The Limits of Merit:

Career Pipelines and Representation in US Law Enforcement

Elisa Maria Wirsching*

May 7, 2026

Abstract

US police departments are systematically Whiter and more Republican than the populations they serve, even within formally meritocratic civil service systems. What drives political sorting in these systems? Existing scholarship and decades of diversification reform have focused on the written entry exam, yet bureaucratic careers extend well beyond entry. Drawing on novel data linking civil service exam records, payroll, promotion, attrition, and voter registrations for over 99,000 entry-level exam takers and 58,000 sworn officers at the New York City Police Department, this study traces how selection processes create and reinforce workplace stratification across bureaucratic careers. While the entry exams produce some sorting by partisanship, they are insufficient to explain the representational gaps that emerge. Even among applicants with comparable exam scores, the post-exam hiring stage favors Republican and White candidates, who are subsequently more likely to be promoted, receive awards, and remain with the agency over time. Equalizing rates at the discretionary hiring stage alone would increase the appointment of Democratic and Black candidates by 78% and 27%, respectively. These findings redirect attention from the entry exam to the bureaucratic pipeline, and reveal how bureaucratic institutions can sustain stratification despite formal meritocratic rules.

*Assistant Professor, London School of Economics and Political Science.
Correspondence: e.m.wirsching@lse.ac.uk
I thank Monika Bauhr, Maria Carreri, Sara Constantino, Sanford Gordon, Kun Heo, Victor Lapuente, Julia Payson, Tara Slough, Arthur Spirling, Carolina Torreblanca, Hye Young You, and participants at APSA, MPSA, the 2024 Gothenburg CoG conference, MZES, the NYU American Politics workshop, the Princeton CSDP conference on Local Bureaucracy, and the Yale 2025 CSAP conference for helpful comments and suggestions. Niklas Hähn, Maxine Lee and Paige Reddick provided excellent research assistance.

1 Introduction

A growing body of evidence shows that many US state and local bureaucracies are unrepresentative of the jurisdictions they serve. Across 50 states and the District of Columbia, the workforce of state and local agencies is systematically Whiter than the resident population, with the gap most pronounced for officials and administrators (a median 14 percentage points), technicians (9 points), and protective service occupations such as policing, firefighting, and corrections (7 points; see Figure 1).[1] Policing has long been singled out for its skew. Across a sample of 98 major US law enforcement agencies, 56% of officers are White compared with only 38% of the population in the relevant jurisdictions (Ba et al., 2025). As policing issues have become increasingly polarized along partisan lines (Grosjean et al., 2023; Gaudette, 2025), scholars and commentators have also paid increasing attention to the fact that law enforcement officers are disproportionately drawn from Republican and conservative bases (Ba et al., 2025; Reny et al., 2025). Across the same sample, 32% of officers identify with the Republican Party compared with 14% of voting-age citizens, while only 31% identify with the Democratic Party compared with 43% of civilians (Ba et al., 2025).

These representational gaps matter for several reasons. First, public employment is a key channel of distributive spending and economic mobility, particularly for historically marginalized groups. In 2023, local governments across the US employed 14.1 million individuals and spent more than $855 billion or 34% of their annual expenses on wages and salaries for employees (US Census Bureau, 2023, 2025).[2] Second, who staffs local government shapes both performance and legitimacy. Bureaucrats from underrepresented groups often deliver more equitable services to those groups, partisan identification can predict how bureaucrats exercise discretion, and the mere presence of demographic representation fosters citizens' trust

[1] By aggregating state and local government employment to the state level, this comparison masks how the gap is distributed across local jurisdictions within a state. Appendix Figure A6 affirms that the same pattern holds at the level of Public Use Microdata Areas (PUMAs, approximately 100,000 residents). Across 2,462 PUMAs in the 2023 American Community Survey, state and local government workers living in a given PUMA are 3.9 percentage points Whiter at the median than that PUMA's residents.

[2] This excludes expenses on retirement funds which also constitute a substantial source of expenditures.

Figure 1: Representativeness in US State and Local Government Employees

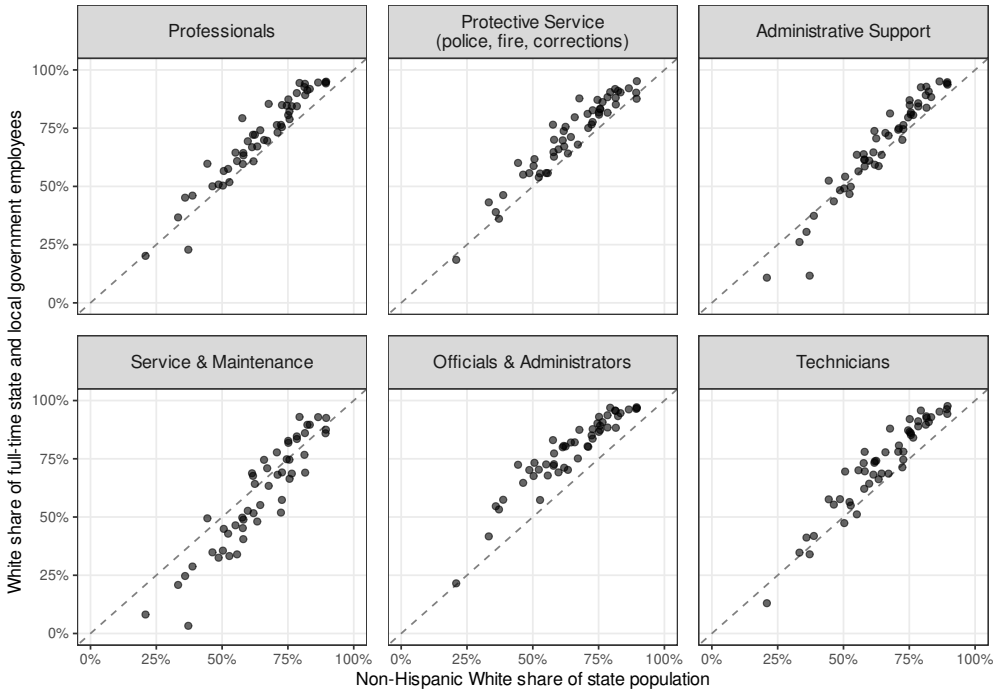

Note: Each point is one of the 50 states or the District of Columbia. The vertical axis is the White share of full-time state and local government employees in the indicated job category; the horizontal axis is the non-Hispanic White share of the state population. Points above the 45-degree line indicate workforces that are Whiter than the state population. Source: 2023 EEO-4 Public Use File (EEOC) and 2023 American Community Survey 1-year estimates. EEO-4 covers state and local government employers with 100 or more employees. Public elementary and secondary school workforces are reported separately under EEO-5 and are not shown. State-level aggregation may mask within-state variation across local jurisdictions.

and cooperation with the state (Bradbury and Kellough, 2011; Nicholson-Crotty et al., 2016; Ba et al., 2021; Harvey and Mattia, 2022; Harris, 2023; Donahue, 2023; Lerman and Page, 2015; Riccucci et al., 2018; Miller and Segal, 2018). The dominance of certain racial and partisan groups is also striking against the institutional design of these bureaucracies. Most US local agencies operate under formalized civil service systems with merit-based examinations, standardized promotions, and tenure protections, which were established precisely to insulate bureaucratic composition from partisan influence and should, at least in theory, enable members from all groups and parties to win employment in the public sector (Skowronek, 1982; Kuipers and Sahn, 2023).

Existing explanations for representational gaps in formally meritocratic systems have

largely treated civil service as a binary system, asking whether merit-based hiring affects selection rather than how meritocratic rules play out within it. Most empirical work in this tradition examines the written entry exam, the most visible gatekeeper of formalized civil service systems. A hidden curriculum literature documents how ostensibly neutral entry exams reproduce existing inequalities through unequal access to education, financial barriers to test preparation, and biased content (Portillo et al., 2020; Moreira and Pérez, 2025; Riccucci and Riccardelli, 2015). A related self-selection literature shows that those drawn to take these exams hold systematically different attitudes before they ever apply (Sidanius et al., 1994; Reny et al., 2025). Yet bureaucratic careers extend well beyond the entry stage. Who is appointed from the pool of exam passers, who advances through the ranks, and who stays long enough to shape the agency also determine its composition. The within-system variation in how meritocratic rules play out across the full career pipeline has received far less attention.

This paper provides a first systematic decomposition of representational disparities across the bureaucratic pipeline of the New York City Police Department, one of the largest local bureaucracies and the largest local police force in the United States. Such an analysis has been hampered by data limitations, since it requires linking individual-level records across hiring, promotion, and attrition over many years. I overcome these constraints by combining civil service exam records, payroll data, promotions, awards and attrition records, and individual voter registrations into a linked dataset that traces the partisan, racial, and demographic composition of more than 99,000 entry-level exam takers and 58,000 sworn officers on the NYC payroll between 2014 and 2021.

The findings cut against the prevailing focus on the entry exam. Although the application and written test stages do produce some sorting by partisanship, they are insufficient to explain the representational gaps that ultimately emerge at the workforce level. Substantial additional gaps emerge and compound at the discretionary stages that follow, including post-exam character investigations and psychological evaluations, discretionary promotions and

assignments, and involuntary terminations. Among applicants with comparable exam scores, Republican and White candidates are about 8 and 3 percentage points more likely to be appointed than Democratic and Black candidates. Additionally, Republican and White officers are more likely to be promoted to prestigious assignments, receive departmental awards, and remain with the agency longer and are more likely to leave voluntarily. Simple counterfactual exercises show that equalizing rates at the discretionary hiring stage alone could increase the appointment of Democratic and Black candidates by 78% and 27%, respectively.

This project is primarily descriptive, following calls for larger-scale quantitative description (Gerring, 2012). My approach documents the cumulative effect of sorting across the bureaucracy, but cannot identify the causal mechanisms producing these disparities. The observed gaps between partisan and racial groups could reflect demand-side explanations (e.g., discrimination in organizational decision-making), supply-side forces (e.g., differential choices by applicants and officers themselves), or factors that correlate with both identity markers and selection outcomes. A key concern is merit: observed disparities could be driven by differences in qualifications or job performance. I address this concern where possible by showing that gaps persist when controlling for exam scores at hiring and measures of on-the-job performance such as disciplinary complaints.

Nonetheless, a richer descriptive understanding of bureaucratic selection processes offers several contributions. First, by documenting the relative magnitude and location of disparities across career stages, it provides essential baseline estimates to guide causal research and can inform reform-minded policymakers on which selection processes to address. Diversification of police recruitment is one of the oldest reform efforts in American policing, and decades of litigation and consent decrees on the written exams have made the entry point the central target of reform (McCrary, 2007; Linos, 2017), thus largely neglecting later and more discretionary stages in the career pipeline. Second, it extends the literature on bureaucratic selection by showing how sorting arises within the rank-and-file of a meritocratic system, absent top-down political appointments (Lewis, 2008) and beyond entry exams (Moreira and

Pérez, 2025). Finally, by applying the pipeline logic of the political selection literature (Fox and Lawless, 2011; Rolke and Rickne, 2025) to bureaucratic careers and extending emerging evidence on partisan sorting in private-sector workplaces (Bonica et al., 2020; McConnell et al., 2018; Chinoy and Koenen, 2024; Colonnelli et al., 2025) to a public-sector setting designed to prevent it, the paper bridges work on political pipelines and the politics of the workplace.

2 Civil Service, Discretion, and Political Sorting in Bureaucratic Careers

Civil service systems — now standard in most modern bureaucracies — rest on two key premises: political neutrality and meritocratic competence. In the US, the federal Pendleton Act of 1883 followed by numerous state and local jurisdictions throughout the nineteenth and early twentieth centuries established competitive exams and civil service commissions so that government jobs were awarded on the basis of merit rather than political connections (Kuipers and Sahn, 2023; Anzia and Trounstine, 2025). These reforms responded to a system widely seen as corrupted by partisan favoritism and machine politics, where political elites distributed coveted government jobs to political loyalists through patronage and direct appointments (Skowronek, 1982; Reid and Kurth, 1992; Ruhil, 2003). By removing staffing decisions from partisan control, civil service systems promise a bureaucracy insulated from the electoral interests of those in power. And by replacing discretionary appointments with standardized examinations, Weberian merit-based systems require that government employees are selected based on their ability and competence rather than other personal characteristics, such as race, gender or class (Weber, 1946).

Together, these premises, at least in theory, imply equal access to government jobs and upward social mobility. Weber himself recognized this function of competitive examinations, noting that they "mean or appear to mean a 'selection' of those who qualify from all social

strata rather than a rule by notables" (Weber, 1946, p. 240). If government employment is determined by competence rather than connections, members of all groups, parties, and classes should be able to compete for and advance through public service on equal terms (Kuipers and Sahn, 2023; Peters, 2015). An early generation of scholarship on US civil service reform largely confirmed that it delivered on this promise, eliminating de jure discrimination on grounds of partisan politics, race, and national origin (Van Riper et al., 1963; Hoogenboom, 1968). Similarly, recent work finds that reform did open avenues to representation for constituencies previously locked out of government jobs, especially in smaller municipal bureaucracies (Kuipers and Sahn, 2023). Nevertheless, in practice, local bureaucracies across US jurisdictions and agencies today remain dominated by certain partisan and racial groups. What explains this persistence of partisan and racial stratification in professionalized bureaucracies governed by civil service rules and political insulation?

Existing scholarship offers two main accounts. The first emphasizes self-selection in recruitment. Individuals whose values align with an organization's mission and culture are more likely to seek employment there (Schneider, 1987; Besley and Ghatak, 2005). In policing, for example, a growing body of work documents that those who aspire to law enforcement careers hold more conservative and authoritarian attitudes before they apply (Sidanius et al., 1994; Reny et al., 2025), and that recruitment messaging can shift who is drawn to apply (Linos, 2017). While this work highlights important supply-side dynamics, it focuses almost exclusively on who seeks to enter a profession, not on how institutional processes shape who is ultimately hired, promoted, or retained.

A second account focuses on the "hidden curriculum" of examination-based systems, showing how ostensibly neutral merit exams can reproduce existing societal inequalities. As Weber (1946, p. 240) put it, "democracy fears that a merit system and educational certificates will result in a privileged 'caste."' Written exams can disadvantage underrepresented groups due to unequal access to education, financial barriers to test preparation, or exam questions that reward general aptitude over job-relevant skills (Riccucci and Ric-

cardelli, 2015; Portillo et al., 2020; Kuipers and Sahn, 2023; Moreira and Pérez, 2025). In fact, several scholars have argued that municipal civil service reforms in the US aimed to advantage upper- and middle-class, native-born whites in the competition for government employment (Hofstadter, 1955; Shefter, 1993). Studying the implementation of civil service exams through the Pendleton Act, Moreira and Pérez (2025) find that the reform significantly reduced the representation of lower-status and immigrant groups while raising the share of employees with higher-status parental occupations in the US federal bureaucracy. These concerns are especially pronounced in uniformed services, where written exams remain a dominant selection tool and have been subject to repeated litigation and consent decrees since the 1960s. As Riccucci and Riccardelli (2015, p. 352) conclude, "it is the use of written exams that has perhaps been the biggest obstacle for police and fire departments to achieve social diversity based on race and ethnicity."

2.1 Decomposing Civil Service Across the Bureaucratic Pipeline

While these accounts offer valuable critiques of exam-based selection, both treat the composition of a bureaucracy as primarily determined at the entry point. Yet, bureaucratic careers extend well beyond the first selection stage. Who is appointed from the pool of exam passers, who advances through the ranks, and who stays long enough to shape the organization also determines the composition of the bureaucracy. The existing literature on the effects of civil service systems has largely treated meritocracy as a binary system-level variable: a jurisdiction either has competitive civil service exams or it does not (Rauch and Evans, 2000; Folke et al., 2011; Portillo et al., 2020; Kuipers and Sahn, 2023; Moreira and Pérez, 2025).[3] In contrast, I argue that civil service is better understood as a pipeline of career stages with varying degrees of formalization. Decomposing this pipeline — rather than focusing on the mere existence of merit exams — is necessary to understand where political sorting emerges and how much each stage contributes.

[3] An important exception is growing work on the interests of bureaucratic management in hiring and their implications for bureaucratic representativeness (Hassan et al., 2024; Brierley et al., 2026).

The key logic of civil service systems is that formalized rules tie career outcomes to codified, measurable criteria. Yet, within most professionalized civil service systems, career stages vary considerably in the degree to which outcomes are determined by such criteria.[4] Some stages are highly formalized. Written entrance exams, for instance, typically define clear cutoff scores necessary to be placed on official civil service lists, and agencies then hire among the top-scoring candidates — a provision commonly known as the "rule of three." Similarly, competitive promotions require candidates to pass exams and be placed on rank-ordered lists. Across these formalized stages, criteria are specified in advance, outcomes are determined by measurable performance, and the process is often administered by an independent body (such as a civil service commission) rather than the hiring agency itself.

Other stages in the bureaucratic pipeline are far less constrained. Performance criteria in civil service systems tend to grow more subjective at higher ranks, as measurement becomes more difficult and more dependent on cultural and social values (Peters, 2015). Assignments to specialized units, nominations for awards, and promotions between grades that do not require competitive exams lack standardized benchmarks. Decisions about discipline and involuntary separation are similarly ambiguous. And even at the entry stage, formalized and discretionary elements often coexist: a standardized written exam may be followed by a character investigation or psychological evaluation with considerable subjective judgment.[5] To be sure, as the hidden curriculum literature has demonstrated, formalized stages are not immune to producing disparate outcomes. But the question this paper addresses is whether the less formalized stages, which have received less scholarly and reform attention, also contribute to sorting and how much. If discretionary stages increase the number of factors that shape career outcomes and these disadvantage underrepresented groups, we

[4] A parallel literature on street-level bureaucracy highlights the importance of discretion in how public servants deliver services to citizens (Lipsky, 1980). This paper examines discretion in personnel processes that determine who those street-level bureaucrats are in the first place.

[5] A large literature on private-sector organizations documents that discretionary stages (e.g., subjective evaluations, unstructured interviews, informal networks) widen the set of factors that shape who is selected and promoted, while formalized processes narrow outcomes to codified criteria (Reskin, 2003; Castilla and Benard, 2010; Rivera, 2012). These dynamics have received little attention in public bureaucracies.

should expect political sorting to concentrate at the less formalized stages of the bureaucratic pipeline.

This expectation is consistent with both demand-side and supply-side mechanisms. On the demand side, discretionary stages give institutional gatekeepers greater freedom to exercise preferences (whether conscious or not) in ways that formalized criteria would constrain. On the supply side, less formalized stages tend to be longer, less transparent, and harder for candidates to navigate, creating more opportunities for differential self-selection. Candidates with weaker networks, fewer outside resources, or less familiarity with organizational norms may be more likely to withdraw or underperform in unstructured settings. While it is hard to empirically adjudicate between these mechanisms, decomposing the civil service pipeline by stages and degree of formalization shows where the system's neutrality premise bites and where it does not, which matters for institutional design regardless of the underlying mechanism.

2.2 Formalization and Discretion in Policing Careers

US law enforcement is a good case to test this argument for several reasons. First, in most local jurisdictions, policing is governed by strong civil service systems (Riccucci and Riccardelli, 2015; Ornaghi, 2019; Anzia and Trounstine, 2025). Most US states decentralize civil service procedures to the local level, and nearly all major cities administer written exams for police hiring and promotion, even where statutes do not formally require them and even in states that have weakened or abolished state-level civil service (e.g., Florida, Georgia, Texas) (Riccucci and Riccardelli, 2015).

Second, US law enforcement remains systematically unrepresentative of the jurisdictions they serve, and a robust empirical literature shows that who officers are shapes how they do their jobs. Ba et al. (2021) find that Black, Hispanic, and female officers in Chicago use force less often than White and male officers on the same shift and patrol assignment,

especially against Black civilians and in majority-Black areas.[6] There is also some evidence that Democratic officers make fewer traffic stops and arrests, use force less often, and exhibit smaller racial disparities than Republican officers (Donahue, 2023).[7]

Third, policing careers run through a structured pipeline with substantial variation in the degree of formalization. For example, standardized written entry exams with pre-specified cutoffs are paired with more discretionary screens, including character investigations, psychological evaluation, and physical tests. Similarly, while promotions to civil service ranks such as sergeant or lieutenant rely on formalized exams, horizontal moves, such as grade changes within the detective track or assignments to elite units, are far more discretionary. This within-agency variation allows me to compare formalized and discretionary career stages while holding organizational culture and institutional context largely fixed.

