Percent of people in group quarters	-143.63	(-417.94, 130.68)
Percent of essential workers	-27.97	(-108.30, 52.36)
Percent of people who are uninsured	-25.64	(-89.86, 38.58)
Percent of people who use public transportation	-292.81	(-662.98, 77.35)
Percent of people who carpool	-259.91	(-766.70, 246.88)
Percent of people who commute by driving alone	-284.29	(-657.04, 88.46)
Percent of people who commute using motorcycle	-338.12	(-2815.56, 2139.32)
Percent of people who commute using bicycle	-9.95	(-548.17, 528.27)
Percent of people who commute by walking	-325.40	(-826.66, 175.87)
Percent of people who work at home	-389.06*	(-791.71, 13.60)
Observations	39	

Notes: The dependent variable is the date of eviction moratorium expiration at the state level. All covariates are predetermined and come from the 2015–2019 five-year American Community Survey. The specification controls for the moratorium start date to account for mechanical differences in policy duration. Reported coefficients are from a second-stage regression of state fixed effects obtained after partialling out contemporaneous COVID-19 conditions. The absence of systematic predictors indicates that moratorium repeal timing is orthogonal to observable state fundamentals. Ninety-five percent confidence intervals are reported in parentheses. ***, **, and * denote significance at the 1%, 5%, and 10% levels.

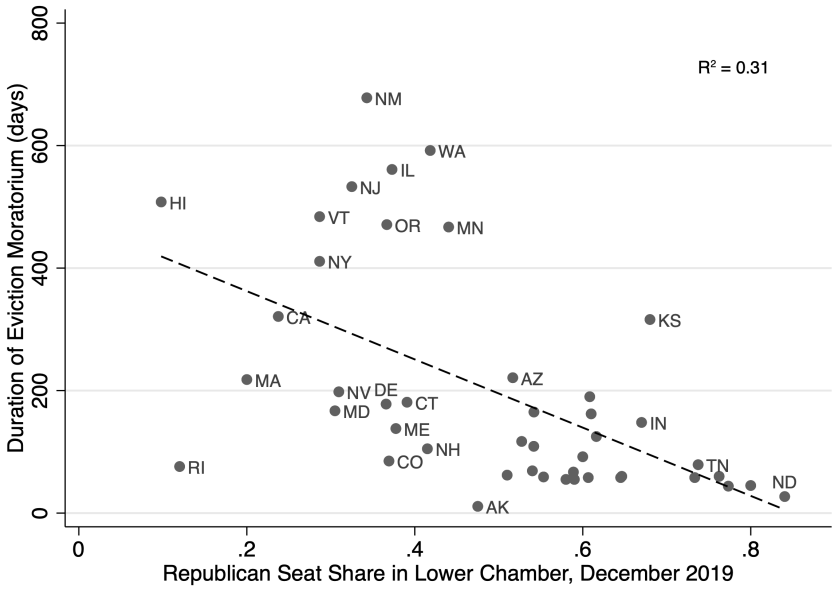


Figure 5: Duration of Eviction Moratoria and the Republican Seat Share in the State Legislature

Notes: Each dot is a state that enacted an eviction moratorium during the pandemic. The horizontal axis is the Republican seat share in the state's lower legislative chamber in December 2019, from the National Conference of State Legislatures; the vertical axis is the number of days the state's eviction moratorium was in effect. The line is the least-squares fit.

Having shown that expiration timing did not respond to time-varying economic conditions, we next ask what does explain the cross-state variation in when moratoria ended. We proceed in two steps: we first show that expiration timing is not predicted by a broad set of predetermined socioeconomic characteristics, and we then show that the variation is instead explained largely by political preferences fixed before the pandemic.

We implement a two-stage procedure. In the first stage, we regress the date of moratorium expiration on state fixed effects and contemporaneous COVID-19 conditions. In the second stage, we regress the estimated state fixed effects on pre-pandemic socioeconomic characteristics from the 2015–2019 American Community Survey, controlling for the moratorium start date to account for mechanical differences in policy duration. As Table 5 shows, none of these predetermined covariates predicts the timing of repeal at the 5% level. Variation in moratorium expiration was therefore not systematically driven by observable state fundamentals, and to the extent that such characteristics matter, they are absorbed by state fixed effects in our main specifications.

In contrast, predetermined political preferences have substantial explanatory power. Re-

gressing the duration of each state's moratorium on the Republican seat share in the state's lower legislative chamber, measured in December 2019, yields a strong negative relationship: states with more Republican legislatures ended their protections sooner, and this single predetermined variable alone explains roughly a third of the cross-state variation in moratorium duration ($R^2 = 0.31$; Figure 5). Adding the seat share to the specification in Table 5 does not change the conclusion.

From the standpoint of identification, this pattern is benign. The threat to parallel trends is selection on shocks that coincide with expiration; a determinant fixed months before the pandemic cannot generate such shocks. Political composition may well correlate with the level of discriminatory screening across states, but level differences are differenced out by the design. What matters is that the timing variation traces to predetermined institutional preferences rather than to conditions evolving at the time of expiration.[5]

6 Alternative Explanations

This section examines whether factors other than the restoration of eviction enforcement could account for the changes in screening behavior documented above. In particular, we consider three potential confounding mechanisms. First, moratorium expiration could coincide with changes in rental market conditions, such as shifts in the number or characteristics of units listed for rent. Second, our estimates could reflect contemporaneous public health dynamics during the COVID-19 pandemic. Third, screening behavior may have changed in response to heightened public awareness of racial discrimination following the murder of George Floyd. Across these tests, the evidence consistently indicates that the observed changes in relative response rates are most consistent with changes in enforcement risk rather than with these alternative explanations.

6.1 Market Conditions and Rental Unit Characteristics

A potential alternative explanation for the observed changes in relative response ratios is that market conditions shifted around moratorium expiration, for example, through changes in the volume or composition of rental units listed for rent. To examine this possibility, we use CoreLogic MLS data to construct a comprehensive panel of rental listings from 2018–2020. Rental units with multiple listing spells are merged following the procedure in Garriga

[5]Predictability by political composition does not make the seat share a valid instrument for expiration timing: political composition may affect screening through channels other than the moratorium, for example through the stringency of fair-housing oversight (Section 8.5), which would violate the exclusion restriction. We therefore use it descriptively rather than as an instrument.

Table 6: Staggered DiD Estimates of the Effect of the End of Moratorium on Housing Characteristics

	# Log(1 + Listings)	Sq. Ft.	Beds	Baths	Log(1 + Age)
			Panel A: 30 days around treatment		
ATT	0.018	0.014	-0.028	0.027	-0.060
	(-0.554, 0.591)	(-0.035, 0.064)	(-0.111, 0.055)	(-0.050, 0.105)	(-0.144, 0.024)
Number of Observations	1042	1009	1009	1009	1009
			Panel B: 45 days around treatment		
ATT	-0.104	0.003	0.020	0.047	-0.084*
	(-0.657, 0.448)	(-0.051, 0.057)	(-0.044, 0.084)	(-0.033, 0.126)	(-0.180, 0.012)
Number of Observations	1674	1626	1626	1626	1626
			Panel C: 60 days around treatment		
ATT	0.213	-0.006	0.034	-0.025	-0.072*
	(-0.340, 0.766)	(-0.067, 0.055)	(-0.046, 0.114)	(-0.120, 0.069)	(-0.155, 0.011)
Number of Observations	2280	2217	2217	2217	2217

Notes: 1) ATT stands for the Average Treatment Effect on the Treated. 2) 95% confidence intervals in parentheses. 4) ***, **, and * denote significance at the 1%, 5%, and 10% levels.

and Hedlund (2020).

Table 6 reports staggered DiD estimates of the effect of moratorium expiration on the number of rental listings and key unit characteristics, including square footage, number of bedrooms and bathrooms, and unit age. Across all event windows, the estimated effects are small in magnitude and statistically insignificant at the 5% level. These results indicate that neither the quantity of rental listings nor their observable characteristics change discontinuously at moratorium expiration, making it unlikely that shifts in the supply of rental units drive the changes in relative response rates documented above.

A related concern is that the composition of the pool of prospective tenants searching for rental housing may change around moratorium expiration, potentially in race-specific ways. Direct micro-level data on the racial composition of searching tenants are not available. However, such changes would threaten identification only under a restrictive set of conditions: applicant composition would need to shift differentially by race, differentially by state, and discontinuously at the moratorium expiration date, through channels unrelated to eviction enforcement. This conjunction of conditions is difficult to reconcile with the sharp timing of the estimated increase in the relative response rates and the lack of the pre-trends.

6.2 Black Lives Matter Movement and George Floyd Murder

Because the experiment was conducted during summer 2020, a potential concern is that heightened awareness of racial discrimination following the murder of George Floyd on May 25, 2020, may confound our estimates. If the intensification of the Black Lives Matter

movement independently altered screening behavior, it could contribute to the observed changes in responses to African American applicants.

Table 7: Estimates from Sample Before and After George Floyd Murder

	Before			After		
	$h = 7$	$h = 10$	$h = 15$	$h = 7$	$h = 10$	$h = 15$
Panel A: 30 days around treatment						
ATT	0.062	0.055	0.037	0.081**	0.042	0.025
	(-0.043, 0.168)	(-0.044, 0.153)	(-0.049, 0.124)	(0.013, 0.148)	(-0.015, 0.098)	(-0.024, 0.075)
Number of Observations	307	307	307	322	322	322
Panel B: 45 days around treatment						
ATT	0.047	0.038	0.024	0.145***	0.096***	0.057**
	(-0.058, 0.151)	(-0.060, 0.136)	(-0.061, 0.109)	(0.071, 0.219)	(0.040, 0.152)	(0.012, 0.102)
Number of Observations	481	481	481	529	529	529
Panel C: 60 days around treatment						
ATT	0.061	0.049	0.033	0.122***	0.083**	0.049*
	(-0.047, 0.169)	(-0.048, 0.147)	(-0.050, 0.116)	(0.036, 0.208)	(0.018, 0.149)	(-0.003, 0.101)
Number of Observations	653	653	653	719	719	719

Notes: 1) ATT stands for the Average Treatment Effect on the Treated. 2) h is the smoothing parameter of the weighted logit, see the text. 3) 95% confidence intervals in parentheses. 4) ***, **, and * denote significance at the 1%, 5%, and 10% levels.

To assess this possibility, we re-estimate our main specifications separately for subsamples before and after May 25, 2020. Table 7 reports the results. Although statistical precision declines due to smaller subsamples, the estimated effects remain economically meaningful and of similar sign across periods. The estimates are not statistically different even at the 10% significance level. Even if they were, persistent discrimination of African Americans during the eviction moratoria after the George Floyd murder goes against the hypothesis that the Black Lives Matter Movement is an explanation of our results.

Overall, the estimated effects across subsamples suggest that contemporaneous changes in racial awareness are unlikely to account for our main results. Instead, the evidence is more consistent with changes in screening behavior associated with the reinstatement of eviction enforcement.

7 Robustness

To strengthen the credibility of our estimates, we conduct a series of robustness checks that examine sensitivity to the definition of the control group, the presence of the federal

CARES Act moratorium, alternative staggered difference-in-differences estimators, and estimation uncertainty arising from the two-stage procedure.

7.1 Removing States with Discretionary Terminations

Table 8: Staggered DiD Estimates of the Effect of the End of Moratorium on the Natural Logarithm of the Relative Response Ratio for African Americans Relative to White Applicants for Subsamples of States

	No Discretionary Terminations			No Discretionary Terminations or Extensions		
	$h = 7$	$h = 10$	$h = 15$	$h = 7$	$h = 10$	$h = 15$
Panel A: 30 days around treatment						
ATT	0.105***	0.074***	0.048**	0.054	0.042	0.021
	(0.053, 0.156)	(0.026, 0.121)	(0.007, 0.089)	(-0.018, 0.127)	(-0.023, 0.107)	(-0.037, 0.079)
Number of Observations	672	672	672	378	378	378
Panel B: 45 days around treatment						
ATT	0.116***	0.091***	0.057***	0.094***	0.077***	0.040**
	(0.061, 0.170)	(0.041, 0.141)	(0.015, 0.100)	(0.047, 0.141)	(0.034, 0.120)	(0.003, 0.077)
Number of Observations	1039	1039	1039	732	732	732
Panel C: 60 days around treatment						
ATT	0.100***	0.084***	0.055***	0.111***	0.102***	0.062***
	(0.043, 0.157)	(0.034, 0.133)	(0.014, 0.095)	(0.066, 0.157)	(0.061, 0.144)	(0.026, 0.097)
Number of Observations	1403	1403	1403	970	970	970

Notes: 1) ATT stands for the Average Treatment Effect on the Treated. 2) h is the smoothing parameter of the weighted logit, see the text. 3) 95% confidence intervals in parentheses. 4) ***, **, and * denote significance at the 1%, 5%, and 10% levels.

Our identification strategy does not require the timing of moratorium expirations to be exogenous, but it does require that states did not end their moratoria in response to trends in discriminatory screening. To assess this concern, we trace the operative legal instrument that ended each state's moratorium and classify terminations into five mutually exclusive categories: court lapse (protection operated through court closures and ended when courts resumed; 12 states), scheduled executive-order lapse (a single order that expired on the end date fixed at issuance; 3 states), legislative sunset (protection that ended on a statutory sunset date; 3 states), extension followed by lapse (an order extended one or more times that expired on the stated end date of its final extension; 15 states), and discretionary termination (the acting authority retained the power to extend but affirmatively ended the moratorium; 8 states).

The first three categories are plausibly unrelated to contemporaneous conditions in the

rental market: the end date was fixed in advance by a court's reopening calendar, the terms of the original order, or statute, and was often mechanical, for example, set in 30-day increments or tied to the beginning or end of a month. Terminations through extension followed by lapse share this feature, since each end date was set when the extension was issued, but the decision to renew repeatedly could itself respond to current market conditions. Discretionary terminations pose the clearest threat to identification, because the end date was actively chosen and could therefore have responded to conditions in the housing market, including discriminatory screening.

We therefore re-estimate our benchmark specification on two subsamples: one that excludes states with discretionary terminations, and one that additionally excludes states whose moratoria were extended before lapsing. Because the not-yet-treated control group is drawn from the same set of states, these restrictions remove the excluded states from both the treated and the control groups.