3 Case and Data

I study the New York City Police Department (NYPD), the largest municipal police force in the United States. NYC operates one of the most formalized civil service systems in the country, with strict rules governing hiring, promotions, and separations. Critically for this paper, the NYPD career pipeline contains both highly formalized and highly discretionary stages, and I am able to observe each in the administrative data. At entry, I observe both exam scores and ultimate hiring outcomes, allowing me to compare who passes the written exam with who is ultimately appointed after the discretionary screening process. For career progression, I can distinguish formalized promotions (exam-based advancement to sergeant, lieutenant, and captain) from discretionary ones (grade changes within the detective track, elite unit assignments, and departmental awards). And for attrition, I observe the reason

[6] This does not imply that increasing passive representation will mechanically reduce racial disparities in policing outcomes. Work on perceptions of representativeness and the role of local context (Brunson and Gau, 2015; Benton, 2020; Socia et al., 2021) suggests the link between representation and performance is complex, and more work is needed on how macro-level socioeconomic and institutional factors interact with officer identity.

[7] Ba et al. (2025) find little difference between Republican and Democratic officers in Chicago and Houston after multiple-comparison corrections.

for each officer's separation, distinguishing voluntary retirements from involuntary exits such as dismissals and terminations. Finally, NYC's elected officials have been overwhelmingly Democratic for decades. Showing the dominance of Republican and White employees in a largely liberal city government makes the case particularly useful for studying how sorting emerges within the bureaucratic pipeline itself rather than through political turnover and influence.

I start with a roster of roughly 300,000 unique employees across the six major agencies of NYC in terms of their staff size, including the Department of Education, the Police Department, the Fire Department, the Department of Correction, the Department of Social Services, and the Department of Sanitation. This data comes from the NYC annual payroll between 2014 and 2021 and covers employees with appointment dates between 1970 and 2021.[8] To identify the various demographic attributes of these employees, I merge the employment records with the 2021 L2 voter file based on employees' last names, first names, and middle initials. I restrict possible matches to registered voters in the city's five boroughs or one of the neighboring counties of New York State since NYC agencies require their employees to reside within these areas.[9] Following related work (Ba et al., 2025), I employ the probabilistic record linkage algorithm by Enamorado et al. (2019) and retain all matches with a posterior match probability of at least 0.7.

To study the specific drivers of selection among NYPD officers, I then add information on the career trajectories of about 58,000 uniformed police employees, including appointments, promotions, and retirements, from official records published in the City Record newspaper since 2014.[10] For a cross-section of 33,000 active officers (as of October 2021), I am further able to add information on their exact assignment, their arrest history, as well as their

[8] https://data.cityofnewyork.us/City-Government/Citywide-Payroll-Data-Fiscal-Year-/k397-673e; I identify unique employees based on the last name, first name, middle initial, and appointment date.

[9] Technically, NYC also allows employees to reside in Putnam and Orange County. Yet, because these counties are further away from NYC boundaries and relatively small, I exclude these to reduce the risk of false positive matches.

[10] https://www.nyc.gov/site/dcas/about/cityrecord-editions.page

departmental awards from the official NYPD profiles.[11] Finally, I obtain data on roughly 99,000 entry-level and 5,700 promotion exams for the NYPD between 2014 and 2021, which I match to the L2 voter file and the NYPD officers on the payroll. This allows me to assess the attributes of both hired and non-hired NYPD aspirants.[12] I link these administrative data probabilistically based on individuals' full names and other employment details, where possible, in all these merging procedures. Appendix A describes these different matching procedures in more detail. As Figure A1 shows, I can correctly match most employees and records with a very high probability—the median posterior probability of a match is above 0.95 across all matching procedures.

To measure the partisanship and race of individuals in my data, I rely on the information in the L2 voter files. For partisanship, I focus on the three main categories included in L2: Democrat, Republican, and Non-Partisan. Together, these comprise 96% of the 7,940,000 voter registrations in NYC and its surrounding counties. Note that the L2 information on partisanship in New York — unlike for other US states — is based on official registration records and does not require imputations. However, I must rely on L2's proprietary imputation algorithm to measure bureaucrats' and citizens' race. To code racial categories of registered voters, L2 combines the given name, surname, and demographics of a voter's census block for their inference. I discuss the validity of the L2 race categories in Section 6 in detail.

4 Political Sorting in NYC's Bureaucracy

I begin by comparing the demographic composition of front-line employees across NYC's six major agencies with the city's voting population. Figure 2 displays results along partisanship, race, and gender. Most agencies diverge from their jurisdictions, albeit in different directions. The police, fire, and sanitation departments have substantially smaller shares of

[11] https://nypdonline.org/link/2

[12] Note that the exam data does not provide information on *all* applicants, but instead contains individuals who successfully passed the exam (i.e., those with a score of 70/100 or above).

Democratic, non-White, and female members than the NYC electorate. The NYPD underrepresents Democrats by 28 percentage points and non-Whites by 15 percentage points. The Department of Correction, in contrast, closely matches the city's voters in partisanship and even *overrepresents* non-Whites — a striking difference given that both are law enforcement agencies operating under the same civil service system. The Department of Education and the Department of Social Services skew more Democratic, non-White, and female than the general population. Representational gaps in NYC's police force also vary geographically: Democrats are underrepresented and Republicans overrepresented across all NYC boroughs, with racial gaps particularly stark in high-crime and majority-Black communities (see Figures A7 and A8).

Figure 2: Share of Demographics, Agencies vs. Registered Voters in NYC

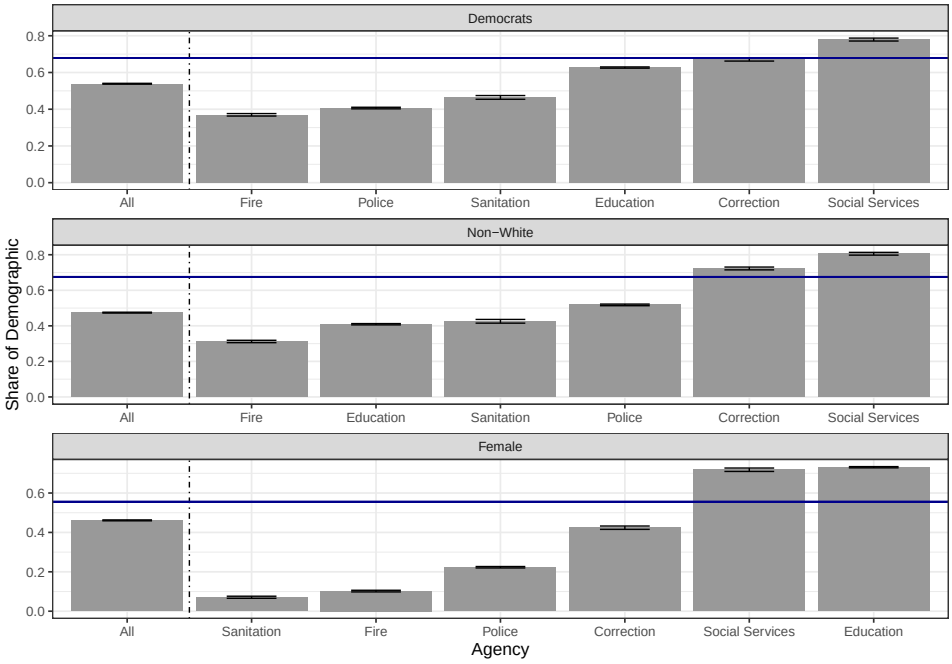

Note: The blue line represents respective share of demographic among NYC voters. Agency employees only include uniformed/public facing employees plus leadership (i.e., I exclude administrative and other agency staff). Agency personnel includes individuals living in NYC's five boroughs and neighboring boroughs, while NYC citizens are restricted to NYC's five boroughs. Agency estimates are weighted by the posterior probability of matches between agency payrolls and the voter file.

5 Tracing Stages in the NYPD's Pipeline

What accounts for the dominance of White and Republican employees at NYC's police department? While I focus on partisan and racial representation gaps across the pipeline, I acknowledge that these are bundled categories. Fixed effects and controls allow me to make more informative comparisons and to account for measurable differences across groups, for instance, by comparing officers within the same entry cohort or with similar exam scores. However, some differences remain unobservable, such as public sector motivation, previous experiences, or biases by selectors. Importantly, these unobserved factors may themselves be part of the sorting process rather than an alternative explanation for it. The goal of the decomposition is not to identify causal mechanisms for representational gaps but to show where gaps concentrate and how much each stage contributes.

5.1 Selection Into the Force

The NYPD hiring process combines both formalized and discretionary stages. To become a police officer at the NYPD, applicants need to be between 21 and 35 years old, have earned 60 college credits with a minimum GPA of 2.0 or 2 years of military service and must live within NYC or one of the neighboring boroughs at the time of appointment. Candidates who satisfy the basic selection criteria first take a written exam offered by the Department of Citywide Administrative Services (DCAS) which covers 9 cognitive abilities. Only candidates with a minimum score of 70% are placed upon an eligible civil service list for appointment, with better performing candidates being placed further up on the list. After an eligible list is certified by the DCAS commissioner, agencies can appoint from the list in ranking order. This regulated process is followed by multiple discretionary hiring stages. Once a list number is reached, candidates advance to a medical exam, written and oral psychological exams, a check for "character and satisfactory background," which screens for arrest records, convictions, and discharge of employment, a physical test, and a drug and

alcohol screening. These additional examinations are all administered by NYPD directly. After successful completion of this process, candidates are then hired as probationary officers and are required to complete the 6-month police academy.

Figure 3: Eligible Citizens, Police Exam Takers, Hired Exam Takers - NYPD

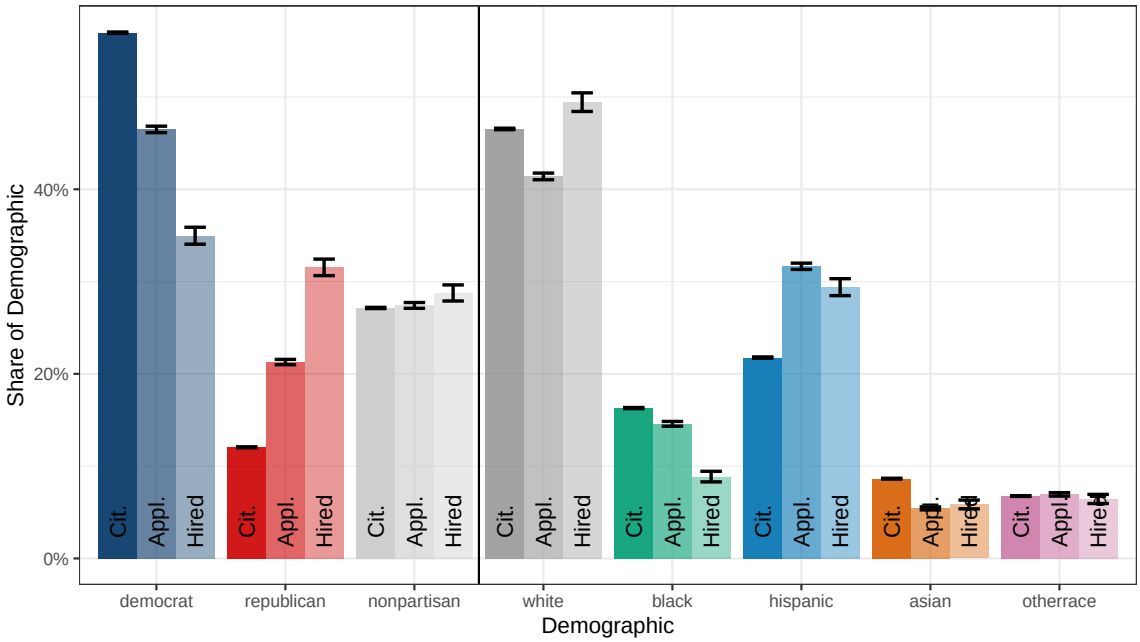

The three bars among each partisan and racial group represent (from left to right) (1) share among NYC voters, (2) share among police exam takers, and (3) share among hired exam takers. Voters and exam takers are matched on age. All estimates are weighted by posterior probability of a match (L2 and exam data).

To follow the process of hiring at the NYPD empirically, Figure 3 shows the share of party and racial categories among three different groups: (1) voters in counties eligible for employment at the NYPD[13], (2) all candidates who passed the NYPD entrance exam[14], and (3) those applicants who were successfully hired and appointed to the police academy. At the exam stage, the applicant pool already underrepresents Democrats relative to eligible voters (47% vs. 57%) and overrepresents Republicans (21% vs. 12%). The racial pattern, however, runs in the opposite direction. The share of Black candidates almost matches the

[13]I rely on age and residency information in the L2 voter data. For education, I use Bayesian imputation with census data from the 2019 American Community Survey based on the census tract.

[14]Unfortunately, the exam data only provides individuals placed on eligible civil service lists and does not provide the pool of unsuccessful candidates.

eligible pool (16% and 15%), Hispanics are overrepresented among exam passers (32% vs. 22%), and White candidates are *underrepresented* (41% vs. 47%). Because Democrats vastly outnumber Republicans among eligible voters, the partisan share gap reflects a Republican-to-Democrat application rate ratio of roughly 2.3 to 1, while the Black-White-gaps at this stage are close to parity in rate-ratio terms. Hence, self-selection into policing and screening through the written entry exam can explain some of the partisan, though not racial, sorting at entry.

The discretionary hiring stage adds important further sorting on both dimensions. Democrats make up only 35% of hired candidates, and Republicans account for 32% of successful applicants. Similar trends are observed for race, where the share of Whites increases by 9 percentage points between the application and the hiring stages. These higher hiring rates for Republican and White applicants are not explained by exam performance.[15] As Table 1 shows, among candidates of the same exam and similar scores, Republicans and White candidates are about eight percentage points and three percentage points more likely to be appointed than Democratic and Black applicants, respectively (Column 3); both gaps persist after observable controls, though the Black-White gap shrinks further to about one percentage point (Column 4). Hispanic and Asian candidates show no within-score hiring disadvantage relative to Whites. For Hispanic candidates in particular, the substantial application advantage seen in Figure 3 does not translate proportionally into hires, but this attenuation seems accounted for by lower exam scores rather than by discretionary screening. Figures A9 and A10 further indicate that these gaps persist across the entire score distribution, and that among hired exam-takers, Blacks and Democrats do not have substantially higher scores. Unlike work on women navigating barriers to political office, where those who break through are often positively selected on observable qualifications (Anzia and Berry, 2011; Ashworth et al., 2024), the underrepresented candidates who clear the NYPD's discretionary screens do not seem stronger on measurable ability than their Republican and White

[15]Note that White candidates achieve slightly higher scores on the exams (see Table A6).

counterparts.

Table 1: Differences in Hiring By Exam Taker Characteristics

	Model 1	Model 2	Model 3	Model 4
Republican	0.09***		0.08***	0.07***
	(0.00)		(0.00)	(0.00)
Non-Partisan	0.04***		0.03***	0.03***
	(0.00)		(0.00)	(0.00)
Black		−0.06***	−0.03***	−0.01***
		(0.00)	(0.00)	(0.00)
Hispanic		−0.02***	0.00	0.01**
		(0.00)	(0.00)	(0.00)
Asian		−0.00	0.01*	0.02**
		(0.01)	(0.01)	(0.01)
Other Race		−0.01**	0.00	0.01
		(0.00)	(0.01)	(0.01)
Exam*Score Bin FE	Yes	Yes	Yes	Yes
Additional controls	No	No	No	Yes
Mean of DV	0.13	0.13	0.13	0.13
Adj. R^2	0.04	0.03	0.04	0.05
Num. obs.	75134	72851	69406	64342

Linear probability regressions, weighted by the posterior probability of a match between exam data and voter file, and between exam data and appointments. Additional controls include gender, age, and county of residence. HC1 standard errors are in parentheses. $^{***}p < 0.001$; $^{**}p < 0.01$; $^{*}p < 0.05$

If exam performance does not explain who clears the post-exam hiring stages, what does? This is fundamentally a causal question, and my setup cannot test the importance of specific mechanisms. The data are, however, consistent with four features that may shape candidates' ability to sustain the process: wait time, outside options, financial constraints, and personal networks. First, the wait time for hiring at the NYPD is unusually long and the process is disqualification-heavy. After passing the entry exam, candidates can wait up to four years to be sworn into the police academy, which is substantially longer than at other NYC agencies, including the Corrections Department (Figure A11). Most of this gap reflects how long NYPD waits before requesting candidates from the eligible list (a median of

948 days vs. 282 days for the Corrections Department), rather than the post-call screening period itself. They are also more likely than applicants at peer agencies to be disqualified during the psychological, medical, and character investigations (Figure A12).

The costs of this lengthy and screening-heavy process are likely to be uneven and might give underrepresented candidates more reason and opportunity to withdraw before a possible appointment. Indeed, Democratic and non-White candidates are more likely than Republicans and Whites to take another civil service exam, and Black and Hispanic candidates in particular are more likely to be hired by the Department of Corrections (Table 2).[16] The DoC pipeline is far from a niche outside option: roughly 2,200 of the 6,500 entry-level Correction Officer appointments between 2014 and 2019 come from candidates who had previously sat the NYPD Police Officer exam, suggesting that outflow from the NYPD pipeline is a meaningful source of recruits for the city's other law enforcement agency. Hence, underrepresented groups gravitate towards outside options during a long, uncertain process, which is consistent with interview-based evidence that minority police applicants withdraw at higher rates than White applicants (Ryan et al., 2000; Ployhart et al., 2002; Kringen and Kringen, 2014).

Third, financial resources differ systematically across exam takers. Non-hired exam takers and those taking another civil service exam cluster at the lower end of the income distribution (Figure A13), with the pattern more pronounced among non-White and Democratic exam takers (Figure A14).[17] Lower income may constrain candidates' ability to absorb a long, uncertain process, particularly when alternatives in the broader civil service offer faster appointment timelines.

Finally, personal networks in law enforcement seem to shape hiring, and the returns to these networks are unequally distributed. Using applicants' home address and surnames, I

[16] To rule out mechanical effects, I exclude candidates who are hired by the NYPD within that year.

[17] Income is estimated from L2's household income measure and per-capita income from the 2019 American Community Survey at the census-tract level. Because L2 income is observed in 2021, hired candidates' income partly reflects post-appointment NYPD salary; the across-group comparisons among *non-hired* candidates are therefore the cleaner test.

Table 2: Differences in Alternative Careers (Among Non-Hired NYPD Exam Takers)

	Other Exam		Hired at DoC	
Republican	-0.010^{*}	0.001	-0.005^{**}	-0.004^{**}
	(0.004)	(0.004)	(0.001)	(0.002)
Non-Partisan	-0.014^{***}	-0.006	-0.004^{**}	-0.003^{*}
	(0.003)	(0.004)	(0.001)	(0.001)
Black	0.084^{***}	0.063^{***}	0.016^{***}	0.016^{***}
	(0.005)	(0.005)	(0.002)	(0.002)
Hispanic	0.032^{***}	0.022^{***}	0.012^{***}	0.012^{***}
	(0.003)	(0.004)	(0.001)	(0.002)
Asian	0.039^{***}	0.024^{**}	0.002	0.002
	(0.006)	(0.008)	(0.002)	(0.002)
OtherRace	0.049^{***}	0.021^{**}	0.005^{**}	0.004
	(0.006)	(0.007)	(0.002)	(0.002)
NYPD Exam*Score Bin FE	Yes	Yes	Yes	Yes
Additional controls	No	Yes	No	Yes
Mean of DV	0.149	0.147	0.019	0.019
Adj. R^2	0.035	0.043	0.020	0.019
Num. obs.	62068	57390	60249	55760

Linear probability regressions, weighted by the posterior probability of a match between the respective data sets. Dependent Variables: Taking another civil service exam within a year of NYPD exam (Model 1,2); Being appointed as corrections officer (Model 3,4). Additional controls include gender, age, and county of residence. HC1 standard errors in parentheses. $^{***}p < 0.001$; $^{**}p < 0.01$; $^{*}p < 0.05$.

measure each candidate's connection to NYPD personnel through the concentrations of current officers or recent appointees in their census tracts and through surname-matches with a retired officer. Police networks are denser among Republicans and Whites. For example, Whites (Republicans) live in census blocks with an average of 31 (33) active police officers, compared to only 24 (23) for Democrats (non-Whites). Additionally, the advantages of personal connections are markedly stronger for White and Republican applicants (Figure 4). For these groups, a doubling of the number of current NYPD officers residing in their neighborhood correlates with a 3-percentage-point increase in hiring probability — a 30% relative gain off a 10% base rate. Democratic and Black applicants show much smaller and often negligible benefits from such neighborhood connections, and instead face marginally lower

Figure 4: Personal Networks and Hiring Probabilities

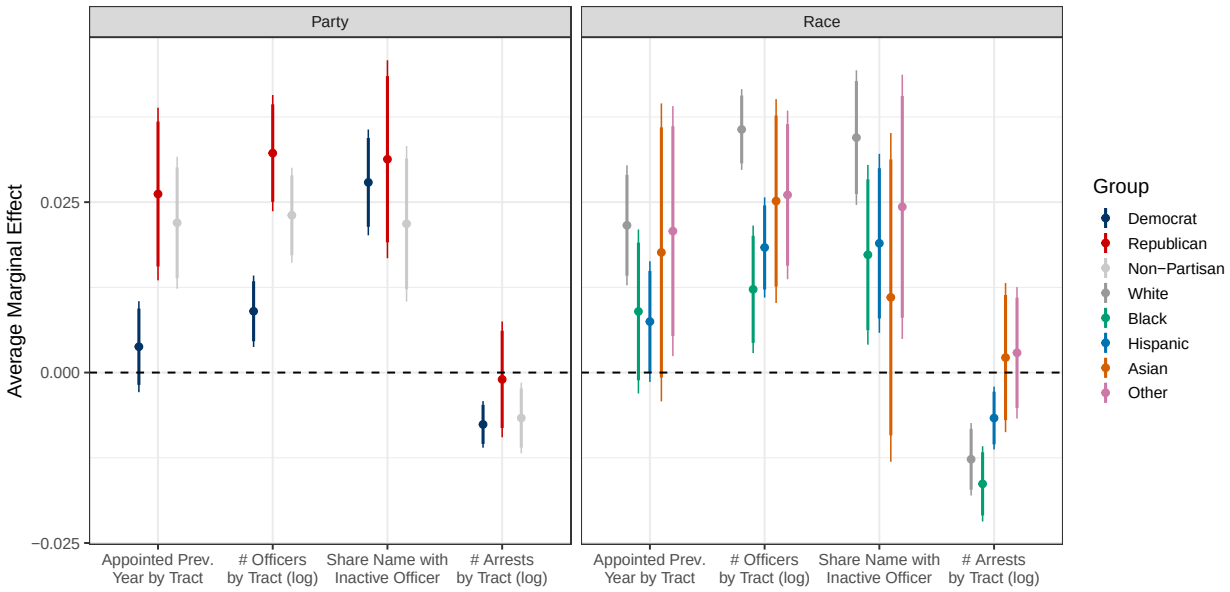

Note: All regressions include exam and score bin fixed effects. Models (1), (2), and (4) further control for the population size per tract. Model (3) controls for name commonality using 2010 census data.

hiring probabilities when living in tracts with higher arrest rates. These findings suggest that informal networks may contribute to the observed disparities in who successfully navigates the NYPD hiring process.

5.2 Career Progression

Career progression at the NYPD, like its hiring process, runs through both formalized and discretionary stages: exam-based promotions to civil service ranks on the one hand, grade changes within the detective track, elite-unit assignments, and departmental awards on the other. The within-pipeline contrast that emerged at entry reappears in career progression.

Figure 5 provides a cross-sectional starting point. Using data on all uniformed employees and traffic enforcement agents on the NYPD payroll (2014–2021), it indicates that Democrats, particularly Black Democrats, are clustered among the lower ranks. For instance, among traffic enforcement agents 33% are Black Democrats, whereas only 6% of Lieutenants

or above are Black and Democratic.[18] For White Republicans, in contrast, these figures amount to 2% and 35%, respectively.

Figure 5: Share of Demographics by Police Rank

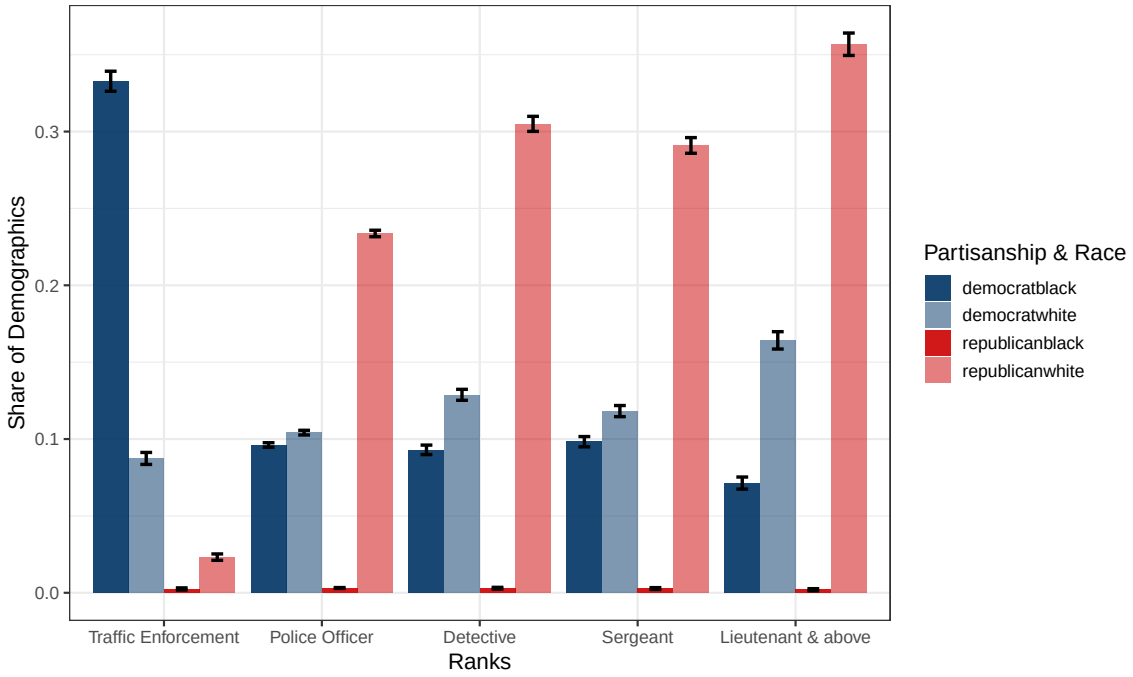

Yet, since these snapshots do not account for compositional differences across cohorts, Table 3 estimates disparities in promotion probabilities across partisanship and race within the same cohorts. Models (1)–(3) include only official promotions that require a promotion exam (i.e., sergeant and above), Models (4)–(6) further include discretionary promotions such as detective grade changes. The estimates do not provide clear evidence that Republicans and Whites are systematically more likely than Democrats or Black officers to receive *official* promotions.[19] Two further checks reinforce this null result at the formalized stage. First, Table A10 shows that scores on the promotion exams are a strong predictor of receiving a promotion, and once exam performance is accounted for the small promotion gaps subside,

[18]Traffic enforcement agents (TEAs) are a separate civil service title with their own entry exam and hiring track, not a lower step on the police-officer career ladder. The TEA-to-lieutenant comparison thus reflects both between-track sorting at hiring and within-track promotion. The regression analyses below restrict to sworn officers.

[19]Coefficients for other racial groups are small but cut in different directions, with Hispanic officers about 1 percentage point less likely and Asian and other-race officers 3–4 percentage points more likely than Whites to receive an official promotion.

Table 3: Differences in Promotions by Officer Characteristics

	Official Promotions			Title Changes		
	Model 1	Model 2	Model 3	Model 4	Model 5	Model 6
Republican	−0.00	−0.00	−0.00	0.01*	0.01	0.00
	(0.00)	(0.00)	(0.00)	(0.00)	(0.00)	(0.00)
Non-Partisan	−0.00	−0.00	−0.00	−0.00	−0.00	−0.00
	(0.00)	(0.00)	(0.00)	(0.00)	(0.00)	(0.00)
Black	−0.01	−0.01	−0.01	−0.03***	−0.02***	−0.03***
	(0.00)	(0.00)	(0.00)	(0.01)	(0.01)	(0.01)
Hispanic	−0.01**	−0.01**	−0.01**	−0.02***	−0.02***	−0.02***
	(0.00)	(0.00)	(0.00)	(0.00)	(0.00)	(0.00)
Asian	0.04***	0.03***	0.04***	−0.00	−0.01	−0.00
	(0.01)	(0.01)	(0.01)	(0.01)	(0.01)	(0.01)
Other Race	0.03***	0.03***	0.03***	−0.00	−0.01	−0.01
	(0.01)	(0.01)	(0.01)	(0.01)	(0.01)	(0.01)
Cohort FE	Yes	Yes	Yes	Yes	Yes	Yes
Exposure FE	Yes	Yes	Yes	Yes	Yes	Yes
Age, Gender	No	Yes	Yes	No	Yes	Yes
CCRB complaints	No	No	Yes	No	No	Yes
Mean of DV	0.09	0.09	0.09	0.17	0.17	0.17
Adj. R^2	0.10	0.10	0.11	0.19	0.19	0.20
Num. obs.	45990	43573	43573	45990	43573	43573

Linear probability regressions, weighted by the posterior probability of a payroll and voter file match. Level of observation: Uniformed employee. Outcome: Dummy for whether the employee received a promotion/title change between 2014 and 2021. All models control for the duration an employee is observed on the payroll since 2014 (Exposure time). Some models also control for age at appointment and gender, as well as total number of CCRB complaints over an officer's full tenure. HC1 standard errors in parentheses. $^{***}p < 0.001$; $^{**}p < 0.01$; $^{*}p < 0.05$

consistent with Republicans and Whites performing slightly better on promotion exams rather than benefiting from differential treatment after the exam. Second, Table A9 indicates no substantial gaps in promotional aspirations, with Republican and White officers no more likely to take a promotion exam throughout their tenure than Democrats and non-Whites.[20]

However, the racial gaps in career progressions are larger and significant when incorporating discretionary promotions and grade changes. Whites are estimated to be 3 and 2 pp more likely to receive both horizontal and vertical career movements than Blacks and

[20]The exception is Asian officers, who are more likely to take promotion exams.

Figure 6: Seniority Gap by Years of Experience

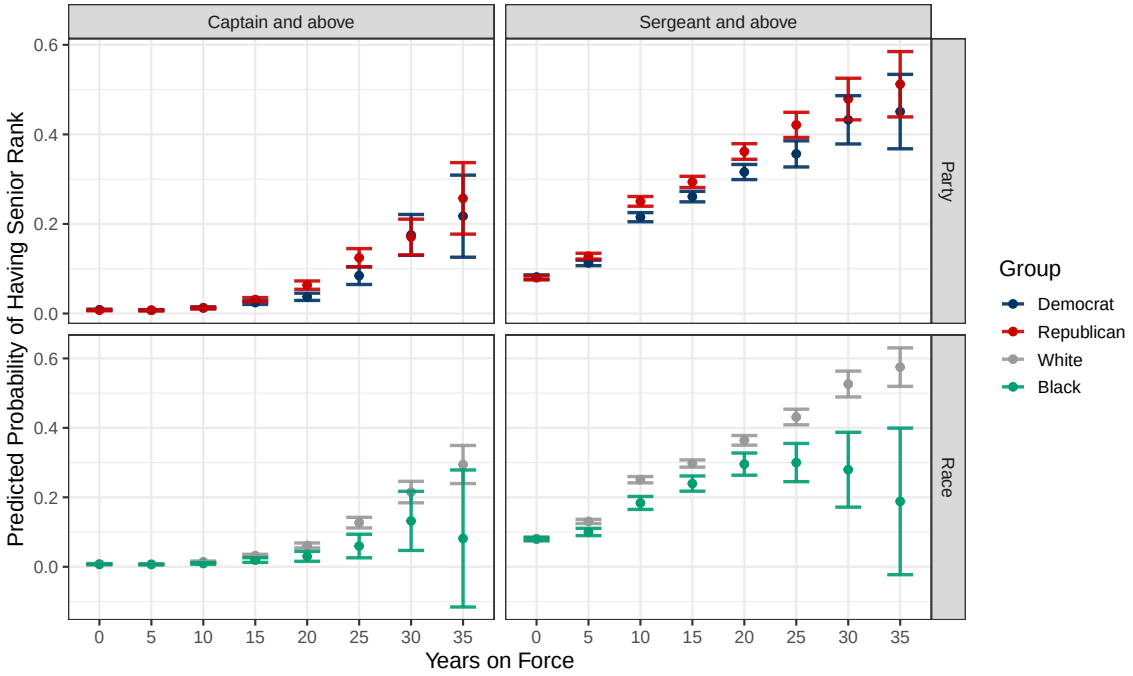

Depicted are predicted probabilities of having a senior rank, with 95% HC1 confidence intervals. All underlying regression models (LPM) include officer cohort fixed effects.

Hispanics, respectively. With an overall promotion rate of 17%, these estimates are non-trivial. The partisan gap in promotions is small and sensitive to age and gender controls, but emerges more clearly in outcomes with even less structure. Republican and White officers are more likely to receive departmental awards and to be assigned to prestigious elite units, including anti-terrorism and special forces (Tables A11 and A12). Similar discretion in team assignments produces parallel patterns of sorting: teams headed by non-White (Democratic) leaders have lower shares of lower-ranked White (Republican) members and higher racial diversity overall (Tables A13 and A14).[21]

Figure 6 further assesses whether this partisan and racial seniority gap persists across officers' tenure. The Black-White gap endures and widens over time, whereas Democrats seem to catch up to the ranks of Republicans after 30 years on the force. Taken together, this suggests that Republican and especially White officers benefit from steeper progress along

[21] These analyses use a cross-section of active officers (as of October 2021) for whom more detailed information on assignments and awards is available.

discretionary parts of the career ladder in law enforcement.

5.3 Attrition

How do these differences in career trajectories translate into attrition? Like earlier stages of the pipeline, attrition has both formalized and discretionary parts. Voluntary retirements at the 20-year mark are tied to formal pension eligibility, whereas dismissals and terminations are discretionary agency actions. Figure 7 depicts the distribution of years on the force at the time officers retire. Evidently, most officers retire around 20 and 25 years of service, when they become eligible for different retirement packages at NYPD. Yet, Republicans and White employees stay on the force slightly longer, often working beyond the retirement age of 20 years on the force. For instance, the median retiring Republican or White officer worked for 22.1 years, compared to 21.3 years for Democratic officers and 21 years for non-White officers. As Table A15 shows, these results hold when accounting for officers' age at appointment.

The discretionary agency side of attrition shows the sharpest demographic gaps. Across all types of officers, retirements and resignations account for the bulk of exits (Figure 8). But involuntary exits (dismissals and terminations), where supervisors and the department decide who is removed, account for significantly higher shares of exits among Democratic, Black, and Hispanic officers than among Republicans and Whites. Specifically, 2% of exiting Democrats, 3.3% of Blacks, and 2.4% of Hispanics leave involuntarily, compared to 1.3% of Republicans and 0.9% of White officers. The pattern is most pronounced among early-career officers (Figures A16 and A17), where tenure protections are weakest.[22]

[22]Voluntary resignations show no comparable gap, suggesting the early-career pattern reflects agency-side removal decisions rather than officer-side decisions to leave.

Figure 7: Distribution of Years on Force at Retirement

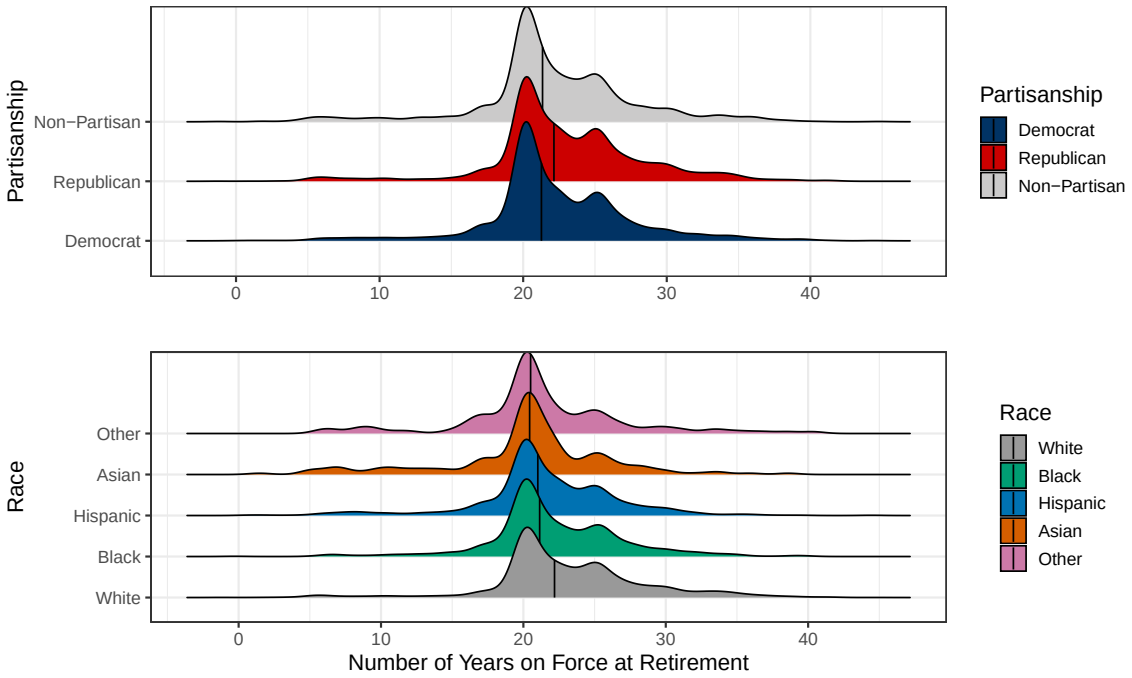

Black lines indicate median years on the force at retirement.

Figure 8: Predicted Probability of Exit Type

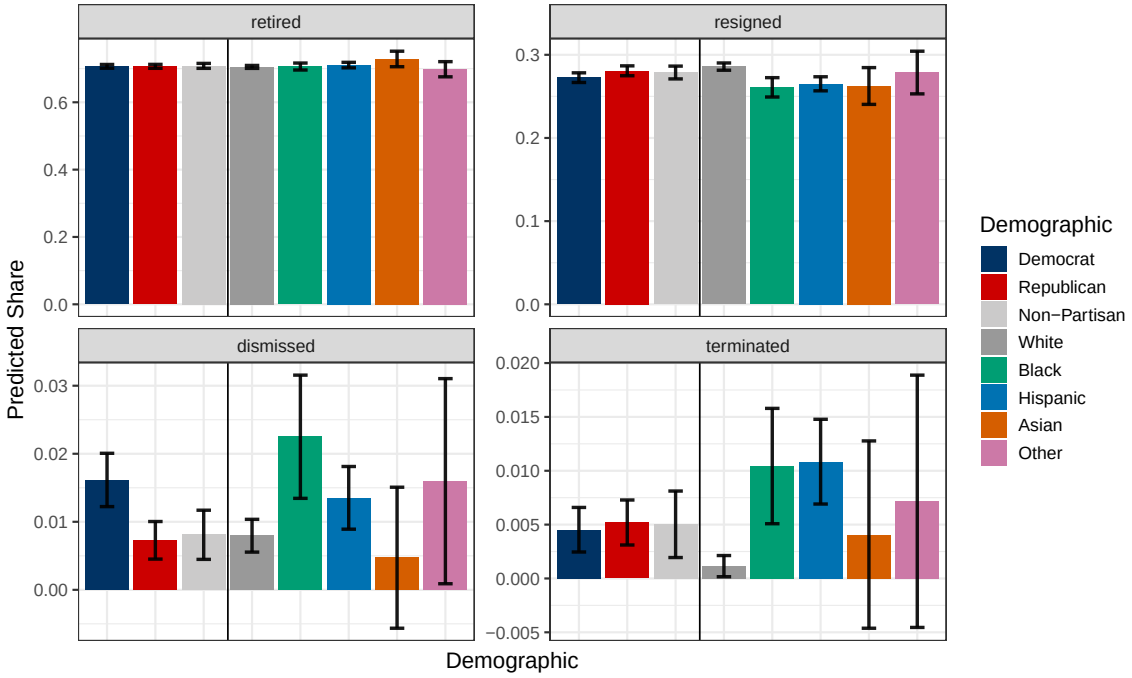

The estimates are obtained from regressions of exit type (conditional on exit) on group demographic, cohort FE and fiscal year FE. The covariates are fixed at their observed values for predictions.