We cannot estimate the model on the complementary subsample of discretionary terminations alone: too few of these states survive the sample restrictions, and each forms a singleton treatment cohort.

Table 8 reports the results. The left panel excludes only states with discretionary terminations; the right panel additionally excludes states with extensions. The estimates remain positive in all specifications and are statistically significant in the 45- and 60-day windows in both subsamples; in the most restrictive subsample, the 30-day estimates are attenuated and imprecise, consistent with the smaller sample. Moreover, none of the subsample estimates is statistically distinguishable from the corresponding benchmark estimate in Table 3 at the 5% level.[6] We conclude that our results are not driven by states that chose when to end their moratoria: the estimated increase in the relative response ratio survives restricting the sample to states whose moratoria ended through a court reopening, a pre-scheduled lapse, or a statutory sunset.

[6]The p-values from two-sample t-tests of equality between each subsample ATT and the corresponding benchmark ATT range from 0.08 to 0.99 across the 18 specifications. Because each subsample is nested in the benchmark sample, the two estimates are positively correlated, and treating them as independent overstates the standard error of the difference, so these tests are conservative.

7.2 Control Group

Table 9: Staggered DiD Estimates of the Effect of the End of Moratorium on the Logarithm of the Relative Response Ratio Relative to White Applicants Using Not Treated States

	Not Treated			Not Treated and Not Yet Treated		
	$h = 7$	$h = 10$	$h = 15$	$h = 7$	$h = 10$	$h = 15$
Panel A: 30 days around treatment						
ATT	0.126***	0.113***	0.095***	0.122***	0.086***	0.047**
	(0.084, 0.167)	(0.078, 0.148)	(0.066, 0.125)	(0.064, 0.180)	(0.038, 0.134)	(0.006, 0.088)
Number of Observations	941	941	941	942	942	942
Panel B: 45 days around treatment						
ATT	0.182***	0.168***	0.141***	0.163***	0.125***	0.078***
	(0.141, 0.222)	(0.134, 0.203)	(0.112, 0.170)	(0.104, 0.223)	(0.076, 0.174)	(0.037, 0.119)
Number of Observations	1381	1381	1381	1381	1381	1381
Panel C: 60 days around treatment						
ATT	0.204***	0.187***	0.157***	0.130***	0.102***	0.066***
	(0.164, 0.245)	(0.152, 0.221)	(0.128, 0.186)	(0.063, 0.197)	(0.050, 0.155)	(0.023, 0.108)
Number of Observations	1789	1789	1789	1790	1790	1790

Notes: 1) ATT stands for the Average Treatment Effect on the Treated. 2) h is the smoothing parameter of the weighted logit, see the text. 3) 95% confidence intervals in parentheses. 4) ***, **, and * denote significance at the 1%, 5%, and 10% levels.

We assess the robustness of our findings to the definition of the control group by re-estimating the staggered Difference-in-Differences specifications using states that never lifted their eviction moratoria during the sample period as the control group. These states provide a natural benchmark, as they remained untreated throughout the study window. Table 9 presents the results using this alternative control group in addition to the not-yet treated states in the right panel or as the only control group in the left panel.

The estimated effects remain positive and are statistically significant in most specifications. Notably, the estimated effects are similar in the magnitudes and are more consistently significant relative to the baseline specification that uses not-yet-treated states as the control group, consistent with the presence of discriminatory behavior during the moratorium period.

These results reinforce our core conclusion: the expiration of eviction moratoria is associated with a meaningful increase in relative response rates to African American rental applicants. The consistency of this pattern across "not-yet-treated", "never-treated", and "not-yet-treated" and "never-treated" control groups strengthens the credibility of our findings and suggests that they are not an artifact of control group selection.

7.3 Federal Eviction Moratorium

Table 10: Estimates During Federal Eviction Moratorium

	Benchmark Estimates			During Federal Moratorium		
	$h = 7$	$h = 10$	$h = 15$	$h = 7$	$h = 10$	$h = 15$
Panel A: 30 days around treatment						
ATT	0.111***	0.076***	0.048**	0.117***	0.083***	0.054**
	(0.053, 0.170)	(0.024, 0.128)	(0.003, 0.093)	(0.056, 0.178)	(0.029, 0.137)	(0.007, 0.101)
Number of Observations	724	724	724	693	693	693
Panel B: 45 days around treatment						
ATT	0.164***	0.120***	0.075***	0.133***	0.102***	0.063***
	(0.102, 0.226)	(0.069, 0.172)	(0.033, 0.118)	(0.075, 0.190)	(0.053, 0.152)	(0.021, 0.104)
Number of Observations	1139	1139	1139	1083	1083	1083
Panel C: 60 days around treatment						
ATT	0.126***	0.099***	0.065***	0.096***	0.083***	0.055***
	(0.057, 0.194)	(0.045, 0.154)	(0.021, 0.109)	(0.036, 0.155)	(0.034, 0.132)	(0.015, 0.096)
Number of Observations	1561	1561	1561	1449	1449	1449

Notes: 1) ATT stands for the Average Treatment Effect on the Treated. 2) h is the smoothing parameter of the weighted logit, see the text. 3) 95% confidence intervals in parentheses. 4) ***, **, and * denote significance at the 1%, 5%, and 10% levels.

The federal CARES Act imposed a nationwide eviction moratorium from March 27 to July 24, 2020, covering properties with federally backed mortgages. Although the policy overlapped with many state moratoria, its institutional structure differed in a key dimension: enforcement. State moratoria operated through local courts and directly suspended eviction filings and proceedings. By contrast, the CARES Act did not create an independent federal enforcement mechanism and did not grant Freddie Mac authority to intervene in state eviction courts.

The oversight reports from the Federal Housing Finance Agency Office of Inspector General underscore this distinction. The report concludes that Freddie Mac "does not have the authority or ability to directly enforce" compliance with the CARES Act and further states that the Enterprise "does not actively monitor its borrowers' compliance with the tenant protections of the CARES Act," instead relying on servicers to investigate alleged violations (FHFA Office of Inspector General, 2022). In other words, the federal moratorium lacked direct monitoring and judicial enforcement.

This institutional contrast allows us to isolate the role of contractual frictions. In our framework, discrimination increases when eviction ceases to be a credible enforcement threat. What matters for screening incentives is the expected enforceability of lease contracts, not

the formal existence of a statutory restriction.

We therefore use the CARES Act period as a cross-regime comparison. Table 10 reports staggered DiD estimates for the full benchmark sample and for the subsample restricted to the CARES Act window. Across our preferred specifications, the estimated effects are economically similar and statistically indistinguishable across regimes.

This pattern is consistent with the contractual frictions mechanism. Because the federal moratorium did not materially alter the enforceability of eviction through local courts, it did not meaningfully change property managers' expected loss from tenant default. Screening behavior responds to changes in enforcement credibility, such as the expiration of state-level court-imposed moratoria, rather than to the mere presence of a statutory restriction without direct enforcement authority.

This cross-regime comparison reinforces the model's central prediction that discrimination intensifies when eviction ceases to function as a credible enforcement mechanism.

7.4 Other Staggered Difference-in-Difference Estimators

Table 11: Estimates the Effect of the End of the Moratorium on the Natural Logarithm of the Relative Response Ratios for an African American Identity Relative to a White Identity from Alternative Staggered Difference-in-Differences Estimators

	Borusyak et al. (2024)			Sun and Abraham (2021)			de Chaisemartin and D'Haultfœuille (2020)		
	$h = 7$	$h = 10$	$h = 15$	$h = 7$	$h = 10$	$h = 15$	$h = 7$	$h = 10$	$h = 15$
Panel A: 30 days around treatment									
ATT	0.063***	0.061***	0.055***	0.242***	0.193***	0.147***	0.146***	0.123***	0.093***
	(0.026, 0.100)	(0.045, 0.077)	(0.044, 0.066)	(0.220, 0.263)	(0.179, 0.207)	(0.137, 0.158)	(0.073, 0.218)	(0.059, 0.187)	(0.052, 0.134)
Number of Observations	1010	1010	1010	941	941	941	1192	1192	1192
Panel B: 45 days around treatment									
ATT	0.071***	0.066***	0.058***	0.241***	0.216***	0.173***	0.208***	0.179***	0.132***
	(0.046, 0.096)	(0.050, 0.082)	(0.039, 0.076)	(0.220, 0.262)	(0.199, 0.233)	(0.160, 0.185)	(0.080, 0.336)	(0.102, 0.255)	(0.034, 0.229)
Number of Observations	1434	1434	1434	1381	1381	1381	1825	1825	1825
Panel C: 60 days around treatment									
ATT	0.056***	0.055***	0.049***	0.174***	0.158***	0.141***	0.194***	0.170**	0.127***
	(0.025, 0.087)	(0.027, 0.084)	(0.018, 0.080)	(0.152, 0.196)	(0.141, 0.174)	(0.130, 0.153)	(0.061, 0.328)	(0.025, 0.314)	(0.038, 0.215)
Number of Observations	1828	1828	1828	1789	1789	1789	2559	2559	2559

Notes: 1) ATT stands for the Average Treatment Effect on the Treated. 2) h is the smoothing parameter of the weighted logit, see the text. 3) 95% confidence intervals in parentheses. 4) ***, **, and * denote significance at the 1%, 5%, and 10% levels.

A standard concern in staggered DiD settings is that two-way fixed-effects style comparisons can be contaminated by comparisons between already-treated and later-treated units, which can attenuate or otherwise distort treatment effect estimates when effects are heterogeneous over event time. This concern is particularly relevant in our setting because states

exited eviction moratoria at different dates. Table 11 therefore re-estimates the effect of moratorium expiration using three alternative estimators proposed to address these issues: Borusyak et al. (2024), Sun and Abraham (2021), and de Chaisemartin and D'Haultfœuille (2020).

Across all three approaches, the estimated effect for African American applicants remains positive in nearly every specification and statistically significant in the large majority of cases. This pattern holds across event windows of 30, 45, and 60 days around treatment and across alternative values of the smoothing parameter h. The consistency in sign across estimators is important because these procedures rely on different comparison sets and weighting schemes. The fact that they all imply an improvement in the relative response rate of African American applicants after the end of the moratorium indicates that our core result is not driven by forbidden comparisons or by a particular implementation of staggered DiD.

At the same time, the magnitude of the estimates varies across estimators, which is expected given differences in estimands, sample composition, and weighting of cohort-time effects. We therefore do not emphasize any single point estimate from Table 11. Instead, we interpret the table as a robustness exercise showing that the qualitative conclusion is stable: discriminatory screening is stronger while eviction enforcement is suspended and declines once moratoria expire.

7.5 Errors from the Two-stage Procedure

Table 12: Bootstrapped Staggered DiD Estimates of the Effect of the End of the Moratorium on the Natural Logarithm of the Relative Response Ratios for an African American Identity Relative to a White Identity

	$h = 7$	$h = 10$	$h = 15$
	Panel A: 30 days around treatment		
ATT	0.140**	0.094*	0.060
95% Confidence Interval	(0.034, 0.238)	(-0.004, 0.151)	(-0.011, 0.101)
Number of Observations	724	724	724
	Panel B: 45 days around treatment		
ATT	0.191***	0.140***	0.087***
95% Confidence Interval	(0.086, 0.307)	(0.063, 0.244)	(0.041, 0.175)
Number of Observations	1139	1139	1139
	Panel C: 60 days around treatment		
ATT	0.128**	0.106**	0.070**
95% Confidence Interval	(0.017, 0.246)	(0.021, 0.212)	(0.007, 0.155)
Number of Observations	1561	1561	1561

Notes: 1) ATT stands for the Average Treatment Effect on the Treated. 2) h is the smoothing parameter of the weighted logit, see the text. 3) 95% bootstrapped bias-corrected confidence intervals in parentheses. 4) ***, **, and * denote significance at the 1%, 5%, and 10% levels.

To evaluate whether our estimated treatment effects are robust to uncertainty introduced by the first-stage estimation of state-by-day discrimination coefficients, we implement a bootstrap procedure that accounts for the two-stage structure of the staggered Difference-in-Differences design. The procedure is designed to propagate first-stage estimation uncertainty through to the second stage and ensure valid inference.

We generate each bootstrap replicate by resampling states with replacement. For each selected state, we retain all observations across the full sample period following its implementation of a moratorium. We then re-estimate the first-stage logit regressions using the original kernel-weighted procedure to recover smoothed estimates of daily relative response ratios. These estimates are subsequently used as inputs to the second-stage Callaway–Sant'Anna estimator to compute the Average Treatment Effect on the Treated (ATT).

This process is repeated 1,000 times, each time drawing a new bootstrap sample at the state level. From the resulting distribution of ATT estimates, we compute bias-corrected point estimates and confidence intervals.

Table 12 presents the bootstrap results for African American identities. The estimated treatment effects remain consistently positive and statistically significant in most specifications. These results confirm that our findings are robust to potential estimation error introduced by the first-stage discrimination coefficients.

8 Mechanisms and Heterogeneity

Having established that racial disparities in inquiry responses widened during eviction moratoria, we turn to the mechanisms driving this effect. The analysis in this section evaluates whether limits on contract enforceability, specifically the suspension of eviction, shifted property managers' screening decisions toward informal mechanisms. We examine heterogeneity by applicant gender, proxies for discriminatory preferences, tenant-protection laws, political affiliation, and neighborhood characteristics. These dimensions allow us to assess whether the observed patterns are more consistent with taste-based discrimination or with risk-based screening, and to evaluate the role of enforcement frictions in shaping discriminatory behavior.