6 Robustness

6.1 Merit and Ability

A key claim of this paper is that the observed representational disparities in the NYPD are not fully explained by differences in ability or merit. However, one may be concerned that unobserved, latent differences across groups could be confounding the estimates of demographic gaps at various stages of the pipeline. For instance, if Black officers have worse on-the-job performance than their White counterparts once hired, this could reflect unobserved differences in latent ability or motivation that also influence promotion and retention outcomes. Yet, several observable implications are inconsistent with this concern. First, recall from Table 2 that among NYPD exam takers who are not hired at NYPD, Democratic, Black and Hispanic candidates are significantly *more* likely to get hired at the Department of Corrections than their Republican and White counterparts, conditional on the same NYPD exam score. This suggests that these candidates are clearly appointable individuals for law enforcement work, and that the same candidates with the same measured ability end up in different agencies depending on which discretionary process they go through. Second, hired Republican and White officers do not have better, if anything slightly *worse*, disciplinary records on the job, as measured by CCRB complaints three years after hiring (Figure A18). While this analyses cannot account for officer assignments or civilian complaint behavior, it cuts against the idea that these officers are higher quality to begin with. Third, one may wonder whether the divergence in promotions across officers' tenures is due to attrition by Black officers with better outside prospects. Yet, Figure A15 suggests that there is no widening gap between Black and White officers across tenure in terms of observables that correlate with quality (i.e., the number of arrests, awards, and disciplinary records). Fourth, the demographic gaps in promotions and exits persist even after controlling for disciplinary records as a rough proxy for ability. Taken together, while no single piece of evidence is conclusive, these patterns make it difficult to attribute the full hiring, promotion

and exit gaps to unobserved ability differences.

6.2 Measurement Concerns

A few additional measurement concerns warrant discussion. First, I match applicants and employees to a single L2 snapshot from 2021, which may not perfectly reflect their partisanship at the time of application or hiring. Becoming a law enforcement officer could itself shift political preferences over time, conflating selection with a treatment effect (Goncalves and Tuttle, 2024). To assess these concerns, I leverage L2's tracking of party changes over the four years preceding the snapshot. Among all NYC voters registered in November 2021, only about 2% switched between the Democratic and Republican parties.[23] Restricting to voters in the 2016–2020 exam cohort, where the change window covers the entire post-exam period, I similarly identify only 2.4% of switchers. Dropping these switchers does not meaningfully alter the partisan gaps reported here (see Table A7 for the hiring analysis).

Second, since New York is a closed-primary state, which restricts primary participation to registered party members, one might worry that Republican registrations among officers are inflated by primary participation incentives in the 2020 election rather than reflecting underlying preferences. For this concern to bias the police-civilian comparison, however, police would need to respond differentially to closed-primary incentives beyond the mobilization alongside the general NY population. To assess this, I restrict the analysis to voters who first registered before 2018 and who have not switched between Democratic and Republican registrations. This combined filter excludes both new registrants who entered the system around the 2020 cycle and long-term registrants who only recently switched into the Republican party. The partisan gap between police and civilians is very similar in magnitude in this restricted sample (see Table A8 for hiring), suggesting the patterns are not driven by strategic 2020 registration.

Third, my measurement of voters' race relies on L2's proprietary imputation algorithm,

[23] L2's change field captures only Democratic-Republican switches; switches involving Non-Partisan or third-party affiliations are not flagged, so the 2% figure is a lower bound on total party switching.

which combines surname information with census-block demographics in the spirit of standard Bayesian Improved Surname Geocoding (BISG) (Imai and Khanna, 2016). To assess its validity, I compare my estimates against (1) official counts of race groups reported by the NYPD to the 2020 Law Enforcement Management and Administrative Statistics (LEMAS) survey and (2) measures of employees' race using Bayesian predictions (Imai and Khanna, 2016) in Appendix A.7. Overall, the results suggest that, if anything, I likely *underestimate* the share of White groups in the NYPD, and the true representational gaps between Whites and non-Whites may be even larger than reported. Among non-White employees, the L2-based estimates have a disproportionately higher share of "Other/unknown race." If these unknowns are predominantly Black, Hispanic, or Asian, the gaps for specific groups (especially Black) may be smaller than reported here, though this does not affect the overall White-vs-non-White comparison.

7 Counterfactual Analysis

The analyses so far reveal that the sorting of Republicans and Whites within the NYPD emerge at various stages of the selection pipeline. What is the relative impact of each stage—attraction, selection, promotion, and retention—on these representational disparities? To what extent could the NYPD hypothetically mitigate these gaps by adjusting their policies at each stage?

To address these questions, I provide a simple counterfactual analysis in the spirit of Chetty et al. (2023). This analysis estimates the potential increase in the number of Democrats and non-Whites on the NYPD payroll under scenarios where disparities in application rates, hiring rates, career progressions, and retention rates are eliminated one at a time. It's important to note that this simulation is not designed to predict actual outcomes of policy interventions, as it does not account for behavioral responses or general equilibrium effects that would likely occur in practice. Additionally, this exercise relies on observed differences in rates across groups that may themselves be influenced by confounding factors

not captured in the data. The analysis should therefore be interpreted as quantifying the relative importance of each selection stage using observed disparities rather than isolating the causal impact of reforming selection procedures. Utilizing the estimates of application rates, hiring rates, rank assignments, and attrition rates from Section 5, I calculate these counterfactual outcomes in the following way.

Applications and Passing of First Entry Exam. Using Equation (1), I calculate the actual and the counterfactual number of appointments from underrepresented groups u (e.g., Democrats or Blacks). I compute counterfactual entry by replacing underrepresented groups' apply or hire rate with the corresponding rate among overrepresented groups, holding the other stage fixed

$$\text{Entry}_u = \text{N Eligible}_u \times \text{ApplyRate}_u \times \text{HiringRate}_u \tag{1}$$

$$\text{Counterfactual Entry}_u = \text{N Eligible}_u \times \text{ApplyRate}_o \times \text{HiringRate}_u \tag{2}$$

This analysis is subject to a few important caveats: First, for partisanship, I have to approximate the number of eligible individuals using information on the number of *NYC voters* in each group who satisfy the minimum application requirements at the NYPD (qualifying age, education, residency). For race, I use estimated counts of eligible citizens by group from the 2019 Census Bureau's American Community Survey (ACS).[24] Additionally, due to the nature of my exam data, the variable *ApplyRate* inherently combines both the willingness to apply and the ability to pass the entry exam.

Hiring. Similarly, I calculate the counterfactual number of appointments in the absence of disparities in hiring rates by assuming that each applicant from underrepresented groups u is hired at the same rate as aspirants from overrepresented demographics o.

$$\text{Counterfactual Entry}_u = \text{N Eligible}_u \times \text{ApplyRate}_u \times \text{HiringRate}_o \tag{3}$$

[24] https://www.census.gov/programs-surveys/acs/microdata.html

Table 4: Counterfactual Analysis - Entry

	N Eligible	**N Applicants**	**Apply Rate**	**Apply Rate CF**	**Hire Rate**	**Hire Rate CF**	**Actual Entry**	**CF Entry Apply**	**CF Entry Hire**
Democrats	1073362	6572	0.006	0.014	0.104	0.185	683	1596	1213
Blacks (L2)	298594	2062	0.007	0.009	0.104	0.132	215	268	273
Blacks (Census)	354659	2062	0.006	0.005	0.104	0.132	215	198	273

The CF rates for Democrats (Blacks) are taken from Republican (White) applicants. *N Applicants* is the total number of applications observed during the sample window divided by the years observed (13,400 per year), times the share of the respective demographic from Figure 3. *ApplyRate* is estimated as *N Applicants*/*N Eligible*. *HireRate* is estimated from predicted hiring probabilities based on Table 1, Column 3.

Table 4 shows the results of this exercise. If application and hiring rates remain unchanged, the NYPD is expected to add an average of 683 Democrats and 215 Black aspirants annually. For partisanship, the application stage is the dominant lever. Equalizing application rates more than doubles Democratic recruitment to 1,596 (+134%), and more than the 1,213 Democrats (+78%) added by equalizing hiring rates. This is because the Republican-to-Democrat rate ratio is bigger at apply (2.3×) than at hire (1.8×). For race, hire-rate equalization is the more reliable lever. Whether application-rate equalization helps depends on how the eligible pool is measured. Under the L2-based estimate, it increases Black recruits by 25% (close to the 27% from equalizing hire rates), while under the more inclusive ACS Census estimate it slightly *worsens* representation because Black applicants already apply at higher rates relative to their eligible pool than Whites. Hire-rate equalization yields a 27% increase across both eligible-pool measures.

Retention. Following a similar logic, I also estimate how many additional underrepresented officers the NYPD could retain by eliminating differences in career progression and attrition rates, and again compare these quantities to the actual predicted number of retained officers. For this, I use the predicted rate of officers from each group to be assigned

to rank r and retained, conditional on officers' years of service (binned) and age.

$$\text{Retention}_u = \text{N Officers}_u \times \sum_r \text{Rank}_{r,u} \times \text{RetentionRate}_{r,u} \tag{4}$$

$$\text{Counterfactual Retention}_u = \text{N Officers}_u \times \sum_r \text{Rank}_{r,o} \times \text{RetentionRate}_{r,u} \tag{5}$$

$$\text{Counterfactual Retention}_u = \text{N Officers}_u \times \sum_r \text{Rank}_{r,u} \times \text{RetentionRate}_{r,o} \tag{6}$$

In contrast, equalizing retention rates and rank distributions has only marginal effects on the composition of the force. Even if the NYPD were to retain Democrats and Blacks at the same rates as Republicans and Whites, the net change would amount to fewer than 20 additional officers per year (Table 5). This is primarily because overall retention rates are already high across groups (approximately 95%), leaving little room for improvement.

Table 5: Counterfactual Analysis - Retention

	N Officers	**Actual Retention**	**CF Retention Ranks**	**CF Retention Rate**
Democrats	11996	11401	11404	11386
Blacks	3427	3259	3263	3239

The CF rates for Democrats (Blacks) are taken from Republican (White) officers. *N Officers* is the average number of uniformed employees on the payroll across years in my sample by the respective demographic.

Overall, this counterfactual exercise illustrates that the most effective lever for enhancing the representativeness of one of America's largest law enforcement agencies is equalizing hiring and applications rates across under- and overrepresented groups.

8 Conclusion

The race and partisanship of street-level bureaucrats are central to issues of representative bureaucracy and crucial determinants of public service provision (Ba et al., 2021, 2025; Donahue, 2023). This article decomposes the selection dynamics of modern civil service systems that shape the political representativeness of city bureaucracies, focusing on the

bureaucratic pipeline of the NYPD. Whereas previous studies focus on political cycles in top-down appointments and the "hidden curriculum" in written entry exams to explain the partisan and racial composition of bureaucracies, this article shows how multiple stages of self-selection, recruitment, promotion, and attrition combine to produce bureaucratic composition, with the largest gaps emerging at the discretionary stages that follow the formalized civil service exams.

These findings reframe a long-standing debate over civil service reform and meritocracy in public hiring. Both reformers and critics of civil service have concentrated their attention on the written entry exam. Reformers see it as the central diversification lever (Van Riper et al., 1963; Hoogenboom, 1968), while critics see it as the principal site at which ostensibly neutral rules reproduce existing inequalities (Portillo et al., 2020; Kuipers and Sahn, 2023; Moreira and Pérez, 2025). The pipeline decomposition presented here suggests that this shared focus has treated civil service as a binary system, missing how meritocratic rules and the discretion they leave intact interact across the full bureaucratic career. Reform efforts that target a single formal procedure, however carefully designed, are unlikely to close gaps that are produced cumulatively across many discretionary decisions over a career.

The findings also reframe how we think about representative bureaucracy. A growing literature studies the consequences of bureaucrats' demographic and partisan identity for service provision (Bradbury and Kellough, 2011; Ba et al., 2021, 2025; Donahue, 2023), but the findings are often inconclusive and the conditions under which representation translates into more equitable outputs remain unsettled (Meier, 2018). Understanding the *origins* of representation is critical to reconcile this ambiguity, because the same level of demographic composition can produce different outcomes depending on how that composition is achieved. At the NYPD, the discretionary stages that drive most of the variation in composition do not appear to select on measurable qualifications. Democratic and non-White officers who successfully navigate these stages may instead be those who most closely align with the prevailing organizational culture. Even those who do not enter pre-selected on alignment may

face strong incentives to conform once inside, since advancement in a predominantly White and Republican environment continues to depend on discretionary judgments by superiors who share that culture. Simply achieving demographic diversity may therefore be insufficient if both the selection processes that admit minority officers and the institutional incentives they face favor those aligned with established norms. To theorize the effects of demographic shifts in police forces, and bureaucracy more broadly, we need to incorporate the mechanisms through which representation arises and is sustained (Ashworth et al., 2024).

Yet, important issues and open questions remain. Although the large administrative data from NYC allows for a rich picture of partisan and racial selection, the analysis is restricted to only one primary US jurisdiction, and it remains to be seen how the results generalize to other city governments. Two features of the case nonetheless suggest that the broader insight may travel. First, the institutional structure that produces the asymmetry is shared across most large US civil service agencies. Formalized exam-based stages coexist with discretionary vetting, promotion, assignment, and termination decisions in police, fire, corrections, and other local bureaucracies across the country (Riccucci and Riccardelli, 2015; Anzia and Trounstine, 2025). Second, NYC is a hard case for substantial political sorting. The NYPD's civil service rules are well-codified and its hiring pipeline includes substantial formalization. If sorting nonetheless concentrates at the discretionary stages within this institutional setting, the same gradient may operate elsewhere and even more sharply where formal rules are thinner.

Additionally, this article is inherently descriptive and cannot uncover the exact mechanisms that lead to the documented disparities in bureaucrats' partisanship and race. For instance, the fact that Republican and White officers are more likely to receive promotions and desirable assignments could be explained by individual-level factors, such as job satisfaction and motivation to advance in the profession, or institutional-level aspects, including support from superiors and the agency more broadly. Future research may seek to use other data sources and methods, such as surveys of street-level bureaucrats and experimental de-

signs, to further explore the mechanisms at play.

References

Anzia, Sarah F. and Christopher R. Berry. 2011. "The jackie (and jill) robinson effect: Why do congresswomen outperform congressmen?". *American Journal of Political Science* 55 (3): 478–493.

Anzia, Sarah F. and Jessica Trounstine. 2025. "Civil Service Adoption in America: The Political Influence of City Employees". *American Political Science Review* 119 (2): 549–565.

Ashworth, Scott, Christopher R. Berry, and Ethan Bueno de Mesquita. 2024. "Modeling Theories of Women's Underrepresentation in Elections". *American Journal of Political Science* 68 (1): 289–303.

Ba, Bocar, Haosen Ge, Jacob Kaplan, Dean Knox, Mayya Komisarchik, Gregory Lanzalotto, Rachel Mariman, Jonathan Mummolo, Roman Rivera, and Michelle Torres. 2025. "Political Diversity in U.S. Police Agencies". *American Journal of Political Science* 69 (4): 1617–1635.

Ba, Bocar A., Dean Knox, Jonathan Mummolo, and Roman Rivera. 2021. "The role of officer race and gender in police-civilian interactions in Chicago". *Science* 371 (6530): 696–702.

Benton, Mark. 2020. "Representation Is Not Enough: Symbolic Representation and Perceptions of the Police". *Administration & Society* 52 (5): 794–822.

Besley, Timothy and Maitreesh Ghatak. 2005. "Competition and Incentives with Motivated Agents". *American Economic Review* 95 (3): 616–636.

Bonica, Adam, Howard Rosenthal, Kristy Blackwood, and David J. Rothman. 2020. "Ideological Sorting of Physicians in Both Geography and the Workplace". *Journal of Health Politics, Policy and Law* 45 (6): 1023–1057.

Bradbury, Mark and J. Edward Kellough. 2011. "Representative Bureaucracy: Assessing the Evidence on Active Representation". *The American Review of Public Administration* 41 (2): 157–167.

Brierley, Sarah, Noah L. Nathan, George Kwaku Ofosu, and Tingxuan Zhu. 2026. "The Seeds of State Capture: Merit and Patrimonialism in the Colonial Bureaucracy". Working Paper.

Brunson, Rod K. and Jacinta M. Gau. 2015. "Officer Race Versus Macro-Level Context: A Test of Competing Hypotheses About Black Citizens' Experiences With and Perceptions of Black Police Officers". *Crime & Delinquency* 61 (2): 213–242.

Castilla, Emilio J. and Stephen Benard. 2010. "The Paradox of Meritocracy in Organizations". *Administrative Science Quarterly* 55 (4): 543–576.

Chetty, Raj, David J. Deming, and John N. Friedman. 2023. "Diversifying Society's Leaders? The Determinants and Causal Effects of Admission to Highly Selective Private Colleges". NBER Working Paper 31492, National Bureau of Economic Research.

Chinoy, Sahil and Martin Koenen. 2024. "Political Sorting in the U.S. Labor Market: Evidence and Explanations". Working Paper. https://sahilchinoy.s3.us-west-1.amazonaws.com/chinoy_politics_work.pdf.

Colonnelli, Emanuele, Valdemar Pinho Neto, and Edoardo Teso. 2025. "Politics at work". *American Economic Review* 115 (10): 3367–3414.

Donahue, Samuel Thomas. 2023. "The Politics of Police". *American Sociological Review* 0 (0): 1–25.

Enamorado, Ted, Ben Fifield, and Kosuke Imai. 2019. "Using a Probabilistic Model to Assist Merging of Large-scale Administrative Records". *American Political Science Review* 113 (2): 353–371.

Folke, Olle, Shigeo Hirano, and James M. Snyder. 2011. "Patronage and Elections in U.S. States". *The American Political Science Review* 105 (3): 567–585.

Fox, Richard L. and Jennifer L. Lawless. 2011. "Gendered perceptions and political candidacies: A central barrier to women's equality in electoral politics". *American Journal of Political Science* 55 (1): 59–73.

Gaudette, Jennifer. 2025. "Polarization in police union politics". *American Journal of Political Science* 69 (3): 961–980.

Gerring, John. 2012. "Mere Description". *British Journal of Political Science* 42 (4): 721–746.

Goncalves, Felipe and Cody Tuttle. 2024. "Police Work and Political Identity". Working Paper. https://codytuttle.github.io/policework_draft.pdf.

Grosjean, Pauline, Federico Masera, and Hasin Yousaf. 2023. "Inflammatory Political Campaigns and Racial Bias in Policing". *The Quarterly Journal of Economics* 138 (1): 413–463.

Harris, Allison P. 2023. "Can Racial Diversity Among Judges Affect Sentencing Outcomes?". *American Political Science Review, conditionally accepted*: 1–16.

Harvey, Anna and Taylor Mattia. 2022. "Reducing racial disparities in crime victimization: Evidence from employment discrimination litigation". *Journal of Urban Economics*: 103459.

Hassan, Mai, Horacio Larreguy, and Stuart Russell. 2024. "Who Gets Hired? Political Patronage and Bureaucratic Favoritism". *American Political Science Review* 118 (4): 1913–1930.

Hofstadter, Richard. 1955. *The Age of Reform: From Bryan to FDR.* New York: Vintage Books.

Hoogenboom, Ari Arthur. 1968. *Outlawing the Spoils: A History of the Civil Service Reform Movement, 1865–1883*. Urbana: University of Illinois Press.

Imai, Kosuke and Kabir Khanna. 2016. "Improving Ecological Inference by Predicting Individual Ethnicity from Voter Registration Records". *Political Analysis* 24 (2): 263–272.

Kringen, Anne Li and Jonathan Allen Kringen. 2014. "Identifying Barriers to Black Applicants in Police Employment Screening". *Policing: A Journal of Policy and Practice* 9 (1): 15–25.

Kuipers, Nicholas and Alexander Sahn. 2023. "The Representational Consequences of Municipal Civil Service Reform". *American Political Science Review* 117 (3): 982–998.

Lerman, Amy E. and Joshua Page. 2015. "Does the Front Line Reflect the Party Line? The Politicization of Punishment and Prison Officers' Perspectives Towards Incarceration". *The British Journal of Criminology* 56 (3): 578–601.