8.1 Gender Discrimination

Table 13: Gender Heterogeneity in the Effect of Moratorium Expiration

	Males			Females		
	$h = 7$	$h = 10$	$h = 15$	$h = 7$	$h = 10$	$h = 15$
Panel A: 30 days around treatment						
ATT	0.111***	0.076***	0.048**	0.092***	0.050*	0.028
	(0.053, 0.170)	(0.024, 0.128)	(0.003, 0.093)	(0.031, 0.153)	(-0.008, 0.108)	(-0.022, 0.079)
Number of Observations	724	724	724	724	724	724
Panel B: 45 days around treatment						
ATT	0.164***	0.120***	0.075***	0.136***	0.090***	0.053**
	(0.102, 0.226)	(0.069, 0.172)	(0.033, 0.118)	(0.078, 0.193)	(0.038, 0.141)	(0.010, 0.095)
Number of Observations	1139	1139	1139	1139	1139	1139
Panel C: 60 days around treatment						
ATT	0.126***	0.099***	0.065***	0.102***	0.071***	0.043**
	(0.057, 0.194)	(0.045, 0.154)	(0.021, 0.109)	(0.047, 0.157)	(0.025, 0.117)	(0.006, 0.081)
Number of Observations	1561	1561	1561	1561	1561	1561

Notes: 1) ATT stands for the Average Treatment Effect on the Treated. 2) h is the smoothing parameter of the weighted logit, see the text. 3) 95% confidence intervals in parentheses. 4) ***, **, and * denote significance at the 1%, 5%, and 10% levels.

Table 13 reports estimates separately by applicant gender. Across specifications, the estimated effects of moratorium expiration are positive for both male and female applicants, indicating higher relative response rates once enforcement is restored. Although the differences across genders are not statistically significant, point estimates are consistently larger for male applicants, particularly in specifications with wider event windows.

Consistent with prior evidence that discriminatory screening in housing markets is more severe for men of color, we focus on male applicants in the remainder of the paper.

8.2 Statistical or Taste-Based Discrimination?

Table 14: Staggered DiD Estimates of the Effect of the End of Moratorium on the Natural Logarithm of the Relative Response Ratio for African Americans Relative to White Applicants by Pre-Pandemic Rent

	Rent Above Median			Rent Below Median		
	$h = 7$	$h = 10$	$h = 15$	$h = 7$	$h = 10$	$h = 15$
Panel A: 30 days around treatment						
ATT	0.186***	0.154***	0.097***	0.053**	0.021	0.008
	(0.076, 0.296)	(0.069, 0.239)	(0.035, 0.158)	(0.004, 0.102)	(-0.016, 0.059)	(-0.025, 0.040)
Number of Observations	197	197	197	491	491	491
Panel B: 45 days around treatment						
ATT	0.186***	0.154***	0.097***	0.130***	0.100***	0.065***
	(0.076, 0.296)	(0.069, 0.239)	(0.035, 0.158)	(0.073, 0.186)	(0.056, 0.143)	(0.030, 0.100)
Number of Observations	263	263	263	808	808	808
Panel C: 60 days around treatment						
ATT	0.186***	0.154***	0.097***	0.096***	0.081***	0.054***
	(0.076, 0.296)	(0.069, 0.239)	(0.035, 0.158)	(0.040, 0.153)	(0.037, 0.124)	(0.019, 0.089)
Number of Observations	312	312	312	1109	1109	1109

Notes: 1) ATT stands for the Average Treatment Effect on the Treated. 2) h is the smoothing parameter of the weighted logit, see the text. 3) 95% confidence intervals in parentheses. 4) ***, **, and * denote significance at the 1%, 5%, and 10% levels.

We next turn to the first test of the mechanism underlying the increase in discrimination documented above. The theoretical framework delivers a distinction between statistical and taste-based discrimination. Under statistical discrimination, the tightening of screening following moratorium expiration should scale with the economic value of the unit, $\bar{V}$, which is increasing in rent (equation (7)). By contrast, under taste-based discrimination, the impact of enforcement constraints on screening need not vary systematically with $\bar{V}$ (equation (6)).

We test these predictions by splitting the sample based on pre-pandemic rents from

Zillow and re-estimating our preferred staggered DiD specification separately for above- and below-median rent markets. Table 14 reports the results for African American applicants.

We find suggestive evidence that treatment effects are larger in higher-rent markets, as implied by statistical discrimination. The difference is concentrated in the 30-day window around moratorium expiration, where point estimates in above-median rent markets exceed those in below-median markets across bandwidths. At the 45- and 60-day windows, estimates for the two groups are of similar magnitude with overlapping confidence intervals. This pattern is consistent with the prediction that weakened enforcement induces the strongest screening responses where the financial stakes are highest.

Taken together, this first mechanism test suggests that the increase in discrimination during eviction moratoria reflects statistically driven screening responses rather than prejudice.

8.3 Racial Animus Index

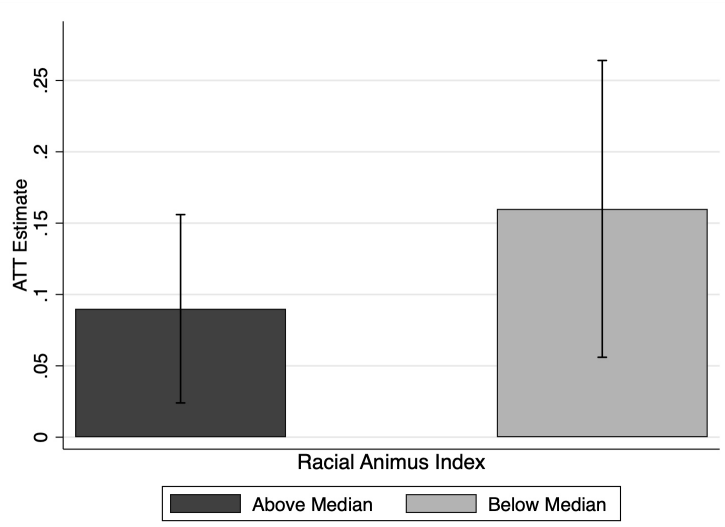


Figure 6: Discriminatory Screening Does Not Increase with the Racial Animus Index

As a direct test of the taste-based discrimination mechanism, we examine whether the response to eviction moratoria varies with measures of racial animus. While the rent-based heterogeneity analysis isolates the role of economic incentives, this exercise instead probes variation in discriminatory preferences across locations.

Specifically, we use the Racial Animus Index developed by Stephens-Davidowitz (2013), which aggregates racially charged internet search activity to capture latent prejudice at the state level. Prior work shows that this index predicts a range of racially discriminatory behaviors, even after controlling for observable socioeconomic characteristics.

Figure 6 plots estimates separately for states with above and below median values of the animus index, with corresponding estimates reported in Appendix Table A7. The estimates for the two groups are statistically indistinguishable over most of the event window. The only exception is the 30 days surrounding treatment, where point estimates are larger in below-median animus states, the opposite of the taste-based prediction. If enforcement constraints amplified pre-existing prejudicial preferences, the screening response should be concentrated where measured animus is highest. Instead, discriminatory screening rises in states with widely different levels of measured racial animus, consistent with statistical discrimination in which screening responds to group-specific beliefs about default risk rather than to local prejudice.

Taken together with the rent-based heterogeneity results, these patterns are more consistent with statistical than with taste-based discrimination. Treatment effects do not vary systematically with measured racial animus, as a preference-based channel would predict, while point estimates are, if anything, larger in higher-rent markets, where the financial stakes of a bad tenant match are greater. We view this evidence as suggestive of statistical discrimination rather than definitive, because the two mechanisms are not separately identified in the long run. Taste-based discrimination can itself produce the belief differences on which statistical discrimination operates: by restricting market access, it delays housing search, limits neighborhood choice, and weakens access to jobs and amenities for minority renters. Over time, these constraints generate real disparities in the outcomes landlords observe, which feed back into beliefs about default risk and can make statistical discrimination self-reinforcing (Phelps, 1972).

8.4 Tenant Protections

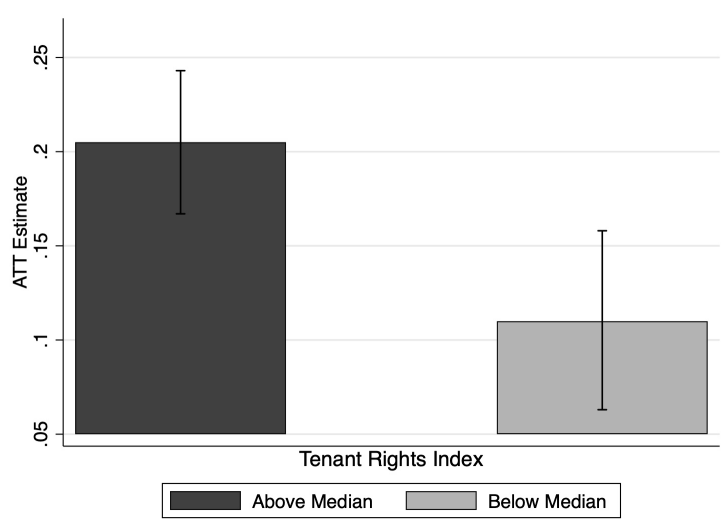


Figure 7: Estimates for States with Tenant Rights Index Above and Below Median

We next examine whether the response to eviction moratoria varies systematically with baseline legal constraints on landlords' ability to enforce rental contracts. To proxy for baseline constraints on formal enforcement, we use the tenant rights protection index developed by Coulson, Le, Ortego-Marti, and Shen (2025), which summarizes state-level legal provisions governing eviction procedures and tenant rights and is described in Section 3.

Figure 7 and Table A8 show that the increase in discriminatory screening during eviction moratoria is larger in states with higher values of this index (the estimate on the left corresponds to a higher value of the index). At first glance, this pattern may appear counterintuitive, as tenant-protective legal environments are often associated with reduced landlord discretion. However, this heterogeneity is consistent with a contractual-enforcement mechanism. In states with stronger tenant protections, landlords already operate with limited scope for formal eviction and longer expected enforcement timelines. In tenant-protective environments, moratoria remove the last remaining formal enforcement option, forcing landlords to rely almost exclusively on ex ante screening. We interpret this pattern as suggestive evidence that discrimination operates as a substitute for weakened contract enforcement.

8.5 Political Affiliation

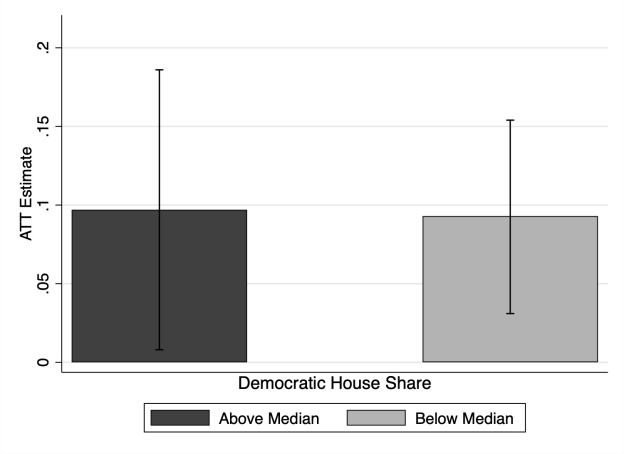


Figure 8: Heterogeneity by Political Affiliation

If the increase in discriminatory screening under eviction moratoria reflected partisan differences in the regulatory environment, such as the stringency of fair-housing oversight, the treatment effect should differ across states with different political compositions. We examine this possibility by splitting states at the median Republican seat share in the lower legislative chamber (Section 3) and re-estimating our staggered Difference-in-Differences specification separately for the two groups. Figure 8 reports the average treatment effects on the treated for the log relative response ratio in each subsample.

The estimated effects are similar across partisan environments. In the specification shown in Figure 8, the ATT is 0.093 (95% confidence interval: 0.031 to 0.154) in states with above-median Republican representation and 0.097 (0.008 to 0.186) in states with below-median Republican representation. The same pattern holds across the remaining specifications, and none of the differences between the two groups is statistically significant even at the 10% significance level.

The absence of partisan heterogeneity carries two messages. First, the screening response to restored eviction enforcement is broad-based rather than confined to a particular regulatory environment, consistent with property managers responding to the change in enforcement risk itself rather than to politically determined differences in fair-housing oversight. Second, the result complements the analysis in Section 5.2: political composition strongly predicts *when* states ended their moratoria, but it does not predict the *size* of the screening response to expiration. Treatment timing correlated with politics combined with treatment effects that do not is precisely the pattern expected when the timing variation is benign, reinforcing our identification strategy.

8.6 Neighborhood Characteristics

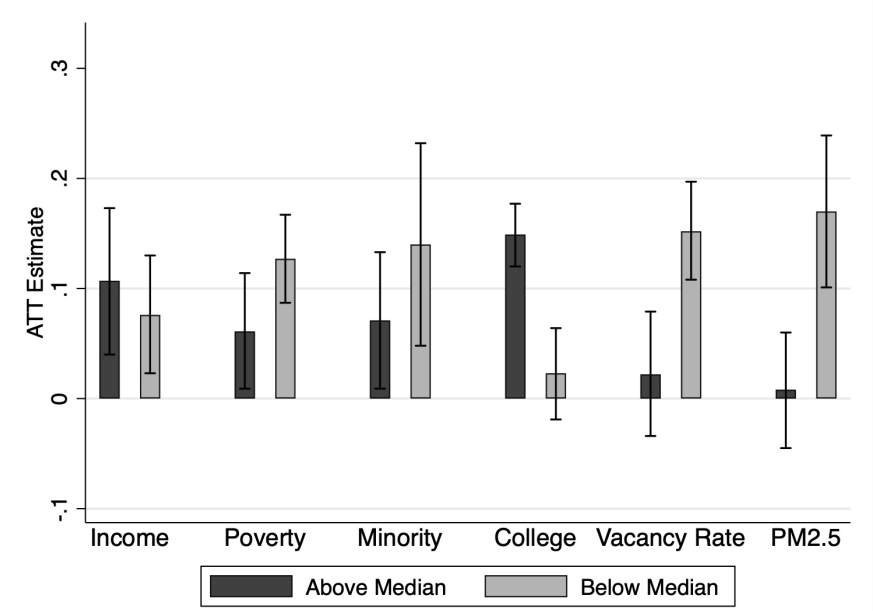


Figure 9: Heterogeneity by Neighborhood Characteristics

This section examines whether the increase in discriminatory screening during the eviction moratoria varies systematically with neighborhood characteristics. This heterogeneity by neighborhood characteristics reveals where the exclusion induced by weakened enforcement is concentrated, and therefore which prospective tenants bear its cost.

Figure 9 reports average treatment effects estimated separately for subsamples above and below the median of six pre-pandemic characteristics: median household income, poverty rate, minority population share, college-educated share, rental vacancy rate, and ambient air pollution, measured by fine particulate matter ($PM_{2.5}$) concentrations. The figure displays our preferred specification (60 days around treatment with $h = 10$); the remaining specifications display the same pattern.

The effects divide sharply along amenity and market-tightness lines. The increase in discriminatory screening is concentrated in areas with a more educated population, cleaner air, and tighter rental markets. In these subsamples, the estimated effects exceed the full-sample benchmark by half or more, while in the complementary subsamples they are small and statistically indistinguishable from zero; the differences are significant at the 1 percent level. To fix magnitudes, the estimate in low-pollution areas is 0.170, compared with 0.008 in high-pollution areas. The split by poverty rate displays the same asymmetry, with larger effects in low-poverty areas, although the difference is statistically significant only in

the shortest event window. In contrast, the effects do not vary significantly with median household income or with the minority population share.

These patterns speak to the mechanism. In the model, the incentive to screen selectively when enforcement is suspended scales with the option value of a vacant unit, which is greatest where rental demand is strong and units re-lease quickly. The concentration of the response in educated, environmentally advantaged, and tight rental markets is consistent with this option-value channel. The air-quality split carries an additional implication for environmental justice: suspending enforcement raised barriers to entry precisely into the neighborhoods with the greatest amenity value, steepening the gradient of access to environmental quality for minority renters.

9 Consequences of Increased Discrimination

The preceding analysis shows that property managers changed their screening behavior when eviction enforcement was suspended. An important question is whether these screening decisions have real consequences for housing allocation and market outcomes. In this section, we examine two implications. First, we test whether differential responses in the correspondence experiment translate into systematic differences in who ultimately rents the units. Second, we evaluate the financial consequences of discriminatory screening for property owners, who rely on property managers as delegated agents. Together, these analyses link the screening behavior documented above to both allocation outcomes in the rental market and the economic returns to housing assets.

9.1 Housing Outcomes

To examine whether the discrimination observed in our correspondence experiment affects renters' subsequent housing outcomes, we link the experimental sample to the residential historical InfoUSA database, compiled by Data Axle. This dataset covers approximately 309 million individuals (effectively the entire U.S. population) and contains address histories along with demographic and economic attributes, including age, gender, race/ethnicity, marital status, number of children, estimated wealth and income, and length of residence at the current address. We locate 87.3% of the addresses from the rental listings in the experiment in the addresses from the InfoUSA database. Appendix D details the procedure.

Using information on race/ethnicity and residential tenure from InfoUSA, we identify newly moved-in renters with varied races including white, African-American, and Hispanic. Specifically, we exclude homeowners and renters who moved in before the inquiry date from

the experiment or renters of "Other" race, yielding a final sample of 3,326 matched addresses. Appendix Table A10 reports the distribution of tenant race at these addresses. African American and Hispanic renters each account for roughly 12 percent of the sample, white renters comprise 55 percent, and 6.6 percent are of other races (e.g., Asian). The final 14 percent of matched addresses contain multiple renter households of different races (e.g., both white and African American families). These addresses are typically multi-family properties that advertise one address but offer multiple rental units for lease.

To examine whether actual move-in patterns are systematically aligned with the racial identity of the fictitious renter in the experiment, we implement an outcome test in the spirit of Becker (1957, 1993). Specifically, we construct the variable Same Race Share, which for each inquiry equals the fraction of tenants at the matched address whose race coincides with that of the fictitious renter. In most cases the measure takes values of zero or one, but for addresses with multiple renter households of different races, it can take intermediate values reflecting the share of tenants of the same race as the inquiry identity.

To assess whether property managers' screening decisions from the experiment affect the observed racial composition of tenants, we estimate the following equation for each state-day cell:

$$\text{Share Same Race}_{ijsm} = \alpha + \beta \text{Response}_{ijst} + X_j' \theta_{s\tau} + \varepsilon_{ijst},$$

where i indexes properties, j fictitious identities, s states, and t dates, consistent with (8). The left-hand side variable $\text{Share Same Race}_{ijsm}$ is measured at the subsequent move-in period $m > t$, which occurs after the inquiry date t. The parameter β describes the effect of receiving a response from a property manager on the inquiry of a particular racial identity on the likelihood of an actual individual with that identity showing up in the apartment.

We estimate this specification using the quasi-maximum likelihood fractional logit procedure of Papke and Wooldridge (1996), weighted as in (9), to accommodate the bounded fractional outcome.

We use the estimated coefficients to predict the share of same-race tenants under two counterfactuals: when a manager responds to an inquiry ($\text{Response}_{ijst} = 1$) and when no response is sent ($\text{Response}_{ijst} = 0$). We then construct the natural logarithm of the relative response ratio for each state s and date t:

$$\ln(RR_{st}) = \ln(\text{Share Same Race}_{st} \big| \text{Response} = 0) - \ln(\text{Share Same Race}_{st} \big| \text{Response} = 1).$$

We take the natural log transformation and clamp predicted shares to the interval $(0.001, 0.999)$ to mitigate instability from extreme values. As a robustness check, we also re-estimate the specification using an alternative dependent variable—the risk difference, defined as the difference in predicted same-race shares between responses and non-responses,

(Share Same Race$_{st}$|Response = 0) − (Share Same Race$_{st}$|Response = 1), and obtain similar results, reported in Table A11 and Figure A3 in the Appendix.

In the second stage, we apply the staggered difference-in-differences estimator of Callaway and Sant'Anna (2021) to the $\ln(RR_{st})$ series, exploiting the staggered expiration of eviction moratoria across states. This design allows us to test whether the likelihood that tenants match on race worsened during moratoria due to screening behavior.

Table 15: Staggered DiD Estimates of the Effect of the End of Moratorium on the Natural Logarithm of the Ratio of the Predicted Share of Same Race after No Response to the Predicted Share of Same Race after a Response

	$h = 7$	$h = 10$	$h = 15$
	Panel A: 30 days around treatment		
ATT	0.547***	0.515***	0.347***
	(0.281, 0.812)	(0.275, 0.756)	(0.209, 0.485)
Number of Observations	709	719	737
	Panel B: 45 days around treatment		
ATT	0.482***	0.414***	0.257***
	(0.259, 0.705)	(0.228, 0.599)	(0.142, 0.372)
Number of Observations	1096	1137	1166
	Panel C: 60 days around treatment		
ATT	0.763***	0.559***	0.298***
	(0.498, 1.029)	(0.389, 0.729)	(0.199, 0.397)
Number of Observations	1485	1539	1598

Notes: 1) ATT stands for the Average Treatment Effect on the Treated. 2) h is the smoothing parameter of the weight in the quasi-maximum likelihood procedure with the logit link, proposed by Papke and Wooldridge (1996), see the text. 3) 95% confidence intervals in parentheses. 4) ***, **, and * denote significance at the 1%, 5%, and 10% levels.

Table 15 and Figure 10 present the results. We find that the expiration of eviction moratoria significantly increased the relative response ratio, indicating that racial disparities in property managers' screening decisions narrowed once enforcement authority was restored. The magnitudes are economically meaningful: for example, the estimate of 0.414 implies that the relative response ratio rose by approximately 51 percent ($100 \times (e^{0.414} - 1)$) following the end of a moratorium. In other words, property managers' non-responses during the moratorium disproportionately reduced the likelihood of same-race matches, and this gap diminished once enforcement constraints were lifted.

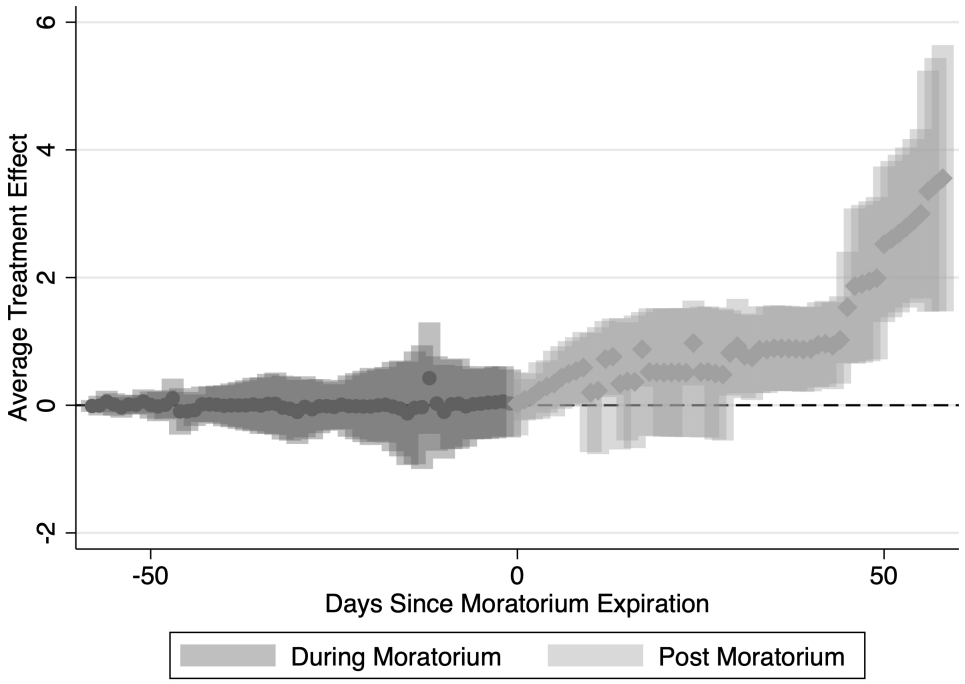


Figure 10: Event Study Coefficients from Callaway and Sant'Anna (2021)'s Estimator for the Natural Logarithm of the Relative Response Ratios for the Same Race without a Response Relative to the Share of Same Race with a Response with the Smoothing Parameter $h = 10$ and $\hat{\tau} = 60$ Days around Treatment

These findings from the outcome test align closely with our benchmark test based on response rates. As emphasized by Gaebler and Goel (2025), when both benchmark and outcome tests point in the same direction, the conclusion of discrimination is statistically robust under mild assumptions. In our setting, the convergence of evidence across both tests implies that property managers not only differentially screened minority applicants during eviction moratoria but that these decisions also translated into systematically different rental outcomes.

9.2 Investor Outcomes

What are the consequences of discriminatory screening for the investors who delegate leasing operations to property managers? If the discrimination we document reflected prejudice, the answer would follow Becker (1957): taste-based discriminators forgo profitable matches, so investors would have lost money whenever their agents screened on race. The evidence instead points to statistical discrimination. Property managers responded optimally to their assessment of the risks and opportunities of leasing a rental unit while eviction was unen-

forceable. If they believed that the risk-adjusted payoff of leasing to an African American applicant was lower than that of leasing to a white applicant, they optimally steered away from African American applicants.[7] Optimal, however, is not legal. Disparate treatment by race in rental housing violates the Fair Housing Act even when justified by a business purpose: the accuracy of the underlying risk assessment is not a defense. The screening therefore produced two things at once: a risk-management benefit for the owner and an exposure to legal sanction. Whether it paid off is a benefit–cost question. This section values both sides of that ledger.

What were the benefits of the discriminatory steering? During the moratoria, the overall responsiveness of property managers did not change, but they shifted responses toward white applicants at the expense of African American applicants (Section 4.4), reducing the response probability to an African American inquiry by 6.9 percentage points.[8] To value this shift, consider the representative property manager from our data and industry benchmarks: 100 units under management, 50 listings per year, 22-30 inquiries per listing, an inquiry-to-lease conversion rate of 0.6 percent, and a monthly rent of $2,000.[9] The moratoria lasted 68.5 days on average, affecting $50 \times 68.5/365 = 9.4$ listings of this manager. Of the inquiries each listing receives, 19.3 percent come from African American applicants if inquiries mirror the racial composition of renter households.[10] Assume the manager assesses default risk using the best contemporaneous public data, the Census Household Pulse Survey: over the April–July 2020 collection waves, the household-weighted rent delinquency rate of African American renter households was 25.3 percent, against 12.2 percent for white renter households, a difference of

[7]Landlords appear to make the same calculation at the other end of the tenancy: Lodermeier (2024) finds that landlords tolerate 4.9 percent less back rent from minority tenants than from white tenants with the same landlord and contract rent before filing an eviction, a pattern consistent with accurate statistical discrimination under a marginal outcome test.

[8]During the moratoria, the relative response ratio was 0.911 and white applicants received responses with probability 0.598. The preferred ATT of 0.120 (Table 3) implies a counterfactual ratio of $0.911 \times e^{0.120} = 1.027$ and hence a counterfactual response probability of $1.027 \times 0.598 = 0.614$ for African American applicants, against 0.545 observed, a difference of $0.614 - 0.545 = 0.069$. Capping the counterfactual ratio at parity yields 5.3 percentage points and proportionally smaller benefits.

[9]The inquiry volume is based on industry reports summarized in ShowMojo's Q2 2025 market overview (https://hello.showmojo.com/blog/the-rental-market-in-q2-2025-its-a-great-time-to-double-down-on-efficiency/). The inquiry-to-lease conversion rate is drawn from industry benchmarks reported by Follow Up Boss (https://www.followupboss.com/blog/real-estate-lead-conversion-rate). The relisting rate is drawn from turnover statistics reported by Nspire Experts (https://nspireexperts.com/apartment-turnover-rate/); the number of units per manager from industry statistics compiled by DoorLoop (https://www.doorloop.com/blog/property-management-industry-stats).

[10]The U.S. Census Bureau reports that 4.6 million Black-alone renter households were cost-burdened, which are 56.2 percent of all Black-alone renter households in 2023, implying $4.6/0.562 = 8.2$ million Black-alone renter households among the 42.5 million renter households, which is a 19.3 percent share (U.S. Census Bureau, "Nearly Half of Renter Households Are Cost-Burdened, Proportions Differ by Race," CB24-150, September 12, 2024; 2023 American Community Survey 1-year estimates).