Lewis, David E. 2008. *The Politics of Presidential Appointments: Political Control and Bureaucratic Performance*. Princeton: Princeton University Press.

Linos, Elizabeth. 2017. "More Than Public Service: A Field Experiment on Job Advertisements and Diversity in the Police". *Journal of Public Administration Research and Theory* 28 (1): 67–85.

Lipsky, Michael. 1980. *Street Level Bureaucracy: Dilemmas of the Individual In Public Services*. Russell Sage Foundation.

McConnell, Christopher, Yotam Margalit, Neil Malhotra, and Matthew Levendusky. 2018. "The economic consequences of partisanship in a polarized era". *American Journal of Political Science* 62 (1): 5–18.

McCrary, Justin. 2007. "The Effect of Court-Ordered Hiring Quotas on the Composition and Quality of Police". *American Economic Review* 97 (1): 318–353.

Meier, Kenneth J. 2018, 06. "Theoretical Frontiers in Representative Bureaucracy: New Directions for Research". *Perspectives on Public Management and Governance* 2 (1): 39–56.

Miller, Amalia R and Carmit Segal. 2018, 09. "Do Female Officers Improve Law Enforcement Quality? Effects on Crime Reporting and Domestic Violence". *The Review of Economic Studies* 86 (5): 2220–2247.

Moreira, Diana and Santiago Pérez. 2025. "Who Benefits from Meritocracy?". Working Paper.

Nicholson-Crotty, Sean, Jason A Grissom, Jill Nicholson-Crotty, and Christopher Redding. 2016, 04. "Disentangling the Causal Mechanisms of Representative Bureaucracy: Evidence From Assignment of Students to Gifted Programs". *Journal of Public Administration Research and Theory* 26 (4): 745–757.

Ornaghi, Arianna. 2019. "Civil Service Reforms: Evidence from U.S. Police Departments". Working Paper. https://websitearianna.s3.eu-west-2.amazonaws.com/190702_ornaghi_civil_service_reforms.pdf.

Peters, B. Guy. 2015. "Political Patronage, Machine Politics and Ethnic Representativeness in the Public Sector". In *The Politics of Representative Bureaucracy*, pp. 113–121. Cheltenham: Edward Elgar.

Ployhart, Robert E., Lynn A. McFarland, and Ann Marie Ryan. 2002. "Examining Applicants' Attributions for Withdrawal From a Selection Procedure". *Journal of Applied Social Psychology* 32 (11): 2228–2252.

Portillo, Shannon, Domonic Bearfield, and Nicole Humphrey. 2020. "The Myth of Bureaucratic Neutrality: Institutionalized Inequity in Local Government Hiring". *Review of Public Personnel Administration* 40 (3): 516–531.

Rauch, James E and Peter B Evans. 2000. "Bureaucratic structure and bureaucratic performance in less developed countries". *Journal of Public Economics* 75 (1): 49 – 71.

Reid, Joseph D. and Michael M. Kurth. 1992. "The Rise and Fall of Urban Political Patronage Machines". In *Strategic Factors in Nineteenth Century American Economic History: A Volume to Honor Robert W. Fogel*, pp. 427–445. Chicago: University of Chicago Press.

Reny, Tyler T., Marcel F. Roman, Benjamin J. Newman, and David O. Sears. 2025. "Late adolescents entering college intending a career as police officers hold more right-leaning views than their peers". *Proceedings of the National Academy of Sciences* 122 (32): e2500220122.

Reskin, Barbara F. 2003. "Including Mechanisms in Our Models of Ascriptive Inequality". *American Sociological Review* 68 (1): 1–21.

Riccucci, Norma M. and Margaret Riccardelli. 2015. "The Use of Written Exams in Police and Fire Departments: Implications for Social Diversity". *Review of Public Personnel Administration* 35 (4): 352–366.

Riccucci, Norma M, Gregg G Van Ryzin, and Karima Jackson. 2018, 06. "Representative Bureaucracy, Race, and Policing: A Survey Experiment". *Journal of Public Administration Research and Theory* 28 (4): 506–518.

Rivera, Lauren A. 2012. "Hiring as Cultural Matching: The Case of Elite Professional Service Firms". *American Sociological Review* 77 (6): 999–1022.

Rolke, Olle and Johanna Rickne. 2025. "The class ceiling in politics". *American Political Science Review* 119 (3): 1375–1392.

Ruhil, Anirudh V. S. 2003. "Urban Armageddon or Politics as Usual? The Case of Municipal Civil Service Reform". *American Journal of Political Science* 47 (1): 159–170.

Ryan, Ann Marie, Joshua M. Sacco, Lynn A. McFarland, and S. David Kriska. 2000. "Applicant Self-Selection: Correlates of Withdrawal From a Multiple Hurdle Process". *Journal of Applied Psychology* 85 (5): 163–179.

Schneider, Benjamin. 1987. "The People Make the Place". *Personnel Psychology* 40 (3): 437–453.

Shefter, Martin. 1993. *Political Parties and the State.* Princeton, NJ: Princeton University Press.

Sidanius, Jim, James Liu, Felicia Pratto, and Julie Shaw. 1994. "Social Dominance Orientation, Hierarchy-Attenuators and Hierarchy-Enhancers: Social Dominance Theory and the Criminal Justice System". *Journal of Applied Social Psychology* 24 : 338–366.

Skowronek, Stephen. 1982. *Building a New American State: The Expansion of National Administrative Capacities, 1877–1920.* Cambridge University Press.

Socia, Kelly M., Melissa S. Morabito, Brenda J. Bond, and Elias S. Nader. 2021. "Public Perceptions of Police Agency Fairness and the Willingness to Call Police". *The American Review of Public Administration* 51 (5): 360–373.

US Census Bureau. 2023. "Annual Survey of State and Local Government Finances". https://www.census.gov/data/datasets/2023/econ/local/public-use-datasets.html.

US Census Bureau. 2025. "Annual Survey of Public Employment & Payroll". https://www.census.gov/data/datasets/2025/econ/apes/annual-apes.html.

Van Riper, Paul, Elder Gunter, Carl B. Barnes, Walter F. Kuiken, and Leroy J. Brenneman. 1963. *The Merit System: Foundation for Responsible Public Management.* Chicago: Chicago Public Personnel Association.

Weber, Max. 1946. "Bureaucracy". In H. Gerth and C. W. Mills (Eds.), *From Max Weber: Essays in Sociology*, pp. 196–244. New York: Oxford University Press.

Appendix: Supporting Information for *The Limits of Merit*

Table of Contents

A Merging Procedures and Quality

In this section, I describe the different merging procedures and the quality of the resulting matches.

A.1 Merge Payroll to L2

As described in the main text, I start with a roster of the six largest agencies in NYC. I identify individual employees by unique combinations of appointment date, last name, first name, and middle initial. I then use a probabilistic record linkage algorithm to match demographic information for all employees from the L2 voter file (Enamorado et al., 2019),

retaining all matches with a posterior match probability of at least 0.7. For first and last names, I allow for partial string distance matches using the Jaro-Winkler distance methods, while for the middle initial I enforce an exact match.

Since L2 does not include information on appointment dates and I need to rely on name information for these matches only, one might be concerned that this introduces a large amount of duplicates in the list of unique employees. However, note that the combination of first name, last name, and middle initial (i.e., excluding appointment date) already uniquely identifies 95% of all bureaucrats. Similarly, 87% of all voter names are unique in the L2 data. Importantly, these shares are very similar across party and race groups: 89% for Democrats, 95% for Republicans, 95% for Independents, 91% for Whites, 91% for Blacks, 85% for Asians, 74% for Hispanics and 92% for other races.

Table A1 shows that I am able to identify more than 80% of bureaucrats across all agencies, and Figure A1 indicates that the median posterior probability for these matches is more than 0.95. Additionally, Table A1 shows that the true match rate is estimated to be at least 69% across the agencies, and with low false discovery rates (FDR) and high false negative rates (FNR) the algorithm clearly errs on the side of not identifying a bureaucrat in the voter file rather than matching the wrong voter to a bureaucrat. More importantly, the FDR is relatively similar across race and party groups, albeit somewhat lower for Republicans, Whites and male bureaucrats across the agencies.[1] It's important to note that I weight by the posterior probability of matches in all analyses to quantify the uncertainty inherent in my merge procedures, and to calibrate and account for the amount of false positives and false negatives across demographic groups in my data.

Figure A1: Boxplots of Posterior Probabilities of Correct Matches across Matching Procedures

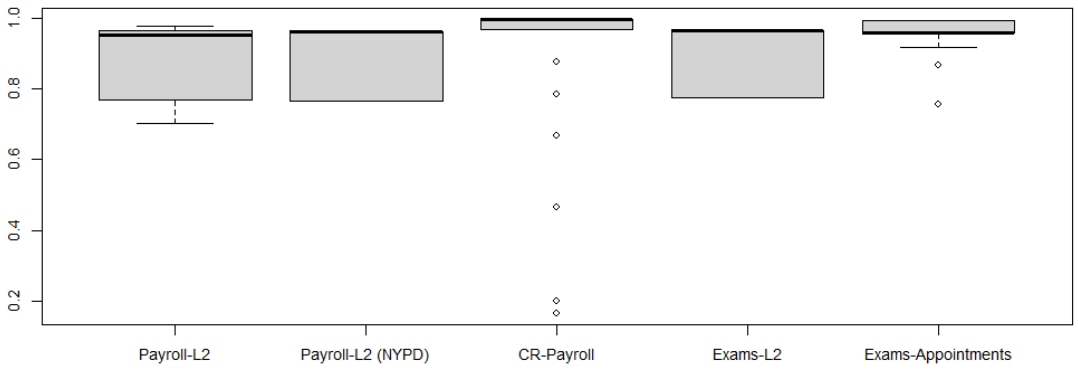

[1]Unfortunately, the FNR is computationally impossible to calculate by race or partisanship, since this would require retaining matches for *all* pairs across the two data sets, i.e. $N_{bureaucrats} \times N_{voters} \geq 168,577,197,264$.

Table A1: Merging of NYC Payroll (2014-2021) to 2021 L2 Voter File (N=7,940,144)

	NYPD (N=91,975)	Sanitation (N=32,468)	FDNY (N=28,016)	Social Services (N=22,386)	Corrections (N=21,231)	Education (N=189,772)
Number of matches	80,661 (88%)	26,618 (82%)	26,172 (93%)	18,909 (85%)	18,345 (87%)	85,798 (45%)
True match rate[a]	78%	73%	87%	69%	74%	
False negative rate (FNR)[b]	94%	95%	96%	92%	93%	
False discovery rate (FDR)[c]	11%	11%	7%	18%	15%	
FDR by groups						
Non-Partisan	11%	11%	7%	20%	15%	
Democrat	12%	12%	8%	17%	14%	
Republican	9%	11%	6%	21%	15%	
White	9%	10%	6%	18%	14%	
Hispanic	13%	13%	9%	20%	17%	
Black	11%	11%	8%	16%	13%	
Other Race	14%	14%	10%	23%	16%	
Asian	15%	16%	11%	22%	18%	
Female	12%	12%	9%	18%	14%	
Male	10%	11%	7%	18%	15%	

[a] Share of matches weighted by their posterior probability of a match; [b] Probability of wrongfully declaring a non-match given the threshold; [c] Probability of wrongfully declaring a match given the threshold

A.2 One-to-Many Matches

An issue when matching administrative records that lack a large number of unique identifiers and matching variables is the possibility of one-to-many matches, i.e., several agency employees match to more than one voter. As noted early, names in both the bureaucrat list and voter files are predominantly unique. As a result, Table A2 shows that the algorithm achieves a very high rate of one-to-one matches, with the share of unique matches being above 90% across agencies (with the exception of the Department of Education). Further inspecting the remaining one-to-many matches suggests that these are due to a re-appointment of employees (e.g., a traffic enforcement agent becomes appointed as a police officer). Since these types of duplicates are warranted, I maintain one-to-many matches without further adjustments.

Table A2: Share of One-to-One Matches: Payroll-L2

	NYPD	Sanitation	FDNY	Social Services	Corrections	Education
Number of matches	80,661 (88%)	26,618 (82%)	26,172 (93%)	18,909 (85%)	18,345 (87%)	85,798 (45%)
Share of one-to-one matches	94.5%	92.8%	92%	96.9%	97.8%	80.1%

Similar levels and reasoning applies to one-to-many matches for the additional merges of NYPD datasets (Table A3). Multiple city records can match the same employee on the payroll, e.g., if the same employee receives multiple promotions or retires for different positions. Similarly, the same voter may take the entry or promotion exam several times, thus leading to a small share of duplicates in the exam-L2 merges. Exceptions are the exam-appointment and exam-promotion match, where the same exam-taker should only be hired/promoted for the same position once. I account for these one-to-many matches in my matching procedures, as outlined in sections A.5 and A.6.

Table A3: Share of One-to-One Matches: Other NYPD matches

	City Record-Payroll	Entry Exams-L2	Entry Exams-Appointments	Promo Exams-L2	Promo Exams-Promotions
Share of one-to-one matches	90.5%	88.9%	64.8%	85%	98%

While these figures alleviate concerns about duplicates in voter names, one may still be concerned that matched names differ in their homogeneity of party and race. That is, if, for example, names in the bureaucrat list matched to Republicans have a larger share of Democrats with the identical names in the voter files than Democratic bureaucrat names have Republicans in the voter file, one-to-one matches could lead to bias and an overestimation of Republicans in the police. To assess this, I estimate group homogeneity in the voter file for all names that I classify as matches in the bureaucrat list. As Table A4 shows, differential name purity is of little concern. If anything, Republican and White names in the officer list have *higher* name purity than Democratic and Black names. Hence, it is unlikely that I artificially inflate the number of Republicans and Whites in the matched set due to mismeasured partisanship or race among their names.

Table A4: Name Purity of Matches: Payroll-L2

	Republican	Democrat	Non-Partisan	Asian	Black	Hispanic	Other Race	White
Purity by Full Names	93%	92%	90%	99.9%	96.3%	100%	99.9%	98.8%

Purity measures are maximum proportions of (partisan/race) groups by names, averaged by matched party or race.

A.3 Merge City Records with NYPD Payroll

For my analysis of the selection dynamics at the NYPD, I collected various additional data sets. Particularly, I obtained information on (1) career milestones, including appointments, promotions, and retirements from official records published in the daily City Record newspaper since 2014[2]; (2) unit assignment, awards, and arrest records for a cross-section of about 33,000 active officers (as of October 2021) from NYPD's official officer profiles[3]; (3) civil service exams, both for entry and promotions.[4] I then match these different data sets to the roster of NYPD bureaucrats. The following sections describe these various matches in more detail. Table A5 illustrates the number of successful matches across data sets, and Figures A2 and A3 illustrate the high posterior probability of a match across the procedures together with slight differences across demographic groups.

I link the 65,856 City newspaper records on appointments, promotions, demotions, retirements, resignations, dismissals, and terminations to the roster of NYPD bureaucrats in the following way: To maximize overlap, I restrict the city records to those with effective date between January 1, 2014, and July 1, 2021. I then match on first name, last name, and middle initial, and retain matches with a minimum posterior of 0.7. I do not deduplicate

[2] https://www.nyc.gov/site/dcas/about/cityrecord-editions.page

[3] https://nypdonline.org/link/2

[4] I carefully combined both active and terminated civil service lists as of July 2022 from https://opendata.cityofnewyork.us/.

matches returned by the algorithm. Instead, I deal with multiple city records that match to the same payroll names in the following way:

- In cases where there is one maximum posterior, we keep the maximum posterior match (equivalent what fastLink enforces as default)
- For cases where there is no unique maximum posterior probability, I use additional information in the data.
- For *appointments*, I identify exact matches using the appointment date provided in both data sets. If there is more than one exact match, I use the maximum posterior probability for these duplicates again and if these maxima are not unique, I retain the payroll entry with the earliest appointment. If there are no exact matches using the appointment date, I retain those with the smallest difference between the appointment date in the payroll and the effective date in the city records. If there is more than one such time distance match, I use the maximum posterior match probability for these duplicates again and if these maximum posteriors are not unique, I retain the payroll entry with the earliest appointment.
- For *attrition*, I again follow the default in fastLink and use the maximum posterior to find the best match in cases of multiple matches based on name. If there are more than one such maximum posterior match, I use the smallest difference between the fiscal year at the date of attrition and the last fiscal year recorded in the payroll. If this still does not give a unique match, I use the employee with the earliest appointment date, i.e. the person who has been on the force the longest before retirement.
- For *promotions*, I use information on changes in job titles on the payroll to adjudicate between multiple maximum-posterior matches. I first retain the maximum posterior among all promotion matches. If there is more than one possible match, I use the year of promotion that matches the promotion in the city records. If this is not the case, I use the minimum difference in the year of the title change on the payroll and the year of promotion in the city records. For the remaining duplicates, I use the employee who was appointed first. For the remaining 4 duplicates, I finally use the first observation in the data.

Overall, I am able to successfully match 65,184 (99%) of 65,856 city records.

A.4 Merge NYPD Payroll to NYPD Online Officer Profiles

I used another iterative process to match the roster of active officers from NYPD online to my payroll roster. I first merge these records based on exact matches using the appointment date, the first name, the last name, and the middle initial. All payroll records I fail to match exactly, I then also match based on a fuzzy name match and extensive manual checks by research assistants.

A.5 Merge Entry Exams to Appointments and L2 Voter File

To match entry-level exams to the roster of NYPD personnel, I collect the information on 96,883 NYC civil service exams for police officer and traffic enforcement agents. I rely on both the active and terminated civil service lists as of July 2022.[5] Following a probabilistic record linkage based on the first name, last name and middle initial, I first retain all matches with a minimum posterior probability of 0.5. I then rely on the following procedure to deduplicate matches of the same exam: First, I restrict matches by enforcing that the police officer title of the exam corresponds to the title of the employee's appointment. I further deduplicate matches by requiring that the date of the exam needs to be before the date of appointment. I further filter duplicates by retaining the maximum posterior probability match by exam. For the remaining duplicates, I use the minimum distance between the appointment date and the exam date. I also account for duplicates in the appointments for the same exam-taker, e.g., because some hired NYPD employees took the exam more than once before getting hired. I again rely on the maximum posterior match and the exam date relative to the appointment date to identify the successful exam for these individuals. Finally, I also match the entry exams to the L2 voter file. For computational reasons, I first estimate the EM step of the fastLink algorithm on a random 50% sample of the exam takers and voters, respectively. I then apply the parameter estimates to all exam-L2 combinations, and block on the first letter of the last name for computational tractability. Also, I restrict the eligible voters to those younger than 45 (in 2020) to account for the age restrictions in the NYPD hiring process.

I am able to match 10% of all exam takers to NYPD appointments and identify 82% of all appointments between 2014 and 2021 among exam takers. The non-matched appointments are mostly from earlier years (e.g., 63% appointments in 2014 remain unmatched to exams) due to truncation and delayed temporal overlap between exams and appointments.

A.6 Merge Promotion Exams to Awarded Promotions and L2 Voter File

Individuals who seek promotions to the ranks of Sergeant, Lieutenant, and Captain need to successfully pass additional civil service exams at the NYPD. I match promotion exams to the promotions recorded in the city records probabilistically based on employee's first name, last name and middle initial. After retaining matches with a minimum posterior match probability of 0.5, I again ensure that exam titles correspond to the titles of the matched promotion records. Similar to the entry exams, I also require that the date of the promotion exam is before the date of the promotion, and use the minimum time difference between the exam date and the promotion date for the remaining duplicates. To ensure that I record only one promotion for a successful exam-taker (i.e., rule out duplicates in promotions for the same exam takers), I follow a similar procedure to the entry exams: I again rely on the maximum posterior match and the exam date relative to the appointment date to identify the successful exam for these individuals. Finally, I also match the promotion exams to the L2 voter file, following exactly the same procedure as for the payroll-L2 merging described

[5] https://data.cityofnewyork.us/City-Government/Civil-Service-List-Active-/vx8i-nprf/about_data; https://data.cityofnewyork.us/City-Government/Civil-Service-List-Terminated-/qu8g-sxqf/about_data

above.

Table A5: Merging of NYPD Data Sets

	L2 Voter File, 2020 (N=7,940,144)	City Records, 2014-2021 (N=65,856)	Active Officer Profiles (N=33,072)	Other Civil Service Exams (N=195,308)
NYPD Payroll*, FY2014-2021 (N=91,975)	80,661 (88%, .)	65,184 (., 99%)	32,632 (., 99%)	
Officer Entry Exams, 2014-2021 (N=99,159)	78,894 (80%, .)	11,763 (12%, 81%[a])		19,651 (20%,.)
Promotion Exams, 2012-2021 (N=5,725)	5,226 (91%,.)	3,501 (61%,.)		

Percentages of matched observations in parentheses (row percentages, colum percentages); * includes both uniformed and civilian personnel; [a] among NYPD POLICE OFFICER and POLICE CADET appointments only (N=14,528)

Figure A2: Average Posterior Probabilities of Correct Matches By Groups across Matching Procedures - Payroll Data

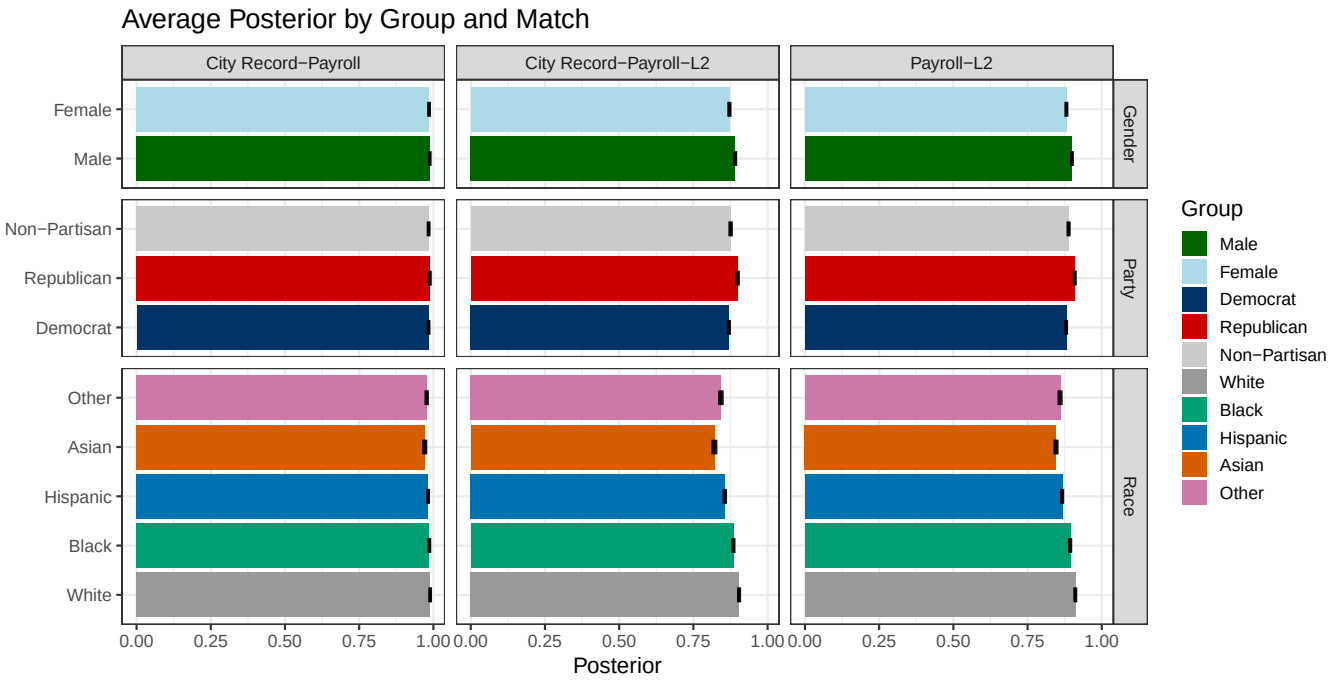

Figure A3: Average Posterior Probabilities of Correct Matches By Groups across Matching Procedures - Exam Data

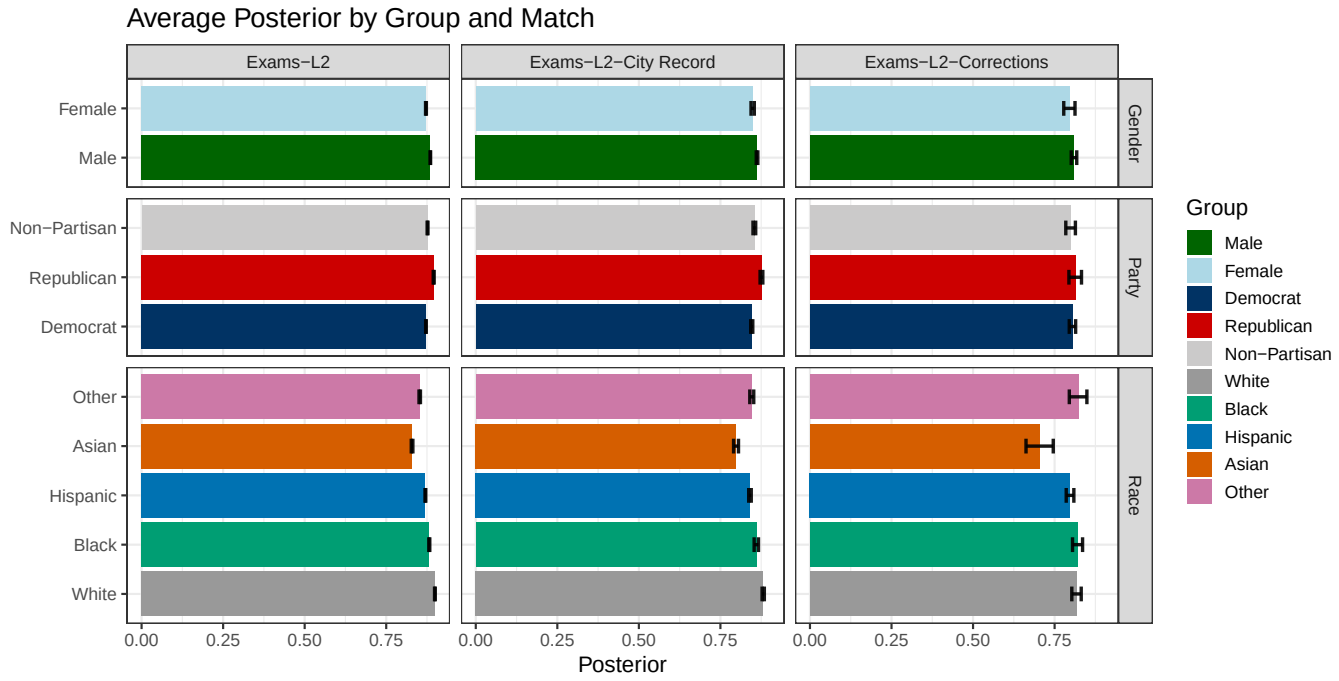

A.7 Measurement Error in L2 Race/Ethnicity

I assess the validity of the L2 race coding in two ways.

First, the LEMAS survey asks police agencies to report the number of sworn full time employees across different demographic groups (particularly, race and gender). I draw on the answers of the NYPD in the 2020 LEMAS survey to assess the validity of the L2 information on race and gender. Figure A4 depicts the proportion of officers in each racial/ethnic category and by gender as measured by L2 vs. LEMAS. To ensure optimal overlap in the populations used in both data sets, I restrict the payroll to only sworn personnel and weight estimates by the posterior probability of a match. As the figure shows, L2 slightly underrepresents the proportion of White officers, and more so, the share of other minorities. This discrepancy largely stems from the overrepresentation of the "other/unknown" category, which makes up 5% in the L2 data, but essentially 0% in the LEMAS data. The L2 data also slightly overrepresents female officers compared to the LEMAS data.

Figure A4: Comparison of 2020 LEMAS and L2 Measures of Officer Race

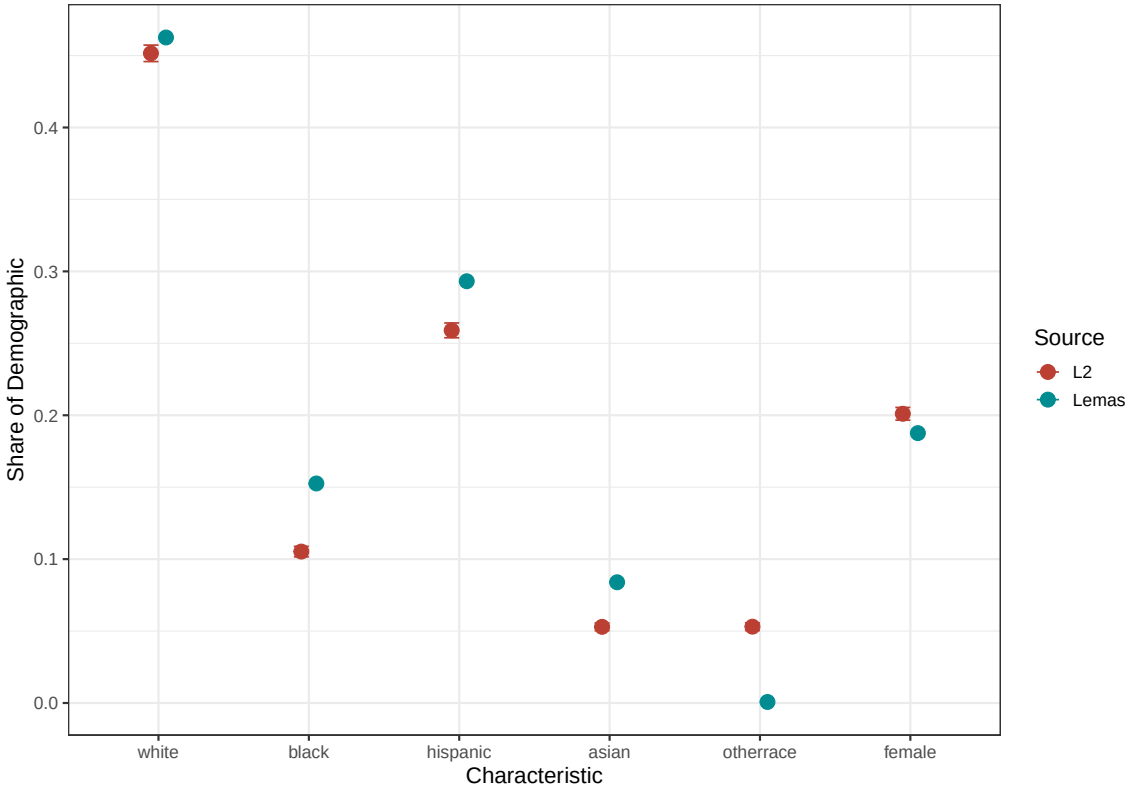

Second, I rely on predictions of individual-level ethnicity of NYC employees using Imai and Khanna (2016). Figure A5 compares the estimated share of non-White employees across agencies by types of measures. Evidently, the share of non-White bureaucrats is estimated to be substantially higher with L2's race measures than alternative predictions across agencies. Taken together, these tests suggest that I slightly *overestimate* the degree of representation for minority groups in NYC's bureaucracy, i.e., the representational gaps may be even starker than I estimate.

Figure A5: Comparison of Bayesian Name Predictions and L2 Measures of Bureaucrat Race

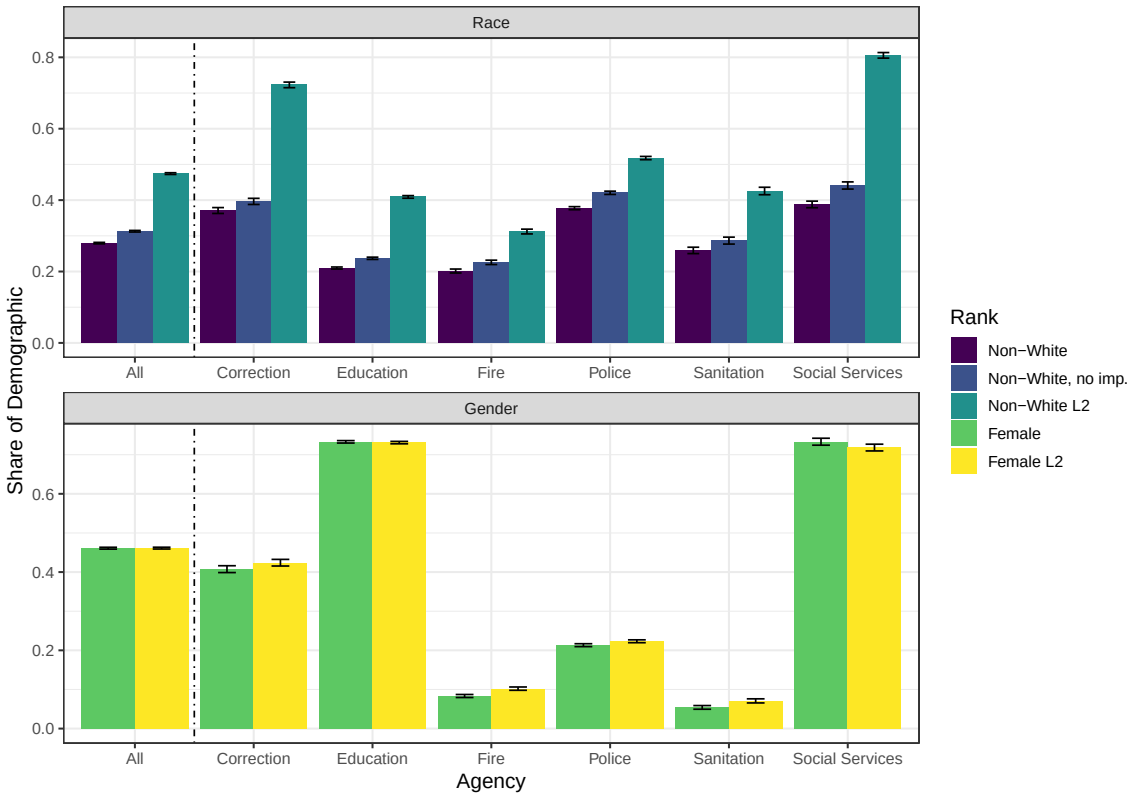

"Non-White" and "Non-White, no imp." use Bayesian surname-based race predictions from the `wru` R package; "no imp." restricts to individuals whose surname matched directly in the Census surname list. "Non-White L2" uses L2's proprietary race imputation based on name and census-block demographics.

B Additional Tables

B.1 Entry Exams

Table A6: Differences in Exam Scores across Demographics

	Model 1	Model 2	Model 3	Model 4
Republican	0.04	−0.01	−0.05	0.36***
	(0.08)	(0.07)	(0.07)	(0.07)
Non-Partisan	0.00	−0.07	−0.06	0.22***
	(0.07)	(0.06)	(0.06)	(0.06)
Black	−1.70***	−1.78***	−1.77***	−2.83***
	(0.09)	(0.09)	(0.09)	(0.09)
Hispanic	−1.15***	−1.28***	−1.25***	−2.28***
	(0.07)	(0.07)	(0.07)	(0.07)
Asian	0.26*	0.15	0.20	−1.19***
	(0.12)	(0.12)	(0.12)	(0.12)
Other Race	−1.83***	−1.75***	−1.71***	−3.02***
	(0.11)	(0.11)	(0.11)	(0.11)
Veteran Credit		6.19***	6.21***	6.27***
		(0.12)	(0.12)	(0.12)
Parent/Sibling Legacy Credit			9.47***	9.56***
			(0.41)	(0.40)
Residency Credit				3.76***
				(0.06)
Exam FE	Yes	Yes	Yes	Yes
Mean of DV	88.82	88.82	88.82	88.82
Adj. R^2	0.03	0.07	0.08	0.13
Num. obs.	69406	69406	69406	69406

Regressions weighted by posterior probability of a match between exam data and voter file, and between the exams and appointments. Dependent variable: Exam score in entry-level exam (police officer exams only). Veteran credits are awarded for veterans and disabled veterans. Parent and sibling legacy credits are additional credits on the exam score to candidates who lost a parent/sibling as a result of 911. Residency credits are given on exams to candidates who maintain a continuous period of residency in NYC. HC1 standard errors in parentheses. $^{***}p < 0.001$; $^{**}p < 0.01$; $^{*}p < 0.05$.

Table A7: Removing Party Switchers For Exams 2016-2020

	Model 1	Model 2	Model 3	Model 4
Republican	0.06***	0.06***	0.05***	0.04***
	(0.01)	(0.01)	(0.01)	(0.01)
Non-Partisan	0.02***	0.02***	0.02**	0.01**
	(0.00)	(0.00)	(0.00)	(0.00)
Black	−0.03***	−0.03***	−0.01	−0.01*
	(0.01)	(0.01)	(0.01)	(0.01)
Hispanic	−0.00	−0.00	0.01	0.01
	(0.00)	(0.00)	(0.01)	(0.01)
Asian	0.00	−0.00	0.01	0.01
	(0.01)	(0.01)	(0.01)	(0.01)
OtherRace	−0.01	−0.01	0.00	−0.00
	(0.01)	(0.01)	(0.01)	(0.01)
Exam*Score Bin FE	Yes	Yes	Yes	Yes
Additional controls	No	No	Yes	Yes
Drop Party Switchers	No	Yes	No	Yes
Mean of DV	0.13	0.13	0.13	0.13
Adj. R^2	0.03	0.03	0.04	0.04
Num. obs.	31035	30106	28608	27755

$^{***}p < 0.001$; $^{**}p < 0.01$; $^{*}p < 0.05$

Table A8: Removing Party Switchers For Exams 2016-2020 and Restricting to Voters Registered Before 2018

	Model 1	Model 2	Model 3	Model 4
Republican	0.06^{***}	0.05^{***}	0.04^{***}	0.03^{***}
	(0.01)	(0.01)	(0.01)	(0.01)
Non-Partisan	0.02^{***}	0.02^{**}	0.01	0.01
	(0.01)	(0.01)	(0.01)	(0.01)
Black	-0.03^{***}	-0.04^{***}	-0.01	-0.01^{*}
	(0.01)	(0.01)	(0.01)	(0.01)
Hispanic	-0.00	-0.01	0.01	0.01
	(0.01)	(0.01)	(0.01)	(0.01)
Asian	-0.00	-0.00	0.00	-0.00
	(0.01)	(0.01)	(0.01)	(0.01)
OtherRace	-0.01	-0.02^{*}	-0.00	-0.01
	(0.01)	(0.01)	(0.01)	(0.01)
Exam*Score Bin FE	Yes	Yes	Yes	Yes
Additional controls	No	No	Yes	Yes
Drop Party Switchers	No	Yes	No	Yes
Mean of DV	0.13	0.12	0.13	0.13
Adj. R^2	0.03	0.03	0.04	0.04
Num. obs.	23196	22451	21460	20768

$^{***}p < 0.001$; $^{**}p < 0.01$; $^{*}p < 0.05$

B.2 Career Advancement and Promotions

Table A9: Difference in Probability of Taking a Promotion Exam (2014-2021)

	Model 1	Model 2	Model 3
Republican	−0.00		0.00
	(0.00)		(0.00)
Non-Partisan	0.00		−0.00
	(0.00)		(0.00)
Black		−0.00	−0.00
		(0.00)	(0.00)
Hispanic		0.00	0.00
		(0.00)	(0.00)
Asian		0.05***	0.05***
		(0.01)	(0.01)
Other Race		0.04***	0.04***
		(0.01)	(0.01)
Cohort FE	Yes	Yes	Yes
Mean of DV	0.09	0.09	0.09
Adj. R^2	0.12	0.12	0.12
Num. obs.	49558	48521	45990

$^{***}p < 0.001$; $^{**}p < 0.01$; $^{*}p < 0.05$

Table A10: Difference in Probability of Receiving Promotion After Promotion Exam (2014-2021)

	Model 1	Model 2	Model 3	Model 4	Model 5	Model 6
Republican	0.03	0.01			0.01	0.01
	(0.01)	(0.01)			(0.01)	(0.01)
Non-Partisan	0.04*	0.01			0.01	0.01
	(0.02)	(0.01)			(0.01)	(0.01)
Black			−0.05*	0.00	0.01	0.01
			(0.02)	(0.02)	(0.02)	(0.02)
Hispanic			−0.02	0.00	0.01	0.01
			(0.02)	(0.01)	(0.01)	(0.01)
Asian			−0.02	−0.03	−0.02	−0.02
			(0.03)	(0.02)	(0.02)	(0.02)
Other Race			0.03	0.01	0.02	0.02
			(0.03)	(0.02)	(0.02)	(0.02)
Examscore (80-90)		0.57***		0.57***	0.57***	
		(0.01)		(0.01)	(0.01)	
Examscore (90-100)		0.73***		0.73***	0.73***	
		(0.01)		(0.02)	(0.02)	
Exam FE	Yes	Yes	Yes	Yes	Yes	Yes
Exam × Score Bin FE	No	No	No	No	No	Yes
Mean of DV	0.61	0.61	0.60	0.60	0.61	0.61
Adj. R^2	0.14	0.51	0.14	0.51	0.51	0.51
Num. obs.	4963	4963	4790	4790	4549	4549