13.1 percentage points. Finally, suppose that a defaulting tenant ultimately costs the owner six months of rent, $6 \times \$2{,}000 = \$12{,}000$.[11] The expected benefit of the discriminatory screening per manager over a moratorium is then

$$\underbrace{9.4}_{\text{listings}} \times \underbrace{22}_{\substack{\text{inquiries}\\ \text{per listing}}} \times \underbrace{0.193}_{\substack{\text{AA share}\\ \text{of inquiries}}} \times \underbrace{0.069}_{\substack{\text{withheld}\\ \text{responses}}} \times \underbrace{0.006}_{\text{conversion}} \times \underbrace{0.131}_{\substack{\Delta \text{ default}\\ \text{risk}}} \times \underbrace{\$12{,}000}_{\substack{\text{loss per}\\ \text{default}}} \approx \$26,$$

and at most \$52 taking the upper end of the inquiry-volume range and the upper confidence bound of the ATT. Twenty-six dollars per moratorium is the benefit side of the ledger.

What were the costs? The screening violated the Fair Housing Act, whose penalty schedule in effect in 2020 set civil penalties of \$21,410 for a first proven violation, \$53,524 with one prior violation, and \$107,050 with two or more, in addition to uncapped compensatory and punitive damages in private actions and pattern-or-practice liability in suits brought by the Department of Justice.[12] These penalties bind only when violations are detected: the cost a property manager expects to pay is the fine multiplied by the probability of a complaint and a finding of liability, and that probability was small. In 2020, 4,821 race-based housing discrimination complaints were filed nationwide across all reporting channels, against 42.5 million renter households, which is a complaint intensity of $4{,}821/42{,}500{,}000 = 0.0113$ percent per household-year.[13] Treating every one of these complaints as a proven first violation, the expected penalty for a moratorium of discriminatory screening was $(4{,}821/42{,}500{,}000) \times \$21{,}410 = \$2.43$, against a benefit of \$26: a net expected gain of roughly \$24. Equivalently, the detection probability at which the expected penalty would offset the benefit is $\$26/\$21{,}410 = 0.121$ percent, roughly one in eight hundred, which is an order of magnitude above the observed complaint intensity.

What was the outcome of the discriminatory screening? A screen that earns \$24 persists, and that is what we observe: the tilt toward white applicants lasted throughout the moratoria and reversed only when eviction enforcement, rather than fair-housing enforcement, was restored. The cost fell on African American applicants, whose reduced access at the inquiry stage carried through to realized move-ins (Section 9.1). The same arithmetic points to the

[11] Ledger-level evidence from low-income rental markets puts the components of this loss at 3.5 months of unpaid rent owed by evicted tenants at move-out, direct eviction filing costs equivalent to two to three months of rent, and roughly two further months of vacancy before the unit is re-let (Humphries et al., 2025); six months is, if anything, below their sum.

[12] 85 FR 13041 (effective April 6, 2020), amending 24 C.F.R. §180.671.

[13] National Fair Housing Alliance, 2021 Fair Housing Trends Report, pp. 6–11. Of the 28,712 complaints across all protected classes, 21,089 were processed by private fair housing organizations, 1,697 by the U.S. Department of Housing and Urban Development (HUD), 5,883 by state and local Fair Housing Assistance Program agencies, and 43 by the Department of Justice; race accounted for 16.79 percent. HUD alone processed 455 race-based complaints.

policy lever: pairing eviction moratoria with expanded fair-housing testing and enforcement, raising the detection probability from 0.0113 percent toward the 0.121 percent break-even, would remove the private return to the screening that the moratoria induced.

10 Conclusion

This paper asks how the enforceability of contracts shapes the way intermediaries screen counterparties, and answers it in a market where enforcement was switched off and back on: U.S. rental housing during the 2020 state eviction moratoria. Combining a hand-collected legal record of the moratoria with a nationwide correspondence experiment and exploiting the staggered restoration of enforcement, we find that discriminatory screening intensified while eviction was unenforceable. Relative response rates for African American applicants rose by 12.8 percent when moratoria expired, and the effect is robust across event windows, control groups, alternative staggered difference-in-differences estimators, and inference that accounts for the two-stage estimation procedure.

Two features of the evidence discipline its interpretation. First, enforcement changed the composition of screening, not its level: property managers responded to as many inquiries during the moratoria as after them, but redirected responses from African American toward white applicants while eviction was suspended. The moratoria reallocated access to vacant units rather than rationing it, a margin invisible in the rent data on which evaluations of eviction policy have focused. Second, the screening mattered: linking listings to residential address histories, non-responses during moratoria translated into systematically different move-in patterns, and the agreement of the benchmark and outcome tests makes the finding of discrimination robust under mild assumptions (Gaebler and Goel, 2025).

The mechanism is statistical rather than taste-based discrimination. The screening response is, if anything, stronger in high-rent markets, where the cost of a bad match is largest; it does not vary with measured racial animus; and it is strongest in states whose tenant protections had already narrowed landlords' enforcement options, so that the moratoria removed the last formal remedy. When the legal technology for managing default risk ex post was suspended, property managers substituted the technology available ex ante: screening on a characteristic they believed predicted default. Consistent with this interpretation, the estimates are unchanged during the federal CARES Act moratorium, which created a statutory restriction without an enforcement mechanism in state courts: screening responds to the credibility of enforcement, not to the statute. However, disparate treatment by race violates the Fair Housing Act regardless of its actuarial content.

Why did an illegal practice persist for the length of the moratoria? Because one enforce-

ment margin was suspended while the other never bound. Our benefit–cost accounting puts the expected risk-management value of the discriminatory tilt at roughly $26 per property manager per moratorium, against an expected fair-housing penalty of $2.43: race-based complaints were filed at 0.0113 percent per renter household-year, an order of magnitude below the 0.121 percent detection probability at which the first-violation penalty would deter. The moratoria thus created an asymmetry between two enforcement systems, the eviction courts that protect owners and the fair-housing apparatus that protects applicants, and property managers optimized against both.

These findings reframe the welfare evaluation of eviction policy. For incumbent tenants, moratoria provide insurance against displacement; for prospective tenants, they raise the screen at the market's front door, and the burden of that screen fell on African American applicants, concentrated in high-amenity neighborhoods. A welfare accounting must therefore weigh incumbent protection against market access, as in the equilibrium analyses of Abramson (2025) and Corbae, Glover, and Nattinger (2024), and our evidence adds a distributional cost those frameworks do not yet price. The policy implication follows from the enforcement asymmetry: if eviction protections are deployed in future crises, pairing them with scaled-up fair-housing testing and enforcement, or with inquiry platforms that withhold identity signals until later in the matching process, would close the gap between the private and social returns to screening. Quantifying the full trade-off in a structural model that joins tenant insurance, landlord learning, search, and discrimination is a natural next step.

The broader lesson is that contract enforcement shapes markets twice: ex post, through the remedies it affords, and ex ante, through the screening it makes unnecessary. When enforcement was withdrawn, screening intensified along the margin that was cheapest and least monitored, and it fell on the applicants the law is designed to protect. Evaluations of policies that suspend creditor or owner remedies, from mandated mortgage forbearance to debtor-friendly bankruptcy exemptions, should anticipate the same substitution.

References

Abramson, B. (2025): "The Equilibrium Effects of Eviction Policies," The Journal of Finance, Forthcoming.

Adelino, M., B. Wei, and F. Zhao (2025): "Screen More, Sell Later: Screening and Dynamic Signaling in the Mortgage Market," Working Paper 2025-3, Federal Reserve Bank of Atlanta, fRB Atlanta Working Paper No. 2025-3.

Ambrose, B. W., J. N. Conklin, and L. A. Lopez (2020): "Does Borrower and Broker Race Affect the Cost of Mortgage Credit?" The Review of Financial Studies, 34, 790–826.

Arellano, C., Y. Bai, and J. Zhang (2012): "Firm dynamics and financial development," Journal of Monetary Economics, 59, 533–549.

Baker, A., B. Callaway, S. Cunningham, A. Goodman-Bacon, and P. H. C. Sant'Anna (2025): "Difference-in-Differences Designs: A Practitioner's Guide," .

Bartlett, R., A. Morse, R. Stanton, and N. Wallace (2022): "Consumer-lending discrimination in the FinTech Era," Journal of Financial Economics, 143, 30–56.

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

——— (1968): "Crime and Punishment: An Economic Approach," Journal of Political Economy, 76, 169–217.

——— (1993): "Nobel Lecture: The Economic Way of Looking at Behavior," Journal of Political Economy, 101, 385–409.

Benfer, E. A., R. Koehler, A. Mark, V. Nazzaro, A. K. Alexander, P. Hepburn, D. E. Keene, and M. Desmond (2023): "COVID-19 Housing Policy: State and Federal Eviction Moratoria and Supportive Measures in the United States During the Pandemic," Housing Policy Debate, 33, 1390–1414.

Benfer, E. A., D. Vlahov, M. Y. Long, E. Walker-Wells, J. Pottenger, G. Gonsalves, and D. E. Keene (2021): "Eviction, health inequity, and the spread of COVID-19: housing policy as a primary pandemic mitigation strategy," Journal of Urban Health, 98, 1–12.

Borusyak, K., X. Jaravel, and J. Spiess (2024): "Revisiting Event-Study Designs: Robust and Efficient Estimation," The Review of Economic Studies, 91, 3253–3285.

Cai, Z., J. Fan, and R. Li (2000): "Efficient Estimation and Inferences for Varying-Coefficient Models," Journal of the American Statistical Association, 95, 888–902.

Callaway, B. and P. H. Sant'Anna (2021): "Difference-in-differences with multiple time periods," Journal of Econometrics, 225, 200–230.

Christensen, P., I. Sarmiento-Barbieri, and C. Timmins (2021): "Racial discrimination and housing outcomes in the United States rental market," Tech. rep., National Bureau of Economic Research.

——— (2022): "Housing Discrimination and the Toxics Exposure Gap in the United States: Evidence from the Rental Market," The Review of Economics and Statistics, 104, 807–818.

Collinson, R., J. E. Humphries, N. Mader, D. Reed, D. Tannenbaum, and W. van Dijk (2023): "Eviction and Poverty in American Cities*," The Quarterly Journal of Economics, 139, 57–120.

Corbae, D., A. Glover, and M. Nattinger (2024): "Equilibrium Evictions," NBER Working Paper 32898.

Coulson, N. E., T. Le, V. Ortego-Marti, and L. Shen (2025): "Tenant rights, eviction, and rent affordability," Journal of Urban Economics, 147, 103762.

Dagher, J. and Y. Sun (2016): "Borrower protection and the supply of credit: Evidence from foreclosure laws," Journal of Financial Economics, 121, 195–209.

de Chaisemartin, C. and X. D'Haultfœuille (2020): "Two-Way Fixed Effects Estimators with Heterogeneous Treatment Effects," American Economic Review, 110, 2964–96.

Desmond, M. (2016): Evicted: Poverty and Profit in the American City, Crown Publishers.

Diamond, D. W. (1984): "Financial Intermediation and Delegated Monitoring," The Review of Economic Studies, 51, 393–414.

FHFA Office of Inspector General (2022): "Oversight of Multifamily Borrowers' Compliance with CARES Act and Freddie Mac Tenant Protections," .

Frame, W. S., R. Huang, E. J. Mayer, and A. Sunderam (2023): "The Impact of Minority Representation at Mortgage Lenders," NBER Working Paper No. w30125, Available at SSRN: https://ssrn.com/abstract=4134921.

Fryer Jr, R. G. and S. D. Levitt (2004): "The causes and consequences of distinctively black names," The Quarterly Journal of Economics, 119, 767–805.

Fuster, A., P. Goldsmith-Pinkham, T. Ramadorai, and A. Walther (2022): "Predictably Unequal? The Effects of Machine Learning on Credit Markets," The Journal of Finance, 77, 5–47.

Gaebler, J. D. and S. Goel (2025): "A simple, statistically robust test of discrimination," Proceedings of the National Academy of Sciences, 122, e2416348122.

Galster, G. and E. Godfrey (2005): "By Words and Deeds: Racial Steering By Real Estate Agents in the U.S. in 2000," Journal of the American Planning Association, 71, 251–268.

Garriga, C. and A. Hedlund (2020): "Mortgage Debt, Consumption, and Illiquid Housing Markets in the Great Recession," American Economic Review, 110, 1603–34.

Ghent, A. C., R. Hernández-Murillo, and M. T. Owyang (2014): "Differences in subprime loan pricing across races and neighborhoods," Regional Science and Urban Economics, 48, 199–215.

Goodman-Bacon, A. (2021): "Difference-in-differences with variation in treatment timing," Journal of Econometrics, 225, 254–277, themed Issue: Treatment Effect 1.

Gromis, A., I. Fellows, J. R. Hendrickson, L. Edmonds, L. Leung, A. Porton, and M. Desmond (2022): "Estimating eviction prevalence across the United States," Proceedings of the National Academy of Sciences, 119, e2116169119.

Guryan, J. and K. K. Charles (2013): "Taste-based or statistical discrimination: the economics of discrimination returns to its roots," The Economic Journal, 123, F417–F432.

Hale, T., N. Angrist, R. Goldszmidt, B. Kira, A. Petherick, T. Phillips, S. Webster, E. Cameron-Blake, L. Hallas, S. Majumdar, and H. Tatlow (2021): "A global panel database of pandemic policies (Oxford COVID-19 Government Response Tracker)," .

Hanson, A. and Z. Hawley (2011): "Do Landlords Discriminate in the Rental Housing Market? Evidence From an Internet Field Experiment in US Cities," Journal of Urban Economics, 70, 99–114.

Hatamyar, J. and C. F. Parmeter (2023): "Local Eviction Moratoria and the Spread of COVID-19," .

Humphries, J. E., S. Nelson, D. L. Nguyen, W. van Dijk, and D. Waldinger (2025): "Nonpayment and Eviction in the Rental Housing Market," Working Paper 2024-150, Becker Friedman Institute for Economics, University of Chicago, december 2024, updated June 2025.

Keys, B. J., T. Mukherjee, A. Seru, and V. Vig (2010): "Did Securitization Lead to Lax Screening? Evidence from Subprime Loans," Quarterly Journal of Economics, 125, 307–362.

Lerner, J. and A. Schoar (2005): "Does Legal Enforcement Affect Financial Transactions? The Contractual Channel in Private Equity," The Quarterly Journal of Economics, 120, 223–246.

Li, D., J. Chen, and J. Gao (2011): "Non-parametric time-varying coefficient panel data models with fixed effects," The Econometrics Journal, 14, 387–408.

Lodermeier, A. (2024): "Racial Discrimination in Eviction Filing," Working Paper, last accessed on July 15, 2024.