Linear probability regressions, weighted by the posterior probability of a match between promotion exams and voter file. Level of observation: Promotion exam taker. Outcome: Dummy for whether the exam taker actually received a promotion between 2014 and 2021. HC1 standard errors in parentheses. $^{***}p < 0.001$; $^{**}p < 0.01$; $^{*}p < 0.05$

Table A11: Differences in (Log) Number of Awards by Officer Characteristics

	Model 1	Model 2	Model 3
Republican	0.20***		0.13***
	(0.01)		(0.01)
Non-Partisan	0.08***		0.04**
	(0.01)		(0.01)
Black		−0.27***	−0.21***
		(0.02)	(0.02)
Hispanic		−0.16***	−0.12***
		(0.01)	(0.01)
Asian		−0.26***	−0.24***
		(0.02)	(0.02)
Other Race		−0.14***	−0.12***
		(0.02)	(0.02)
Cohort FE	Yes	Yes	Yes
Mean of DV	0.92	0.91	0.91
Adj. R^2	0.21	0.21	0.22
Num. obs.	28422	27609	26222

Regressions weighted by posterior probability of a match between payroll and voter file. The information on awards received is only available for active officers (here as of October 2021). Level of observation: Employee. Outcome: Log number of departmental awards since their appointment date at NYPD. $^{***}p < 0.001$; $^{**}p < 0.01$; $^{*}p < 0.05$

Table A12: Differences in Command Assignments for Active Officers (as of 10/20/2021)

	All Elite	Terrorism	Drugs	Special Forces
Republican	0.03^{***}	0.00	0.00	0.02^{***}
	(0.00)	(0.00)	(0.00)	(0.00)
Non-Partisan	0.02^{***}	0.00	0.00	0.01^{***}
	(0.00)	(0.00)	(0.00)	(0.00)
Black	-0.03^{***}	-0.01^{**}	-0.00	-0.01^{**}
	(0.01)	(0.00)	(0.00)	(0.00)
Hispanic	-0.01	-0.00	0.00	-0.00
	(0.00)	(0.00)	(0.00)	(0.00)
Asian	-0.03^{***}	-0.00	-0.01^{***}	-0.01^{**}
	(0.01)	(0.01)	(0.00)	(0.00)
Other Race	-0.02^{**}	-0.00	-0.00	-0.02^{***}
	(0.01)	(0.01)	(0.00)	(0.00)
Cohort FE	Yes	Yes	Yes	Yes
Mean of DV	0.10	0.04	0.02	0.04
Adj. R^2	0.05	0.02	0.01	0.02
Num. obs.	26222	26222	26222	26222

Regressions weighted by posterior probability of a match between payroll data and voter file. The command information is only available for active officers (here as of October 2021). Level of observation: Employee. HC1 standard errors in parentheses. $^{***}p < 0.001$; $^{**}p < 0.01$; $^{*}p < 0.05$

Table A13: Correlation of Team Leadership and Team Composition

	Share of Party			Share of Race				
	Republican	Democrat	Non-Partisan	Hispanic	White	Asian	Black	Other
Republican leader	0.05*	−0.04*	−0.01	−0.03	0.02	−0.01	0.01	0.01
	(0.02)	(0.02)	(0.01)	(0.03)	(0.03)	(0.01)	(0.02)	(0.01)
Non-Partisan leader	−0.02	−0.03	0.05**	−0.01	0.01	−0.01	0.00	−0.00
	(0.02)	(0.02)	(0.02)	(0.02)	(0.03)	(0.01)	(0.02)	(0.01)
Asian leader	0.04	−0.06	0.02	−0.02	−0.07	0.08***	−0.02	0.04*
	(0.04)	(0.04)	(0.03)	(0.04)	(0.04)	(0.01)	(0.02)	(0.02)
Black leader	−0.02	0.04	−0.02	−0.01	−0.13***	0.03*	0.08***	0.03**
	(0.03)	(0.03)	(0.02)	(0.03)	(0.03)	(0.01)	(0.02)	(0.01)
Hispanic leader	−0.03	0.02	0.00	0.06**	−0.09***	0.01	0.01	0.01
	(0.02)	(0.02)	(0.02)	(0.02)	(0.03)	(0.01)	(0.01)	(0.01)
Other Race leader	−0.03	0.05	−0.02	0.02	−0.10	0.02	−0.01	0.07***
	(0.05)	(0.05)	(0.03)	(0.04)	(0.05)	(0.02)	(0.03)	(0.02)
Mean DV	0.38	0.39	0.23	0.27	0.52	0.05	0.11	0.05
R^2	0.05	0.03	0.03	0.06	0.11	0.10	0.05	0.05
Adj. R^2	0.03	0.01	0.01	0.04	0.09	0.08	0.03	0.03
Num. obs.	644	644	644	644	644	644	644	644

Cross-sectional OLS; The information on team assignment is only available for active officers (here as of October 2021). Level of observation: Team. Outcome: Share of relevant demographic per team. Regressions also control for 5 bins of team size, precinct team dummy, special operations team dummy and the highest rank of the leader. $^{***}p < 0.001$; $^{**}p < 0.01$; $^{*}p < 0.05$

Table A14: Correlation of Team Diversity and Team Leadership

	Gini Coefficient	
	Party	Race
Republican	0.03	0.03
	(0.02)	(0.02)
Non-Partisan	−0.01	0.03
	(0.02)	(0.02)
Asian	−0.03	−0.11***
	(0.04)	(0.03)
Black	0.06*	−0.05*
	(0.03)	(0.02)
Hispanic	0.03	−0.01
	(0.02)	(0.02)
Other Race	0.01	−0.10**
	(0.05)	(0.03)
Adj. R^2	0.13	0.22
Num. obs.	644	644

$^{***}p < 0.001$; $^{**}p < 0.01$; $^{*}p < 0.05$. Regressions also control for 5 bins of team size, precinct team dummy, special operations team dummy and the highest rank of the leader. Level of observation: Team. Dependent variable: Gini index.

B.3 Exits

Table A15: Differences in Years on Force at Retirement, by Characteristics and Rank

	Model 1	Model 2	Model 3	Model 4
Republican	0.52^{***}		-0.04	-0.06
	(0.13)		(0.15)	(0.15)
White		1.48^{***}	1.50^{***}	1.48^{***}
		(0.13)	(0.14)	(0.14)
Age at appointment	-0.12^{***}	-0.11^{***}	-0.11^{***}	-0.11^{***}
	(0.01)	(0.01)	(0.01)	(0.01)
CCRB complaints				0.08^{***}
				(0.01)
Mean of DV	22.41	22.40	22.40	22.40
Adj. R^2	0.05	0.07	0.07	0.07
Num. obs.	6624	6613	6613	6613

OLS, weighted by probability of matches between payroll and voter file and payroll and retirement records. Level of observation: Retiree. Outcome: Time since appointment date at retirement. HC1 standard errors in parentheses. $^{***}p < 0.001$; $^{**}p < 0.01$; $^{*}p < 0.05$

C Additional Figures

C.1 Degree of Representativeness

Figure A6: PUMA-level (Un)representativeness of State and Local Government Workforces

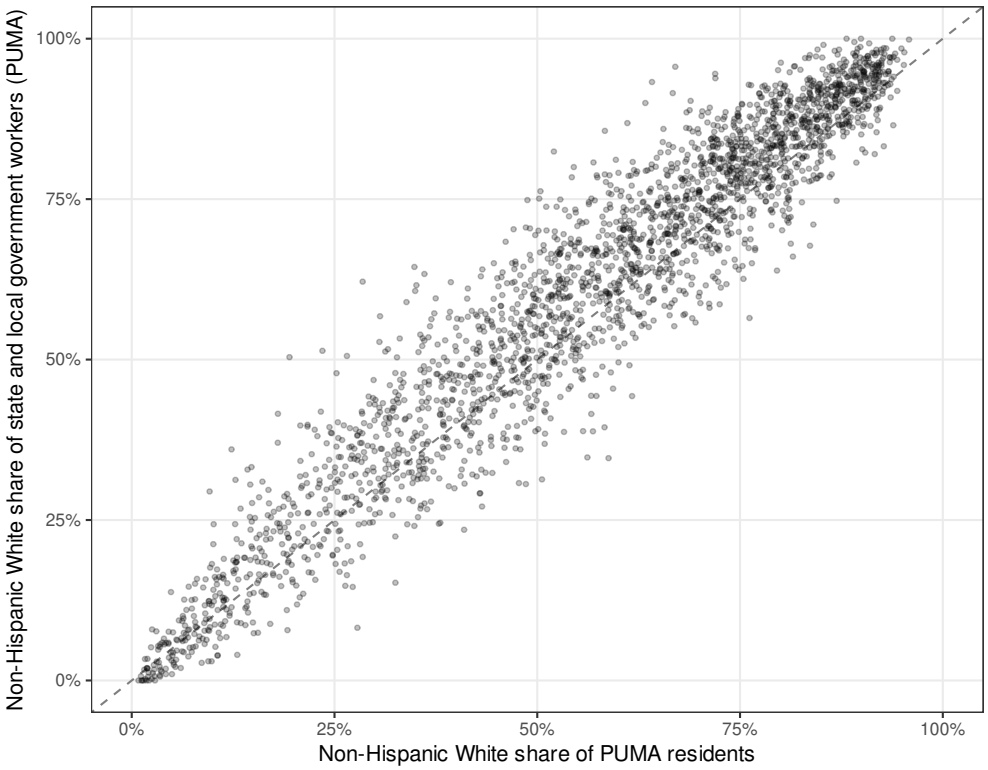

Each point is one of 2,462 Public Use Microdata Areas (PUMAs) in the 2023 1-year ACS, units of approximately 100,000 residents. The vertical axis is the non-Hispanic White share of state and local government workers living in the PUMA; the horizontal axis is the non-Hispanic White share of all PUMA residents. Points above the 45-degree line indicate PUMAs in which the public workforce is Whiter than the resident population. Source: 2023 ACS 1-year Public Use Microdata Sample. The median PUMA-level gap is +3.9 percentage points.

C.2 Entry

Figure A7: Correlation between Precinct Crime Rates and Representation Gaps

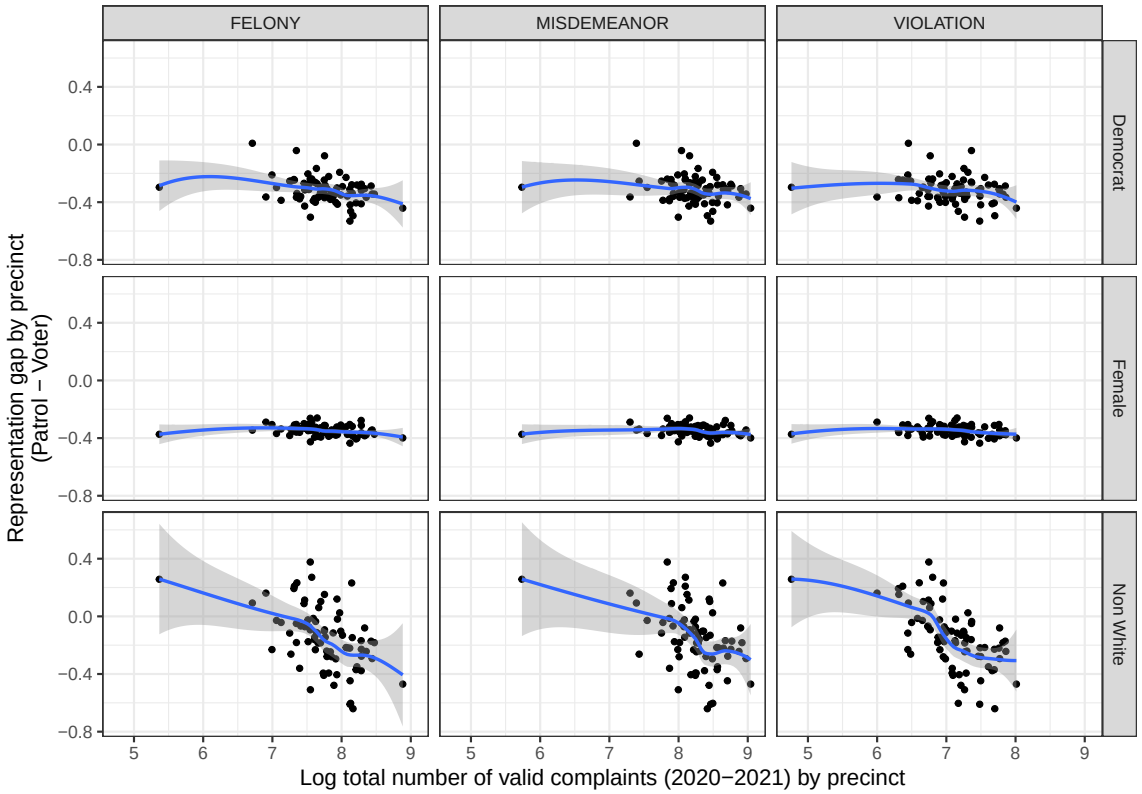

The x-axis shows the log total number of valid complaints by precinct between 01/01/2020 and 30/06/2021. The y-axis is the representation gap by precinct (share of voter - share of officers), where lower values indicate an underrepresentation of the characteristics. The officer data only includes officers that are identifiable in my snapshot of the NYPD online profiles (see Appendix A.4).

Figure A8: Share of Demographics, Police vs. Registered Voters in NYC

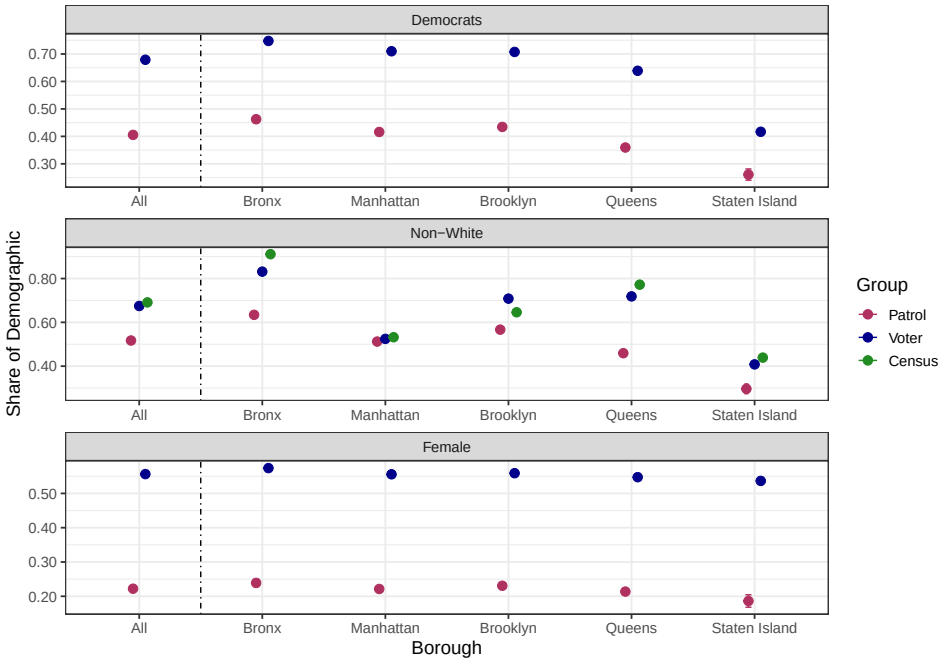

Uniformed police employees and traffic enforcement agents are assigned to boroughs based on their work location in their final year on the payroll. Agency estimates are weighted by the posterior probability of matches between agency payrolls and the voter file.

Figure A9: Hiring Differences Across Exam Scores

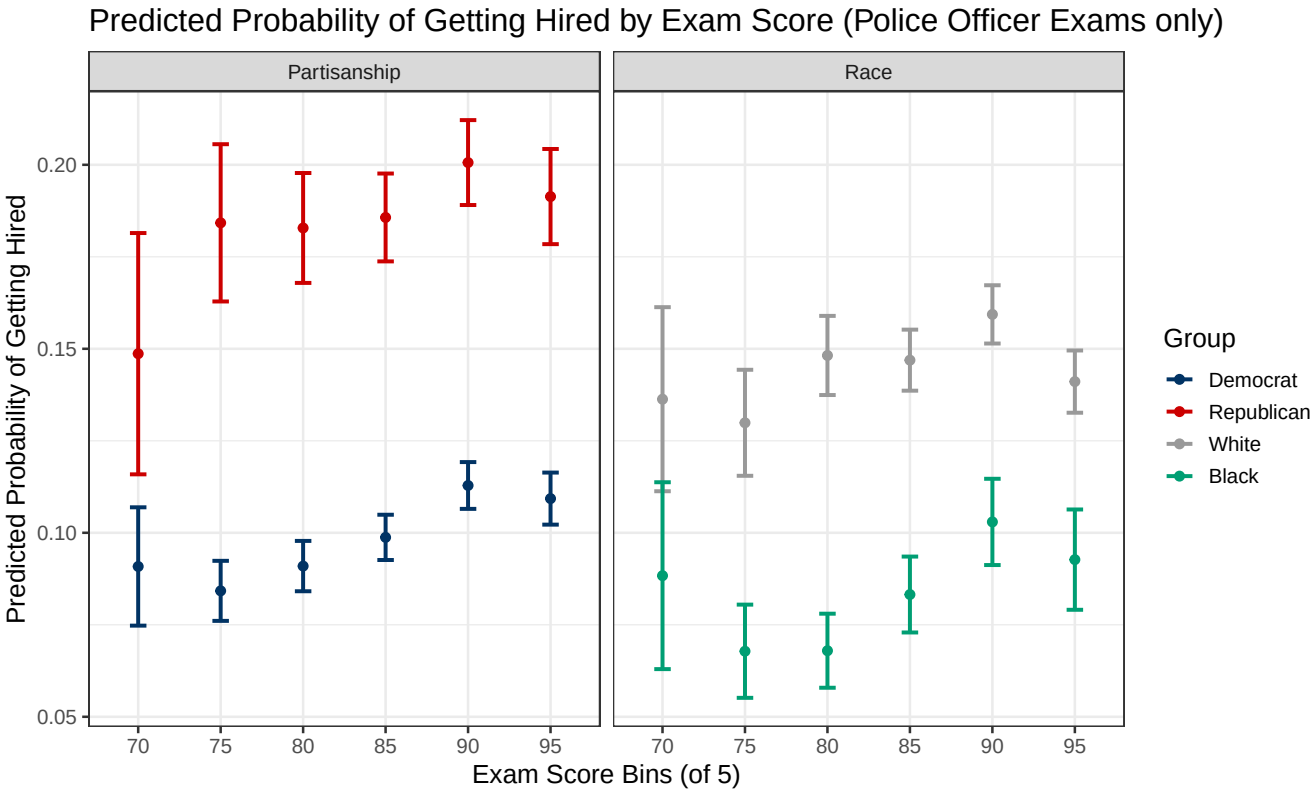

Figure A10: Score Differences by Hiring Groups

Figure A11: Length of Hiring Process

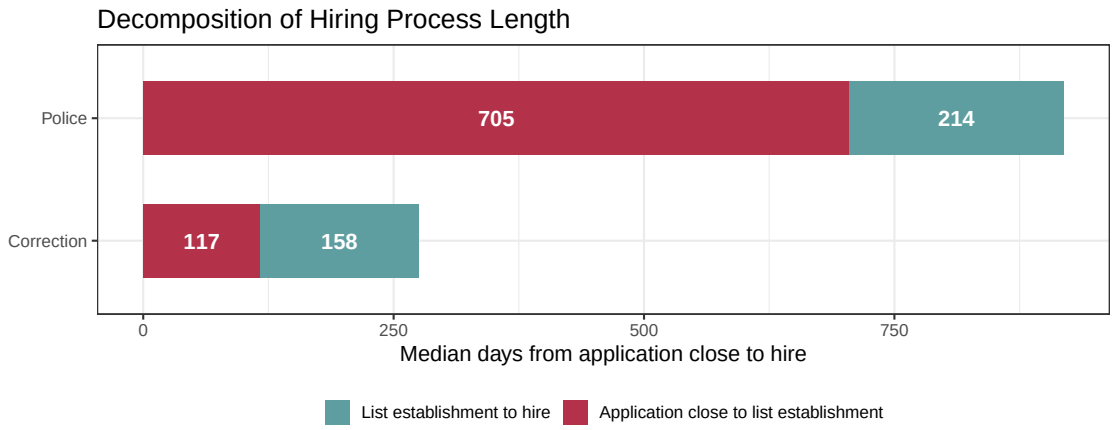

Note: Median number of days, decomposed into two stages. Application close to list establishment captures the period before NYPD/DOC requests certification of an eligible list from DCAS; under NYC civil service rules, certifications are initiated by the appointing agency to fill specific vacancies, so this stage mostly reflects each agency's pace of vacancy-driven requests rather than DCAS processing time. List establishment to hire captures the post-certification screening period administered by the agency itself (psychological, medical, character investigation, etc.). Sample restricted to candidates with non-missing application-period dates from the FOIA-released DCAS schedules and the DCAS open-data Annual Examination Schedule (police: 7,349 hires; corrections: 4,855 appointments).