Papke, L. E. and J. M. Wooldridge (1996): "Econometric methods for fractional response variables with an application to 401(k) plan participation rates," Journal of Applied Econometrics, 11, 619–632.

Pence, K. M. (2006): "Foreclosing on Opportunity: State Laws and Mortgage Credit," The Review of Economics and Statistics, 88, 177–182.

Phelps, E. S. (1972): "The Statistical Theory of Racism and Sexism," The American Economic Review, 62, 659–661.

Quintin, E. (2008): "Limited enforcement and the organization of production," Journal of Macroeconomics, 30, 1222–1245.

Roth, J., P. H. Sant'Anna, A. Bilinski, and J. Poe (2023): "What's trending in difference-in-differences? A synthesis of the recent econometrics literature," Journal of Econometrics, 235, 2218–2244.

Staniswalis, J. G. (1989): "The Kernel Estimate of a Regression Function in Likelihood-Based Models," Journal of the American Statistical Association, 84, 276–283.

Stephens-Davidowitz, S. (2013): "Racial Animus and Voting: Using Google Search Data to Find What Surveys Miss," Working Paper.

Stiglitz, J. E. and A. Weiss (1981): "Credit Rationing in Markets with Imperfect Information," The American Economic Review, 71, 393–410.

Sun, L. and S. Abraham (2021): "Estimating dynamic treatment effects in event studies with heterogeneous treatment effects," Journal of Econometrics, 225, 175–199, themed Issue: Treatment Effect 1.

Tibshirani, R. and T. Hastie (1987): "Local Likelihood Estimation," Journal of the American Statistical Association, 82, 559–567.

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

Yinger, J. (1995): Closed Doors, Opportunities Lost: The Continuing Costs of Housing Discrimination, Russell Sage Foundation.

Appendix

A Proofs

The value of searching for a tenant for an empty unit is

$$\begin{aligned} V &= \mathbb{E}\max\{u_M - \psi, u_W\} = \mathbb{E}[\mathbb{1}_{\{u_M-\psi\geq u_W\}}(u_M-\psi) + \mathbb{1}_{\{u_M-\psi<u_W\}}u_W] = \\ &= P(u_M-\psi\geq u_W)u_M - \mathbb{E}[\psi|u_M-\psi\geq u_W] + (1-P(u_M-\psi\geq u_W))u_W = \\ &= u_W + \Delta u F(\Delta u) - \int_{\psi_{\min}}^{\Delta u}\psi dF(\psi), \end{aligned}$$

where $\Delta u \equiv u_M - u_W$. To simplify, use integration by parts to rewrite the last term as

$$\int_{\psi_{\min}}^{\Delta u}\psi dF(\psi) = \psi F(\psi)|_{\psi_{\min}}^{\Delta u} - \int_{\psi_{\min}}^{\Delta u}F(\psi)d\psi = \Delta u F(\Delta u) - \int_{\psi_{\min}}^{\Delta u}F(\psi)d\psi,$$

where $\lim_{\psi\to\psi_{\min}}\psi F(\psi) = 0$ by assumption. Then the property manager's value of searching for a tenant for an empty unit is

$$V = u_W + \int_{\psi_{\min}}^{\Delta u}F(\psi)d\psi.$$

To derive the difference in the differences of the utilities, calculate the utility from leasing of an applicant i during and before the moratorium from (1) and (4) as

$$\begin{aligned} \bar{u}_i &= \frac{\pi_i R}{1-\beta\pi_i} + \frac{(1-\pi_i)\beta^2}{1-\beta\pi_i}\bar{V} - \frac{((1-\pi_i)\beta+1)}{1-\beta\pi_i}\kappa_i, \\ u_i^* &= \frac{\pi_i R}{1-\beta\pi_i} + \frac{(1-\pi_i)\beta}{1-\beta\pi_i}V^* - \frac{1}{1-\beta\pi_i}\kappa_i. \end{aligned}$$

To derive $\Delta^2 u = (\bar{u}_M - \bar{u}_W) - (u_M^* - u_W^*) = (\bar{u}_M - u_M^*) - (\bar{u}_W - u_W^*)$, we start with calculating $\bar{u}_i - u_i^*$:

$$\bar{u}_i - u_i^* = \beta\frac{(1-\pi_i)}{1-\beta\pi_i}(\beta\bar{V} - V^*) - \frac{\beta(1-\pi_i)}{1-\beta\pi_i}\kappa_i. \tag{15}$$

Then the difference in the differences of the utilities is

$$\begin{aligned} \Delta^2 u &= (\bar{u}_M - u_M^*) - (\bar{u}_W - u_W^*) = \beta(\frac{(1-\pi_M)}{1-\beta\pi_M} - \frac{(1-\pi_W)}{1-\beta\pi_W})(\beta\bar{V} - V^*) - \frac{\beta(1-\pi_M)}{1-\beta\pi_M}\kappa \\ &= \frac{\beta(1-\beta)(\pi_W-\pi_M)}{(1-\beta\pi_M)(1-\beta\pi_W)}(\beta\bar{V} - V^*) - \frac{\beta(1-\pi_M)}{1-\beta\pi_M}\kappa, \end{aligned} \tag{16}$$

where the first equality uses normalization $\kappa_M = \kappa$ and $\kappa_W = 0$. The second equality uses

$$\frac{(1-\pi_M)}{1-\beta\pi_M} - \frac{(1-\pi_W)}{1-\beta\pi_W} = \frac{1-\beta\pi_W-\pi_M+\pi_M\beta\pi_W-1+\beta\pi_M+\pi_W-\pi_W\beta\pi_M}{(1-\beta\pi_M)(1-\beta\pi_W)}$$
$$= \frac{-\beta\pi_W-\pi_M+\beta\pi_M+\pi_W}{(1-\beta\pi_M)(1-\beta\pi_W)} = \frac{(1-\beta)(\pi_W-\pi_M)}{(1-\beta\pi_M)(1-\beta\pi_W)}.$$

To assess how the option value to lease changes, use

$$V = u_W + \int_{\psi_{\min}}^{\Delta u} F(\psi)d\psi,$$
$$\bar{V} - V^* = \bar{u}_W - u_W^* + \int_{\psi_{\min}}^{\Delta\bar{u}} F(\psi)d\psi - \int_{\psi_{\min}}^{\Delta u^*} F(\psi)d\psi.$$

We use the first-order Taylor expansion to approximate $\int_{\psi_{\min}}^{\Delta\bar{u}} F(\psi)d\psi - \int_{\psi_{\min}}^{\Delta u^*} F(\psi)d\psi = (\Delta\bar{u} - \Delta u^*)\cdot F(\Delta u^*) = \Delta^2 u \cdot F(\Delta u^*)$ and $\bar{u}_W - u_W^*$ from (15):

$$\bar{V} - V^* = \beta\frac{(1-\pi_W)}{1-\beta\pi_W}(\beta\bar{V} - V^*) + F(\Delta u^*)\Delta^2 u.$$

Use the relationship above to find the option value to lease during the eviction moratorium:

$$\bar{V}(1 - \frac{\beta^2-\beta^2\pi_W}{1-\beta\pi_W}) = (1 - \frac{\beta-\beta\pi_W}{1-\beta\pi_W})V^* + F(\Delta u^*)\Delta^2 u,$$

where $(1-\beta\pi_W-\beta^2+\beta^2\pi_W) = ((1-\beta^2)-\beta\pi_W(1-\beta)) = (1-\beta)(1+\beta-\beta\pi_W)$. Thus,

$$(1-\beta)(1+\beta(1-\pi_W))\bar{V} = (1-\beta)V^* + (1-\beta\pi_W)F(\Delta u^*)\Delta^2 u,$$
$$V^* = (1+\beta(1-\pi_W))\bar{V} - \frac{1-\beta\pi_W}{1-\beta}F(\Delta u^*)\Delta^2 u.$$

To finish calculation of $\Delta^2 u$ from (16), we need $\beta\bar{V} - V^*$:

$$\beta\bar{V} - V^* = \beta\bar{V} - (1+\beta(1-\pi_W))\bar{V} + \frac{1-\beta\pi_W}{1-\beta}F(\Delta u^*)\Delta^2 u =$$
$$= (\beta - 1 - \beta + \beta\pi_W)\bar{V} + \frac{1-\beta\pi_W}{1-\beta}F(\Delta u^*)\Delta^2 u = -(1-\beta\pi_W)\bar{V} + \frac{1-\beta\pi_W}{1-\beta}F(\Delta u^*)\Delta^2 u.$$

Using the above change in the option value to lease in (16), we get

$$\Delta^2 u = \frac{\beta(1-\beta)(\pi_W - \pi_M)}{(1-\beta\pi_M)(1-\beta\pi_W)}(\frac{1-\beta\pi_W}{1-\beta}F(\Delta u^*)\Delta^2 u - (1-\beta\pi_W)\bar{V}) - \frac{\beta(1-\pi_M)}{1-\beta\pi_M}\kappa,$$

$$(1 - \frac{\beta(1-\beta)(\pi_W - \pi_M)}{(1-\beta\pi_M)(1-\beta\pi_W)}\frac{(1-\beta\pi_W)}{(1-\beta)}F(\Delta u^*))\Delta^2 u =$$

$$= -\frac{\beta(1-\beta)(\pi_W - \pi_M)}{(1-\beta\pi_M)(1-\beta\pi_W)}(1-\beta\pi_W)\bar{V} - \frac{\beta(1-\pi_M)}{1-\beta\pi_M}\kappa,$$

$$(1 - \frac{\beta(\pi_W - \pi_M)}{(1-\beta\pi_M)}F(\Delta u^*))\Delta^2 u = -\frac{\beta(1-\beta)(\pi_W - \pi_M)}{(1-\beta\pi_M)}\bar{V} - \frac{\beta(1-\pi_M)}{1-\beta\pi_M}\kappa.$$

We now can assess how the moratorium affects the difference in utilities, i.e. determine the sign of $\Delta^2 u$. The right-hand side is negative under discrimination of any type including a mix of taste-based and statistical discrimination. The multiplier of $\Delta^2 u$ is positive because $F(\Delta u^*) < 1$ and $\beta\pi_W - \beta\pi_M < 1 - \beta\pi_M$.

We can further analyze special cases. If we have statistical discrimination $\pi_W > \pi_M$ and $\kappa = 0$, the moratorium increases discrimination:

$$(1 - \frac{\beta(\pi_W - \pi_M)}{(1-\beta\pi_M)}F(\Delta u^*))\Delta^2 u = -\frac{\beta(1-\beta)(\pi_W - \pi_M)}{(1-\beta\pi_M)}\bar{V}.$$

$$\Delta^2 u = -\frac{\beta(1-\beta)(\pi_W - \pi_M)}{(1-\beta\pi_M - \beta(\pi_W - \pi_M)F(\Delta u^*))}\bar{V} < 0.$$

In a special case of taste-based discrimination, we have $\pi_M = \pi_W = \pi$, $\kappa > 0$, and

$$\Delta^2 u = -\frac{\beta(1-\pi)}{1-\beta\pi}\kappa < 0,$$

arriving at the same conclusion.

B Baseline Discrimination Specification

Table A1: Estimates from the Baseline Discrimination Specification on the Full Sample

	(1)	(2)	(3)	(4)	(5)	(6)
Model	Linear	Linear	Linear	Linear	Probit	Logit
African American	-0.057***	-0.057***	-0.057***	-0.057***	-0.145***	-0.233***
	(0.005)	(0.005)	(0.005)	(0.005)	(0.020)	(0.031)
Hispanic	-0.028***	-0.028***	-0.028***	-0.028***	-0.073***	-0.118***
	(0.005)	(0.005)	(0.005)	(0.005)	(0.020)	(0.032)
Male		-0.032***	-0.032***	-0.032***	-0.089***	-0.143***
		(0.005)	(0.005)	(0.005)	(0.016)	(0.026)
Less Than High School			-0.002	-0.002	-0.000	-0.001
			(0.007)	(0.007)	(0.020)	(0.032)
High School Graduate			-0.023***	-0.023***	-0.080***	-0.129***
			(0.007)	(0.007)	(0.020)	(0.031)
Inquiry Order = 2				-0.029***	-0.075***	-0.120***
				(0.005)	(0.020)	(0.032)
Inquiry Order = 3				-0.050***	-0.129***	-0.207***
				(0.005)	(0.020)	(0.031)
Constant	0.605***	0.621***	0.630***	0.656***	0.407***	0.653***
	(0.004)	(0.005)	(0.006)	(0.007)	(0.023)	(0.037)
Observations	25,055	25,055	25,055	25,055	25,055	25,055
R-squared	0.007	0.009	0.010	0.015		

Notes: 1) Table reports coefficients from a within-property linear regression model in columns (1)-(4), probit model in column (5), and logit model in column (6). 2) The outcome variable is an indicator of whether a response was received from the property manager. 3) The mean response to a white identity is 0.605 over the whole sample. 4) Standard errors in parentheses. 5) ***, **, and * denote significance at the 1%, 5%, and 10% levels.

In this section, we document the presence of racial and ethnic discrimination in our sample. It is a necessary condition for such disparities to be exacerbated by eviction moratoria, according to our theoretical analysis in Section 2.

The experimental design described in Section 3 generates a series of binary decisions j for each property i, where the manager chooses whether to respond to an inquiry ($Response_{ij} = 1$) or not ($Response_{ij} = 0$), with $j = 1, 2, 3$. To quantify baseline disparities in manager

responses, we estimate the following model with property fixed effects:

$$Response_{ij} = \delta_i + \beta^{AA} African\ American_j + \beta^{H} Hispanic_j + X_j'\theta + \epsilon_{ij}, \tag{17}$$

where $African\ American_j$ and $Hispanic_j$ are indicator variables for the race or ethnicity of applicant j. The vector X_j includes identity-specific characteristics: gender, maternal education, and the order in which the inquiry was sent. Property fixed effects δ_i absorb all time-invariant listing-specific variation, ensuring identification from within-property differences in responses across applicants. Because fictitious names were randomly assigned and balanced across covariates, estimates of β^{AA} and β^{H} are robust to the inclusion or exclusion of X_j.

We estimate equation (17) using the full sample of inquiries across all weeks and states. The number of observations is slightly smaller than the original 25,428 due to the exclusion of some inquiries, e.g., prior to the start of a state's moratorium, as detailed in Section 4.1. Table A1, Columns (1)–(4), report results from the linear probability model. Columns (5) and (6) present estimates from corresponding Probit and Logit specifications.

Across all models, we find statistically significant evidence of lower response rates to African American and Hispanic applicants relative to white applicants, consistent with the findings in Christensen, Sarmiento-Barbieri, and Timmins (2022). These results confirm the presence of discriminatory behavior in our sample.

C Difference-in-Differences Specification

Table A2: Impact of the End of a Moratorium on Likelihood of Receiving a Response

	Dependent Variable: Response						
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
Treatment	-0.059***	-0.073***	-0.008	-0.007	-0.007	-0.009	0.034
	(0.018)	(0.019)	(0.060)	(0.060)	(0.048)	(0.048)	(0.052)
African American	-0.056***	-0.058***	-0.055***	-0.055***	-0.055***	-0.056***	-0.058***
	(0.007)	(0.008)	(0.007)	(0.007)	(0.012)	(0.011)	(0.013)
African American x Treatment	0.016	0.035*	0.015	0.015	0.015	0.017	0.033
	(0.017)	(0.019)	(0.017)	(0.017)	(0.022)	(0.022)	(0.024)
Hispanic	-0.033***	-0.032***	-0.033***	-0.033***	-0.033***	-0.033***	-0.035***
	(0.007)	(0.008)	(0.007)	(0.007)	(0.008)	(0.007)	(0.009)
Hispanic x Treatment	0.022	0.038**	0.023	0.023	0.023	0.024	0.038*
	(0.017)	(0.019)	(0.017)	(0.017)	(0.021)	(0.021)	(0.022)
#Evictions in 2018, thousands		-0.001***					
		(0.000)					
Stringency Index	0.002***	0.002***					
	(0.001)	(0.001)					
COVID Cases per 1k				0.004	0.004	-0.002	0.004
				(0.024)	(0.026)	(0.026)	(0.027)
Constant	0.544***	0.555***	0.643***	0.622***	0.637***	0.676***	0.628***
	(0.040)	(0.043)	(0.018)	(0.126)	(0.138)	(0.129)	(0.152)
Observations	17,763	15,573	17,912	17,912	17,912	18,713	15,620
R-squared	.	.	0.025	0.025	0.025	0.025	0.024
Number of address_id	5,988	5,252	6,036	6,036	6,036	6,303	5,261
Gender	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Educational Level	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Inquiry Order	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Weekly FEs	No	No	Yes	Yes	Yes	Yes	Yes
Property FEs	No	No	Yes	Yes	Yes	Yes	Yes
Clustered at State-level	No	No	No	No	Yes	Yes	Yes
Control Group	Not yet	Not yet	Not yet	Not yet	Not yet	Not treated	Not yet
Moratoria Data							Benfer et al.

Notes: 1) The outcome variable is an indicator of whether a response was received from the property manager. 2) COVID Cases per 1k is the number of COVID-19 infections in a state on a day per 1,000 people. 3) Row "Control Group" specifies whether the control group consists of not-yet-treated states ("Not yet") or not-yet-treated and never-treated states ("Not treated"). 4) The last column shows estimates using the eviction-moratoria data from Benfer et al. (2023). 5) Standard errors are in parentheses. 6) ***, **, and * denote significance at the 1%, 5%, and 10% levels.

We estimate the following DiD specification:

$$Response_{ijt} = \delta_i + \beta^{AA} African\ American_j + \beta^{H} Hispanic_j + \beta^{T} Treatment_{jt} \tag{18}$$
$$+ \beta^{AAT} Treatment_{jt} \times African\ American_j + \beta^{HT} Treatment_{jt} \times Hispanic_j + X_j'\theta + \epsilon_{ijt},$$

where i indexes rental properties, j indexes applicant identities, and t indexes time (in days). The dependent variable $Response_{ijt}$ equals one if identity j received a response from property i on day t, and zero otherwise. The treatment variable $Treatment_{jt}$ equals one if the inquiry occurs after the expiration of the moratorium in the relevant state.

The variables $African\ American_j$ and $Hispanic_j$ are indicators for the racial or ethnic category of the inquiring identity. The vector X_j includes controls for gender, maternal education, and the order in which the inquiry was sent. Rental property fixed effects δ_i absorb all time-invariant listing-specific variation.

Table A2 presents the results. Column (1) includes controls for the local eviction policy stringency index (Hale et al. (2021)), Column (2) adds historical eviction rates from 2018, but omit week and property fixed effects.[14] Column (3) introduces week and property fixed effects. Column (4) adds daily COVID-19 cases per 1,000 residents at the state level. Column (5) clusters standard errors by state – the level at which eviction moratoria are enacted and repealed – to account for within-state correlation in unobserved shocks. Column (6) retains the full specification and redefines the control group to include never-treated states. Column (7) shows estimates when we use the eviction moratoria data from Benfer et al. (2023) instead of our hand-collected dataset.

The estimated interaction terms—$Treatment_{jt} \times African\ American_j$ and $Treatment_{jt} \times Hispanic_j$—are positive but not statistically significant. The staggered timing of moratorium terminations across states during the spring and summer of 2020 may introduce bias into the canonical DiD estimator. When treatment varies across time, early-treated units serve as controls for later-treated units, potentially contaminating post-treatment comparisons and attenuating estimates toward zero (Goodman-Bacon, 2021). To address this concern, we use estimator that explicitly accounts for variation in treatment timing, as described in the main text.

[14] Historical eviction rates are obtained from Gromis et al. (2022), the most recent available data prior to the pandemic.

D Matching Experimental Addresses to the InfoUSA Database

We standardized all experimental rental addresses to a common format containing the house number, street name (including directional suffixes), and unit number when available. Six addresses were excluded from the procedure because they were too ambiguous to match (e.g., listings containing only a street name or anonymized house number). For the remaining addresses, we extracted the tracts containing experimental listings from the InfoUSA database and implemented a multi-step matching procedure, summarized in Table A9.

In the first step, we identified 6,077 exact matches between experimental addresses and InfoUSA records, requiring complete agreement on the full address, including directional suffixes. The second step relaxed this criterion by ignoring suffixes and matching on house number and street name, while requiring agreement on unit number when present, yielding an additional 1,199 matches. The third step expanded this search to all counties containing experimental tracts, adding 62 further matches.

Subsequent steps employed approximate matching techniques within the experimental tracts. First, we implemented a fuzzy string match between standardized experimental and InfoUSA addresses. Second, we conducted a spatial match using geocoded coordinates of the experimental listings, requiring agreement on the street name. These approaches were particularly useful for listings reported as intersections (e.g., "Main St and University Ave") rather than full addresses.

In the final step, we queried InfoUSA records directly on house number, street name, and zip code for any remaining unmatched addresses, obtaining 30 additional matches.

Altogether, we successfully matched 7,393 experimental listings (87.3%) to InfoUSA addresses.

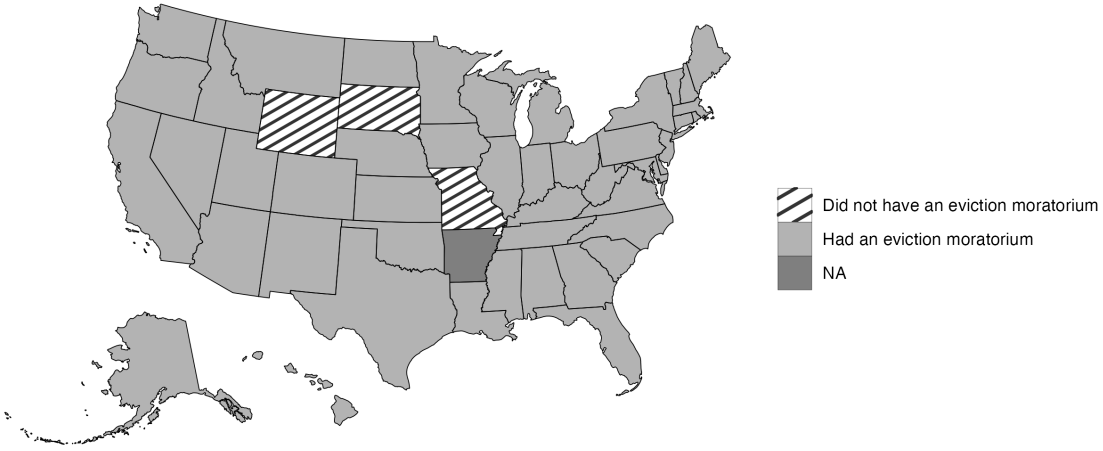
E Tables and Figures

Table A3: Metropolitan Statistical Areas in the Experiment

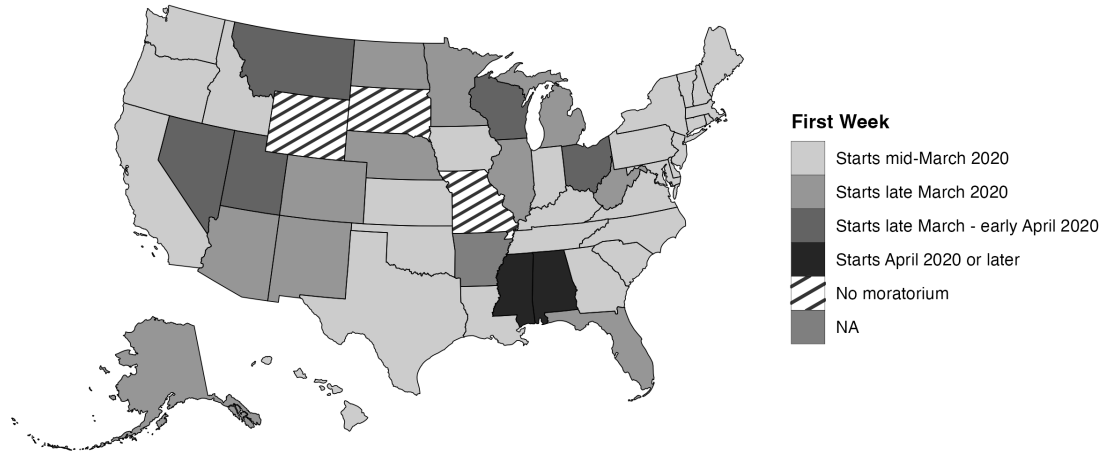
Atlanta-Sandy Springs-Roswell, GA	Minneapolis-St. Paul-Bloomington, MN-WI
Austin-Round Rock, TX	Nashville-Davidson–Murfreesboro–Franklin, TN
Baltimore-Columbia-Towson, MD	New Orleans-Metairie, LA
Birmingham-Hoover, AL	New York-Newark-Jersey City, NY-NJ-PA
Boston-Cambridge-Newton, MA-NH	Oklahoma City, OK
Buffalo-Cheektowaga-Niagara Falls, NY	Orlando-Kissimmee-Sanford, FL
Charlotte-Concord-Gastonia, NC-SC	Philadelphia-Camden-Wilmington, PA-NJ-DE-MD
Chicago-Naperville-Elgin, IL-IN-WI	Phoenix-Mesa-Scottsdale, AZ
Cincinnati, OH-KY-IN	Pittsburgh, PA
Cleveland-Elyria, OH	Portland-Vancouver-Hillsboro, OR-WA
Columbus, OH	Providence-Warwick, RI-MA
Dallas-Fort Worth-Arlington, TX	Raleigh, NC
Denver-Aurora-Lakewood, CO	Richmond, VA
Detroit-Warren-Dearborn, MI	Riverside-San Bernardino-Ontario, CA
Hartford-West Hartford-East Hartford, CT	Sacramento–Roseville–Arden-Arcade, CA
Houston-The Woodlands-Sugar Land, TX	Salt Lake City, UT
Indianapolis-Carmel-Anderson, IN	San Antonio-New Braunfels, TX
Jacksonville, FL	San Diego-Carlsbad, CA
Kansas City, MO-KS	San Francisco-Oakland-Hayward, CA
Las Vegas-Henderson-Paradise, NV	San Jose-Sunnyvale-Santa Clara, CA
Los Angeles-Long Beach-Anaheim, CA	Seattle-Tacoma-Bellevue, WA
Louisville-Jefferson County, KY-IN	St. Louis, MO-IL
Memphis, TN-MS-AR	Tampa-St. Petersburg-Clearwater, FL
Miami-Fort Lauderdale-West Palm Beach, FL	Virginia Beach-Norfolk-Newport News, VA-NC
Milwaukee-Waukesha-West Allis, WI	Washington-Arlington-Alexandria, DC-VA-MD-WV

Table A4: States in the Experiment

	Kentucky	Ohio
Alabama	Louisiana	Oklahoma
Alaska	Maine	Oregon
American Samoa	Maryland	Pennsylvania
Arizona	Massachusetts	Puerto Rico
Arkansas	Michigan	Rhode Island
California	Minnesota	South Carolina
Colorado	Mississippi	South Dakota
Connecticut	Missouri	Tennessee
Delaware	Montana	Texas
Florida	Nebraska	Utah
Georgia	Nevada	Vermont
Guam	New Hampshire	Virgin Islands
Hawaii	New Jersey	Virginia
Idaho	New Mexico	Washington
Illinois	New York	West Virginia
Indiana	North Carolina	Wisconsin
Iowa	North Dakota	Wyoming
Kansas	Northern Mariana Islands	