Figure A12: Disqualification Appeals by Agency

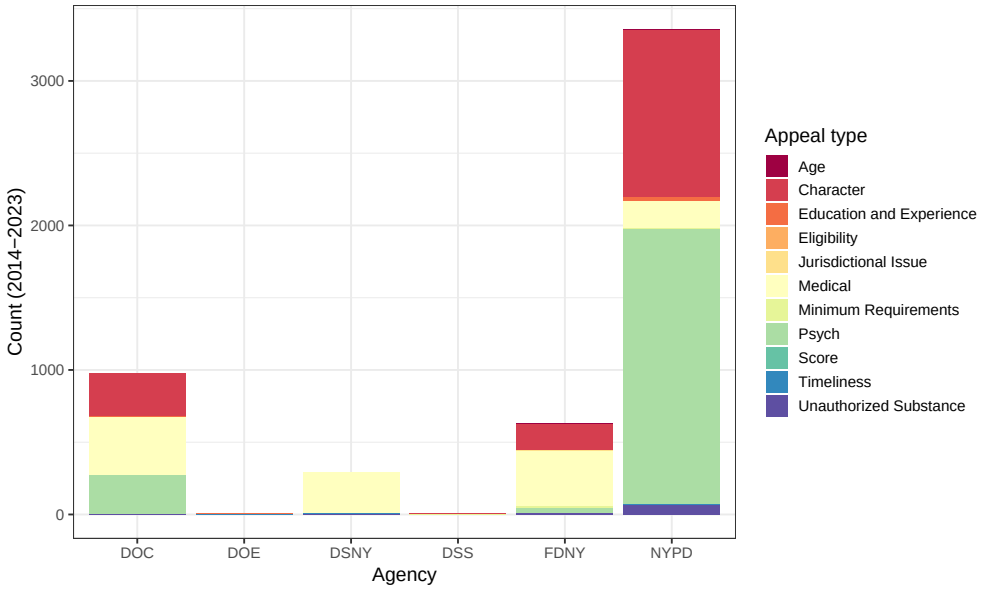

Figure A13: Income Distributions by Exam Taker Type

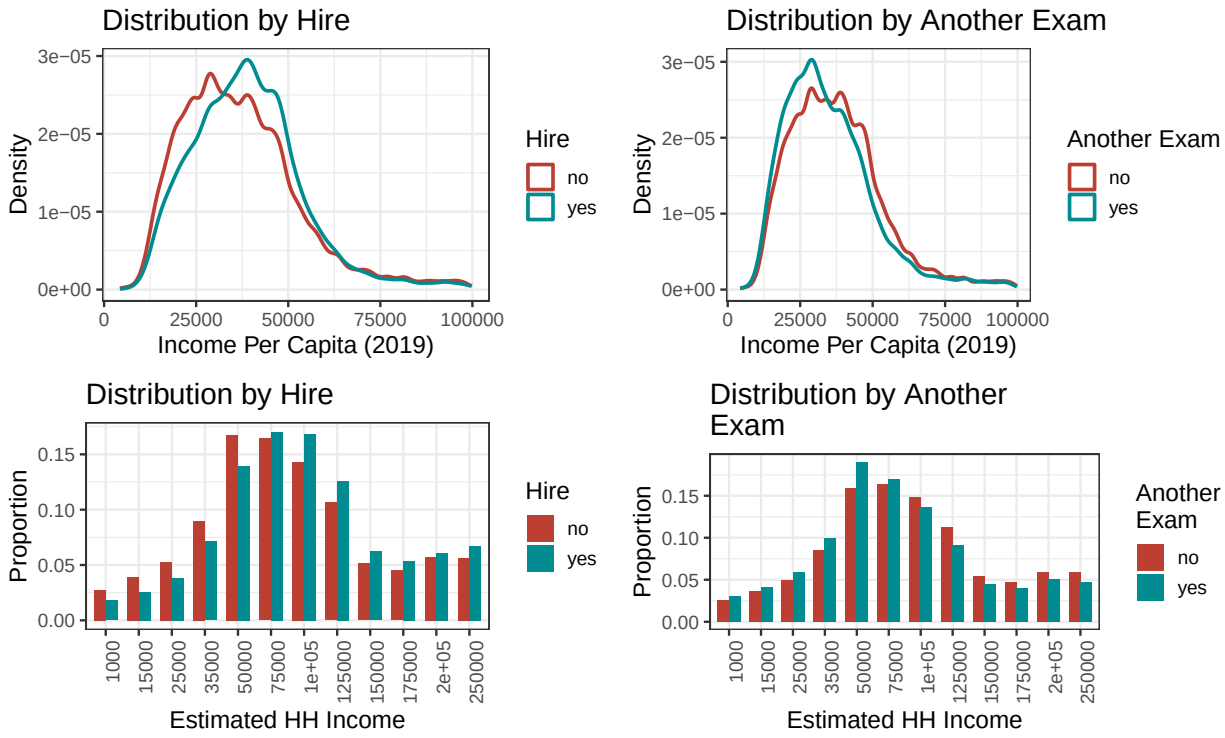

Note: The top panels use income per capita from the 2019 American Community Survey (matched by census tract), and the bottom panels use L2's estimated household income.

Figure A14: Predicted Household Income by Hiring Status and Groups

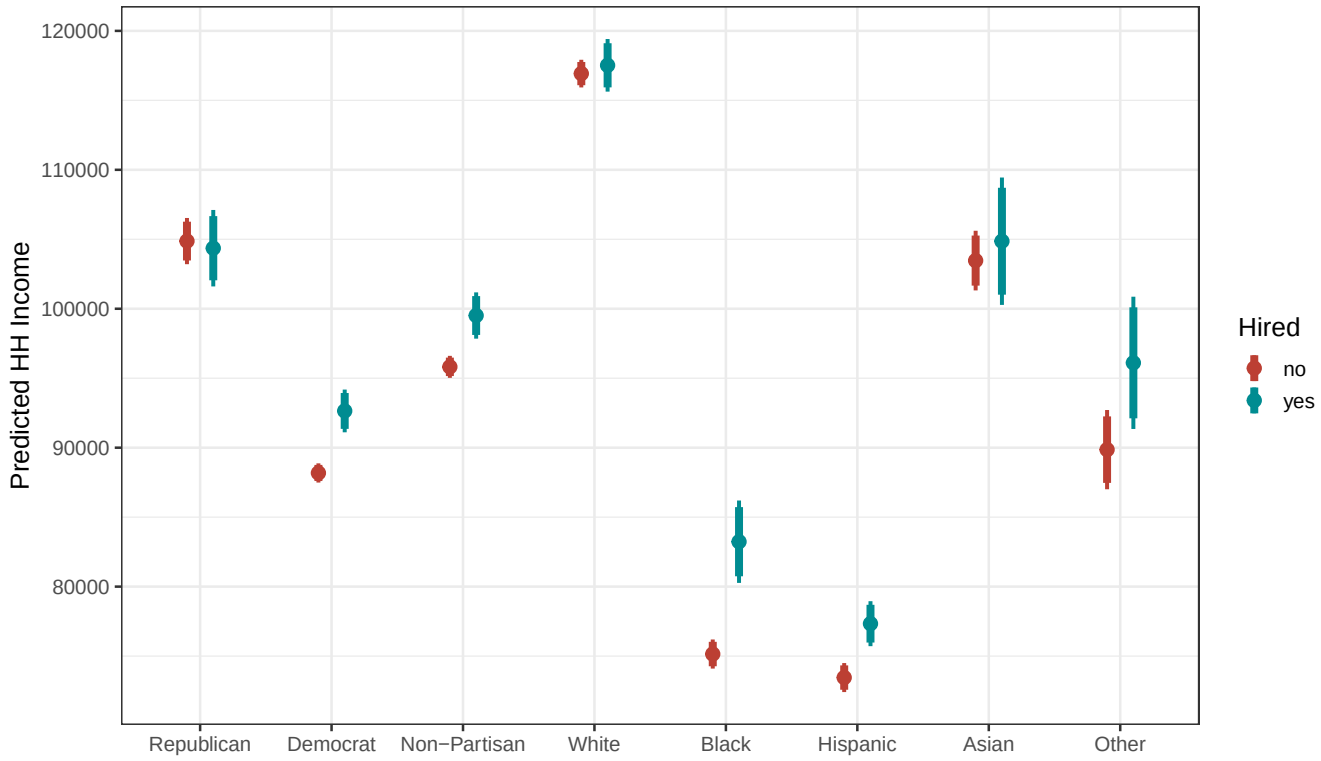

C.3 Career Progression

Figure A15: Difference in Behavior on the Job by Demographic and Tenure

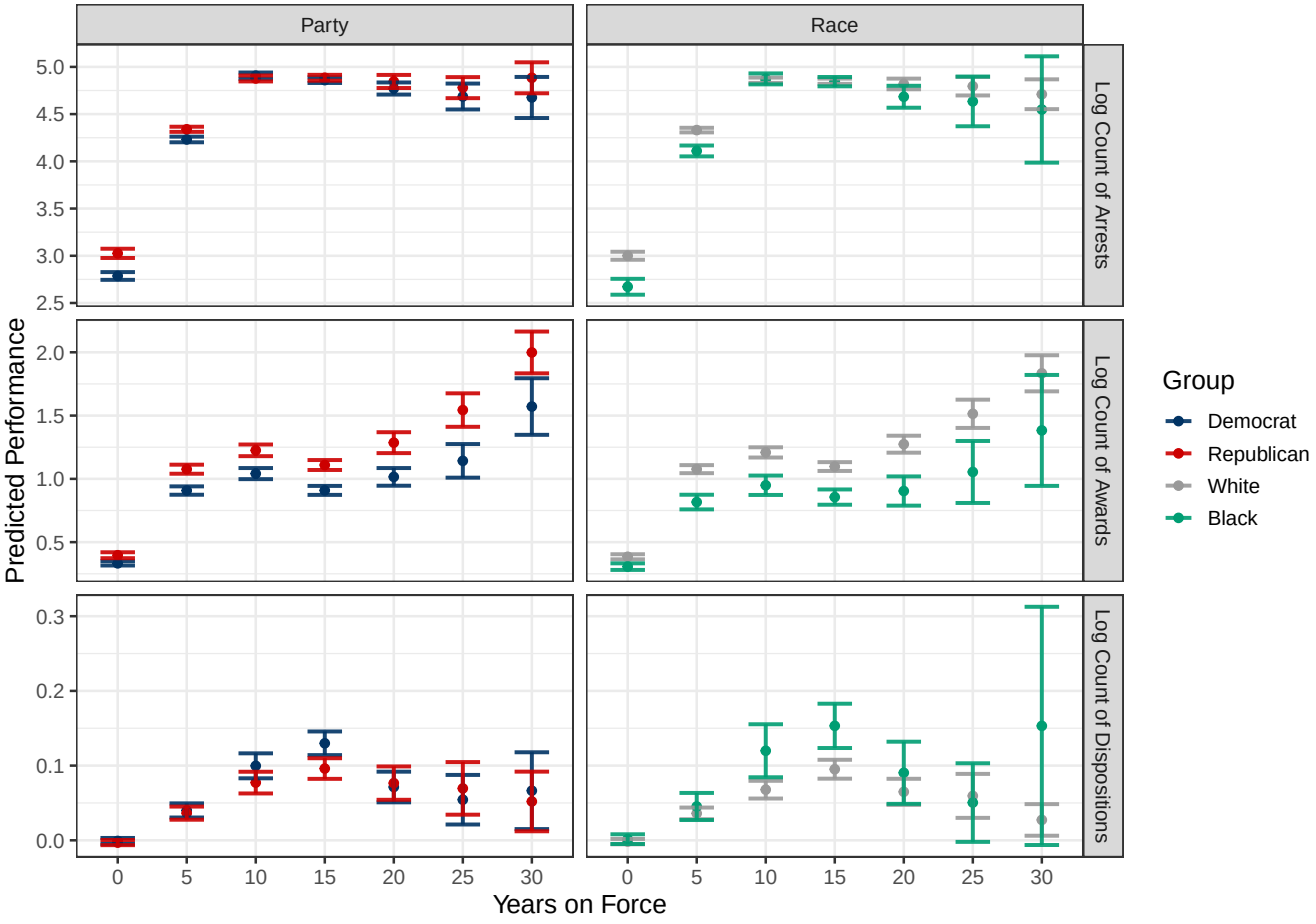

C.4 Exits

Figure A16

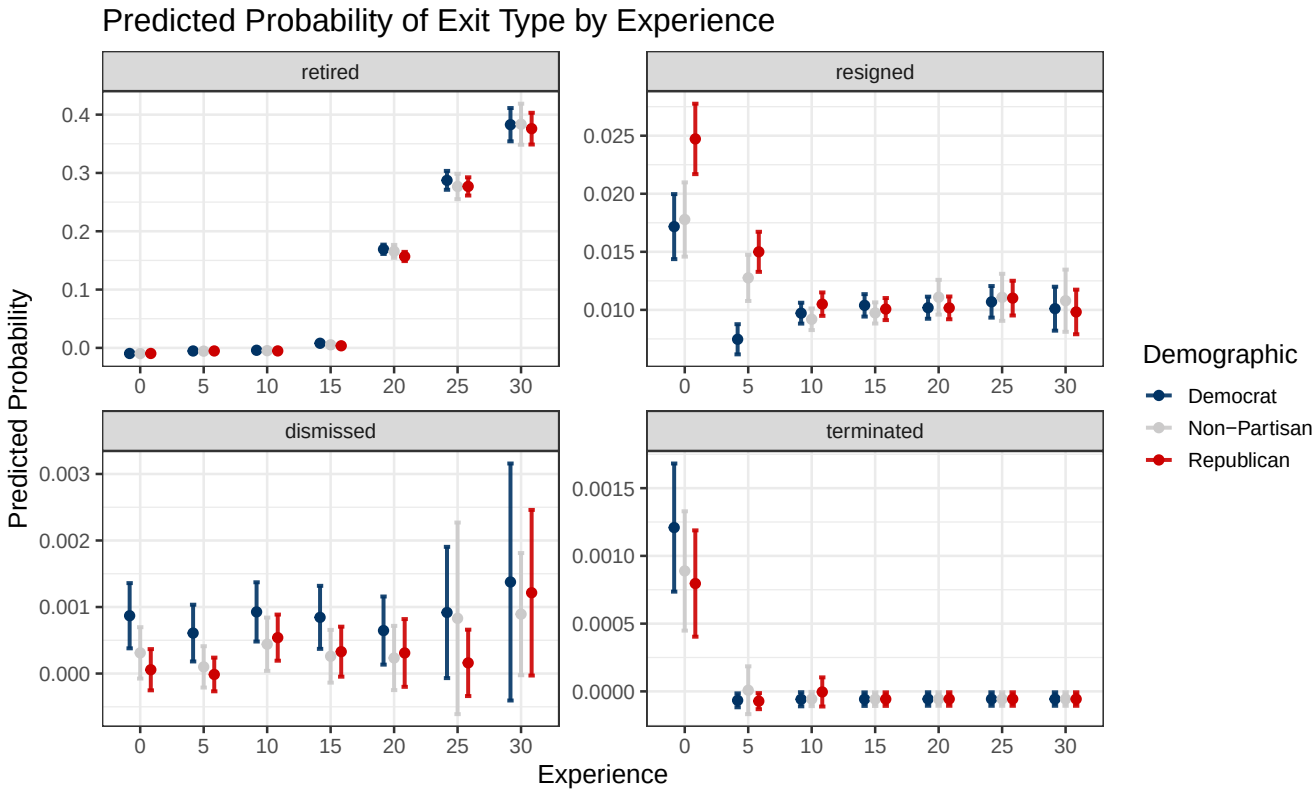

Figure A17

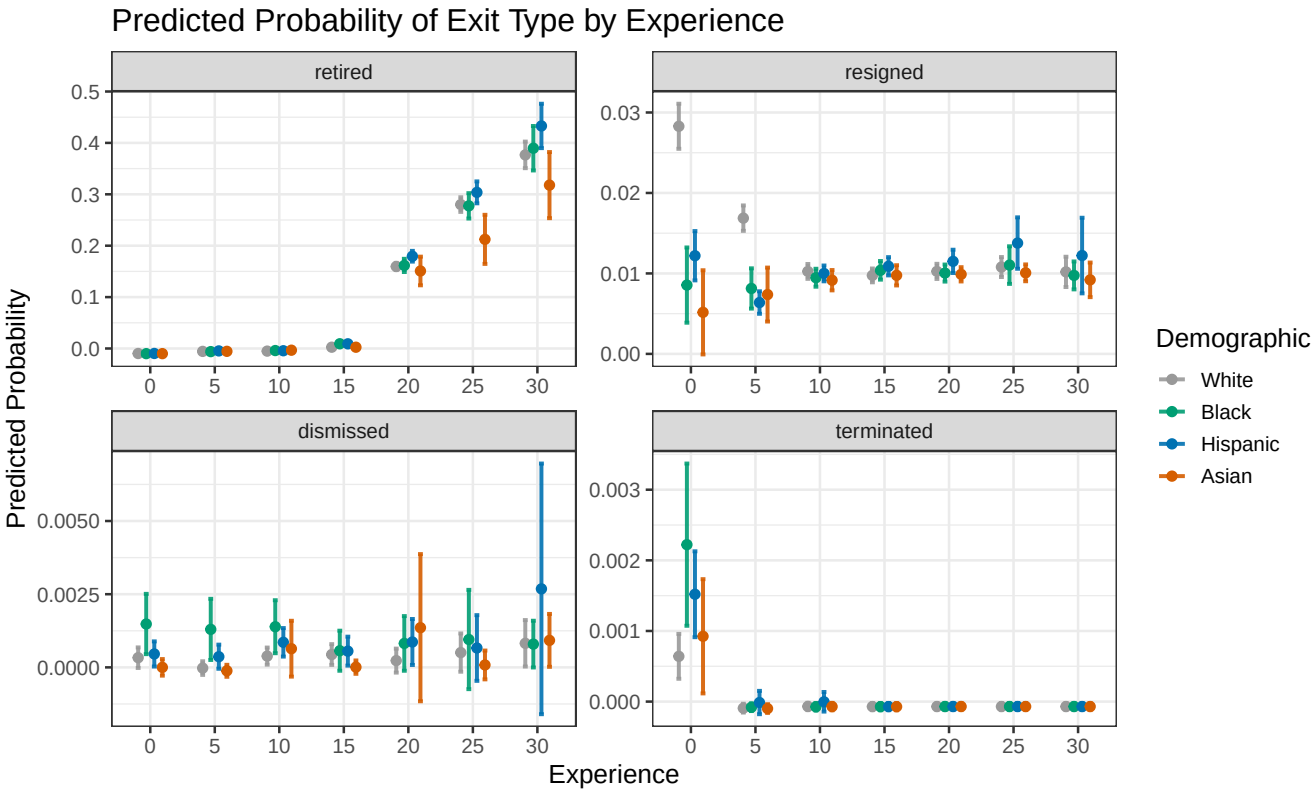
Predicted Probability of Exit Type by Experience
retired
resigned
dismissed
terminated
Predicted Probability
Experience
Demographic
White
Black
Hispanic
Asian

Figure A18: CCRB Records among HIred Police Officers 3 Years After Hiring