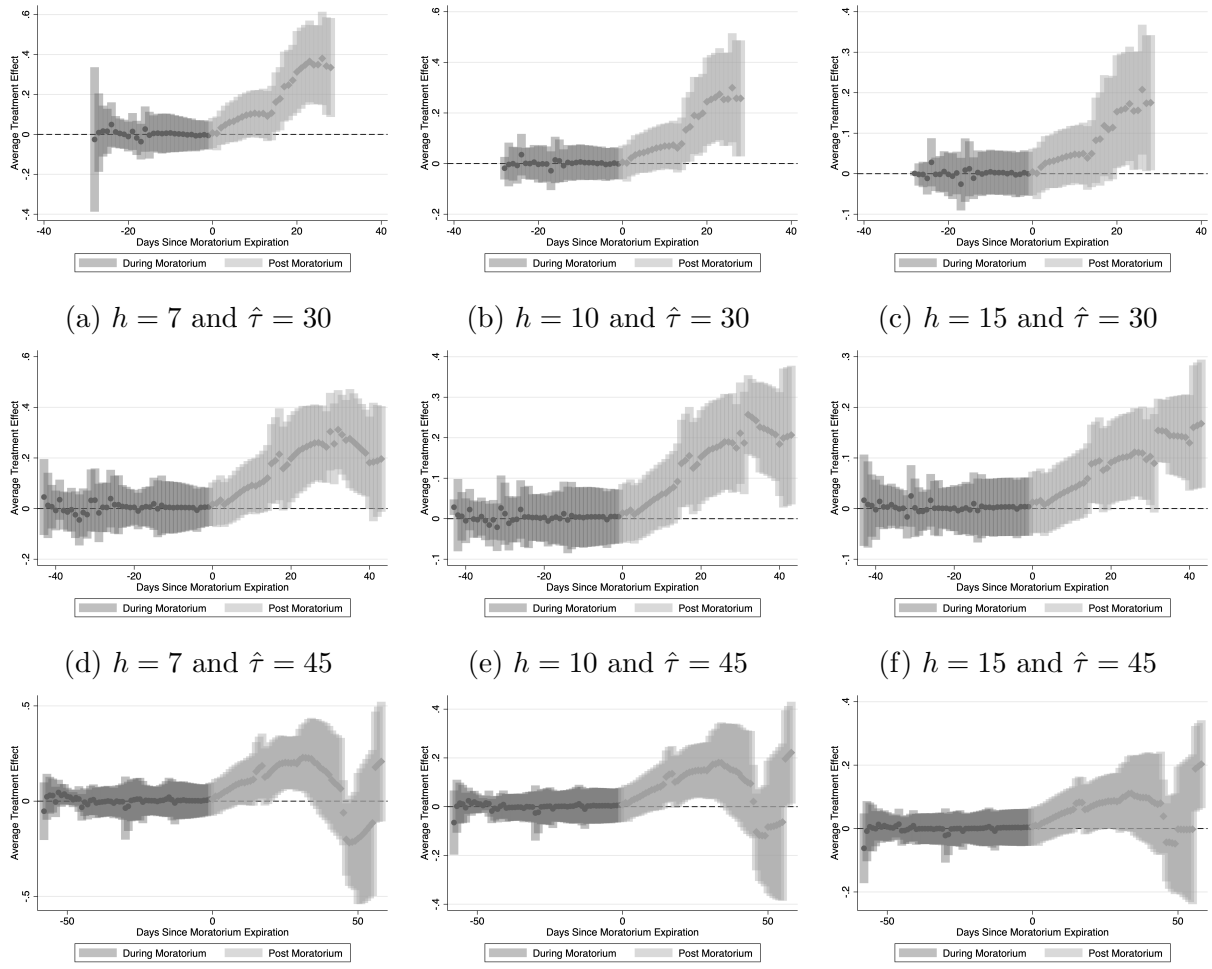


(a) States that Enacted Moratoria



(b) The First Week of the Eviction Moratorium across U.S.

Figure A1: Eviction Moratoria across the U.S.



(a) $h = 7$ and $\hat{\tau} = 30$ (b) $h = 10$ and $\hat{\tau} = 30$ (c) $h = 15$ and $\hat{\tau} = 30$

(d) $h = 7$ and $\hat{\tau} = 45$ (e) $h = 10$ and $\hat{\tau} = 45$ (f) $h = 15$ and $\hat{\tau} = 45$

(g) $h = 7$ and $\hat{\tau} = 60$ (h) $h = 10$ and $\hat{\tau} = 60$ (i) $h = 15$ and $\hat{\tau} = 60$

Figure A2: The Event Study Estimates of the Effect of the End of a Moratorium on the Natural Logarithm of the Relative Response Ratios for an African American Identity Relative to a White Identity from the Staggered DiD for Different Smoothing Parameters h and $\hat{\tau}$ Days Around Treatment

Table A5: Staggered DiD Estimates of the Effect of the End of Moratorium on the Relative Response Ratio Relative to White Applicants

	African Americans			Hispanics		
	$h = 7$	$h = 10$	$h = 15$	$h = 7$	$h = 10$	$h = 15$
Panel A: 30 days around treatment						
ATT	0.101***	0.069***	0.043**	0.028	0.015	0.005
	(0.050, 0.151)	(0.025, 0.114)	(0.003, 0.083)	(-0.021, 0.077)	(-0.027, 0.057)	(-0.029, 0.040)
Number of Observations	724	724	724	724	724	724
Panel B: 45 days around treatment						
ATT	0.139***	0.104***	0.066***	0.053**	0.042*	0.027
	(0.089, 0.188)	(0.061, 0.148)	(0.028, 0.103)	(0.002, 0.105)	(-0.001, 0.084)	(-0.007, 0.061)
Number of Observations	1139	1139	1139	1139	1139	1139
Panel C: 60 days around treatment						
ATT	0.124***	0.094***	0.059***	0.130***	0.106***	0.078***
	(0.072, 0.177)	(0.050, 0.138)	(0.022, 0.096)	(0.069, 0.192)	(0.056, 0.156)	(0.038, 0.118)
Number of Observations	1561	1561	1561	1561	1561	1561

Notes: 1) ATT stands for the Average Treatment Effect on the Treated. 2) h is the smoothing parameter of the weighted logit, see the text. 3) 95% confidence intervals in parentheses. 4) ***, **, and * denote significance at the 1%, 5%, and 10% levels.

Table A6: CoreLogic MLS Rental Listings: Sample Construction

	Sample	# Observations	# Dropped
1	All listings from January 2018 to August 2020	22,096,504	—
2	Rental listings	3,223,536	18,872,968
3	After removing real estate owned (REO)	3,217,397	6,139
4	After removing short sales	3,216,669	728
5	After removing foreclosures	3,212,948	3,721
6	After removing non-residential real estate properties	3,068,476	144,472
7	After removing outliers with extreme characteristics	2,998,364	70,112
8	After removing listings without unique identification #	2,415,059	583,305
9	After removing listings with days on market ≤ 1 or > 4 years	2,367,397	47,662
10	After dropping duplicated listings with the same lease date/price	2,090,292	277,105
11	After merging overlapping listings of the same rental property	1,937,337	152,955
12	After dropping missing or invalid dates	1,069,639	867,698
13	After removing duplicated listings	1,064,237	5,402

Table A7: Staggered DiD Estimates of the Effect of the End of Moratorium on the Natural Logarithm of the Relative Response Ratio for African Americans Relative to White Applicants by Racial Animus Index

	Racial Animus Index Above Median			Racial Animus Index Below Median		
	$h = 7$	$h = 10$	$h = 15$	$h = 7$	$h = 10$	$h = 15$
Panel A: 30 days around treatment						
ATT	0.048*	0.010	-0.001	0.174***	0.160***	0.106**
	(-0.006, 0.102)	(-0.031, 0.051)	(-0.036, 0.033)	(0.060, 0.288)	(0.056, 0.264)	(0.019, 0.193)
Number of Observations	443	443	443	254	254	254
Panel B: 45 days around treatment						
ATT	0.180***	0.129***	0.086***	0.174***	0.160***	0.106**
	(0.103, 0.256)	(0.074, 0.184)	(0.043, 0.129)	(0.060, 0.288)	(0.056, 0.264)	(0.019, 0.193)
Number of Observations	734	734	734	329	329	329
Panel C: 60 days around treatment						
ATT	0.116**	0.090***	0.062**	0.174***	0.160***	0.106**
	(0.026, 0.207)	(0.024, 0.156)	(0.011, 0.113)	(0.060, 0.288)	(0.056, 0.264)	(0.019, 0.193)
Number of Observations	1021	1021	1021	389	389	389

Notes: 1) ATT stands for the Average Treatment Effect on the Treated. 2) h is the smoothing parameter of the weighted logit, see the text. 3) 95% confidence intervals in parentheses. 4) ***, **, and * denote significance at the 1%, 5%, and 10% levels.

Table A8: Staggered DiD Estimates of the Effect of the End of Moratorium on the Natural Logarithm of the Relative Response Ratio for African Americans Relative to White Applicants by Tenant Rights Protection Index

	Tenant Rights Protection Index Above Median			Tenant Rights Protection Index Below Median		
	$h = 7$	$h = 10$	$h = 15$	$h = 7$	$h = 10$	$h = 15$
Panel A: 30 days around treatment						
ATT	0.047***	0.055***	0.040**	0.052*	0.034	0.018
	(0.028, 0.067)	(0.029, 0.081)	(0.002, 0.078)	(-0.010, 0.114)	(-0.024, 0.091)	(-0.031, 0.067)
Number of Observations	92	92	92	481	481	481
Panel B: 45 days around treatment						
ATT	0.192***	0.141***	0.084***	0.087***	0.074***	0.050**
	(0.148, 0.235)	(0.110, 0.172)	(0.057, 0.111)	(0.034, 0.139)	(0.026, 0.122)	(0.010, 0.090)
Number of Observations	269	269	269	783	783	783
Panel C: 60 days around treatment						
ATT	0.289***	0.205***	0.131***	0.147***	0.110***	0.067***
	(0.238, 0.340)	(0.167, 0.243)	(0.097, 0.164)	(0.095, 0.200)	(0.063, 0.158)	(0.028, 0.106)
Number of Observations	403	403	403	1035	1035	1035

Notes: 1) ATT stands for the Average Treatment Effect on the Treated. 2) h is the smoothing parameter of the weighted logit, see the text. 3) 95% confidence intervals in parentheses. 4) ***, **, and * denote significance at the 1%, 5%, and 10% levels.

Table A9: Experiment to InfoUSA Address Match Rates

	Matching Step	New Matches	Total Matches	Share
1	Exact full-string address match in tracts from the experiment	6,077	6,077	71.2%
2	Exact house number + street name within a tract	+1,199	7,276	85.9%
3	Exact house number + street name within a county within tracts from the experiment	+62	7,338	86.6%
4	Fuzzy match on cleaned addresses in tracts from the experiment	+22	7,360	86.9%
5	Spatial match using lat/lon with street-name match	+12	7,372	87.0%
6	House number + street name + zip code	+30	7,393	87.3%
	Total matched		7,393 / 8,470	87.3%

Notes: Six addresses were disqualified from the matching procedure and are excluded from all figures above. The original experimental address count is 8,476; after excluding the 6 disqualified addresses, the matching sample is 8,470.

Table A10: Distribution of Tenant Races

Race	Number of Tenants	Percent
African American	400	12.03
Hispanic	405	12.18
Other	219	6.58
White	1,835	55.17
Multiple	467	14.04
	3,326	100.00

Table A11: Staggered DiD Estimates of the Effect of the End of Moratorium on the Risk Difference: the Predicted Share of Same Race after No Response Minus the Predicted Share of Same Race after a Response

	$h = 7$	$h = 10$	$h = 15$
	Panel A: 30 days around treatment		
ATT	0.071***	0.058**	0.052***
	(0.030, 0.111)	(0.013, 0.103)	(0.013, 0.092)
Number of Observations	709	719	737
	Panel B: 45 days around treatment		
ATT	0.045*	0.042*	0.043**
	(-0.005, 0.095)	(-0.002, 0.086)	(0.008, 0.078)
Number of Observations	1096	1137	1166
	Panel C: 60 days around treatment		
ATT	0.083***	0.080***	0.066***
	(0.030, 0.137)	(0.038, 0.121)	(0.037, 0.095)
Number of Observations	1485	1539	1598

Notes: 1) ATT stands for the Average Treatment Effect on the Treated. 2) h is the smoothing parameter of the weighted logit, see the text. 3) 95% confidence intervals in parentheses. 4) ***, **, and * denote significance at the 1%, 5%, and 10% levels.

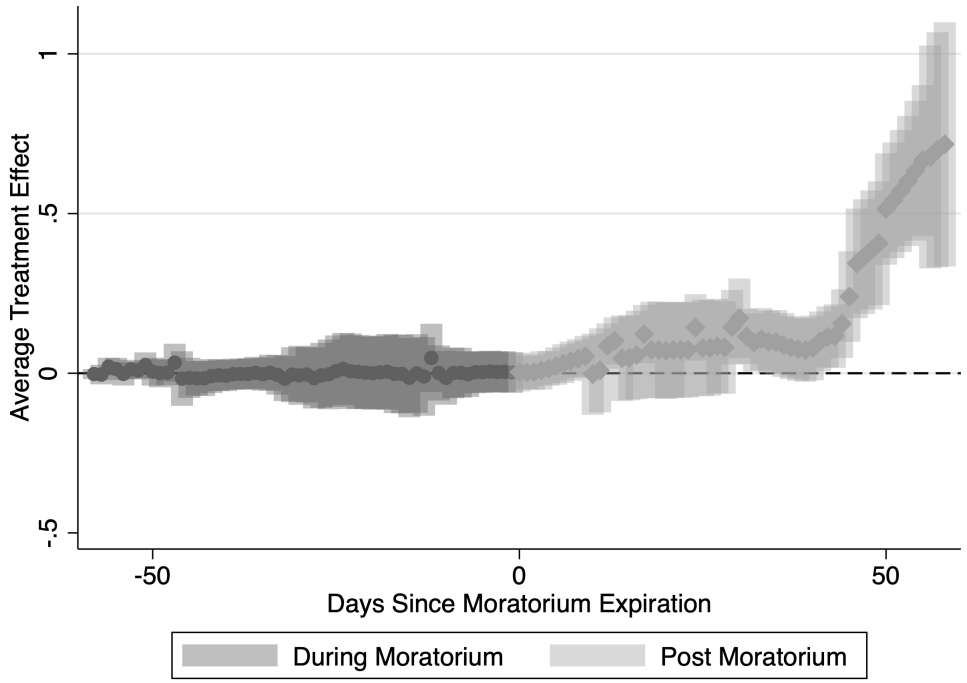


Figure A3: Event Study Coefficients from Callaway and Sant'Anna (2021)'s Estimator for the Risk Difference – the Predicted Share of Same Race after No Response Minus the Predicted Share of Same Race after a Response – with the Smoothing Parameter $h = 10$ and $\hat{\tau} = 60$ Days around Treatment