

Ballot Caps and Candidate Fields: Completion in New York City Ranked Choice Voting

Kolbe A. Dumas

Department of Political Science

University of Houston

Working paper, October 2026

Abstract

Ranked choice voting gives voters several opportunities to express preferences, although jurisdictions often limit how many candidates they may rank. I examine ballot completion around New York City's five-rank limit using Democratic council primaries in 2021, 2023, and 2025. Completion requires ranking every named candidate when fewer than five run and five candidates in larger fields. Across 68 district-elections, completion declines as fields grow from three through five candidates and increases above five. Election-level models with year controls and district-clustered uncertainty retain this pattern. Exploratory precinct models show lower completion in contexts with larger White population shares and a smaller positive association with college education. Adding candidate racial composition on an identical sample changes these average differences little, while evidence that demographic gradients change specifically at the cap remains uncertain. Raw evidence from Minneapolis and San Francisco provides descriptive context under different ranking limits. The findings document how use of the available ranking capacity varies around an institutional threshold, while leaving candidate-entry endogeneity, individual behavior, and the effects of changing the cap unresolved.

Keywords: ranked choice voting; ballot completion; electoral institutions; political participation; candidate fields; precinct demographics

1 Introduction

Ranked choice voting (RCV) allows voters to express preferences for several candidates, but the number of rankings available often depends on a jurisdiction’s rules. In New York City, a voter can rank up to five candidates in a municipal primary. When three candidates run, using every available ranking requires three choices; when seven run, it requires five. The difference matters for how ballot completion is understood, because larger candidate fields increase the completion requirement only until the five-rank limit binds.

Research on RCV distinguishes forced truncation from voluntary stopping. Forced truncation occurs when a rank limit prevents voters from ranking the full candidate field, while voluntary stopping occurs when voters leave some of their available rankings unused. Work from Kilgour et al. (2020) argues that rank limits can impose forced truncation and recommends dropping such restrictions where possible. Legal scholarship similarly argues that jurisdictions should generally allow voters to rank every candidate unless voting technology creates a practical constraint (Helgaas, 2025). These arguments establish why ranking capacity matters, while raising a further question about how voters use that capacity as candidate fields grow.

This paper examines completion in New York City’s Democratic council primaries in 2021, 2023, and 2025. I define a completed ballot as one that ranks every named candidate when fewer than five appear and five distinct named candidates when the field contains five or more. Completion records whether voters use the full opportunity to rank named candidates under the existing rule. Studying its prevalence helps connect the choices an electoral system offers to the preferences voters actually record. Partial ranking can reflect limited information, indifference toward remaining candidates, or a deliberate choice to support only some candidates, so completion alone cannot distinguish these explanations.

The primary analysis uses one observation per district-election, giving each of 68 elections equal weight. This places the candidate-field relationship at the level where candidate counts vary, while precinct-level models examine differences across demographic contexts within elections. I estimate separate slopes below and above five candidates, with district-clustered uncertainty and

adjustments for election year and incumbency. Realized first-choice margin enters a separate specification because candidate entry may itself affect the distribution of votes.

Completion falls from 40.4% in three-candidate elections to 22.1% in five-candidate elections, averaging each election equally. It then rises to 30.2% in eight-candidate elections. The election-level model retains a negative below-cap slope and a positive above-cap slope. The decline thus gives way to an increase across fields where the required number of rankings stays at five. Candidate appeal, campaign information, and the composition of the electorate may also vary across these elections, leaving the source of the reversal unresolved.

I treat demographic heterogeneity as exploratory precinct-context evidence. Within elections, larger White population shares are associated with lower completion, while BA+ share has a smaller positive association. Adding the racial composition of candidate fields changes these average differences little on an identical sample. Evidence that either demographic relationship changes specifically at five candidates remains uncertain. Raw evidence from Minneapolis and San Francisco provides further descriptive context under different ranking limits.

The paper contributes evidence on how fully voters use the rankings available to them across different candidate fields. In New York City, completion is lowest near the point at which voters must rank every named candidate to reach five choices. In larger fields, voters can reach the same depth while leaving some candidates unranked. This pattern makes the relationship between candidate supply and ranking requirements central to interpreting completion, including comparisons that might otherwise attribute lower completion to less political engagement.

2 The New York City Ballot and Raw Completion

New York City provides a useful setting for examining how a fixed ranking limit interacts with different candidate fields. The five-rank limit applies across the included elections, while the number of named candidates ranges from three to fifteen. The sample contains 42 elections in 2021, 11 in 2023, and 15 in 2025. These elections provide different ranking tasks under the same limit: com-

pletion requires every named candidate in a three- or four-candidate field, five choices when five candidates run, and five choices from a larger set once the field exceeds five. The institutional rule therefore remains fixed while the requirement for a completed ballot changes with the candidate field.

Figure 1 presents the raw completion rates before any electoral or demographic adjustment. Averaging elections equally, completion declines from 40.4% with three candidates to 31.5% with four and 22.1% with five. It then increases to 24.4% with six, 26.6% with seven, and 30.2% with eight. The individual election observations show considerable variation around these averages, making the pattern a starting point for analysis rather than a complete account of ballot use. Below five, larger fields require another distinct ranking for completion; above five, the required ranking depth stays fixed even as voters have more candidates from which to choose.

The choice of average matters for interpreting these numbers. Giving elections equal weight describes the completion rate in the average included contest, so a larger electorate does not give a single candidate field more influence. Pooling ballots instead describes the completion rate across participating records and gives larger elections more weight. Under the latter approach, completion is 35.7% in three-candidate fields and 20.6% in five-candidate fields, compared with the respective equal-election means of 40.4% and 22.1%. Both summaries concern the same completed ballots, but answer different descriptive questions. I foreground the equal-election means because the candidate-field comparison is organized around elections, while the supplement retains the ballot-pooled rates.

The available comparisons also vary across election years. All included 2023 fields contain three or four candidates, while the largest fields are concentrated in 2021. Average completion across elections is 30.4% in 2021, 43.8% in 2023, and 26.5% in 2025. Year controls therefore matter for the candidate-field comparison, particularly because the 2023 contests contribute no observations at or above five. Support immediately around five is also uneven: 15 elections have four candidates, 6 have five, and 10 have six. Several of the largest candidate counts contain only one election. The raw figure makes this support visible before the model summarizes the

relationship and helps distinguish a recurring pattern from the experience of any one contest.

3 Ranking Capacity and Ballot Use

3.1 Completion, Truncation, and Political Expression

Ballot completion matters because ranking additional candidates can preserve a ballot's role in later counting rounds. An exhausted ballot can no longer be transferred because none of the candidates it ranks remains eligible to receive it (Burnett & Kogan, 2015). A voter who stops early may therefore lose an opportunity to express a preference between the eventual finalists. Exhaustion also qualifies majoritarian claims about RCV, since the winner may receive a majority of continuing ballots while falling short of a majority of the original votes cast (Graham-Squire & McCune, 2025).

The connection between completion and exhaustion nevertheless depends on which candidates are ranked. A partially ranked ballot can remain active through the decisive round, while a ballot using the full five-rank allowance can exhaust if the field contains more than five candidates. Ranking everyone in a small field also requires including the voter's least-preferred candidate, whose ranking may never affect the winner. Completion therefore captures use of an available institutional capacity rather than a sufficient measure of electoral influence. This distinction is especially relevant when a rank limit changes how much capacity the ballot offers.

Boudreau et al. (2026) examine political expression in New York City elections through voters' willingness to rank multiple candidates and include a finalist, showing how these outcomes vary across local offices and with endorsements. They question the emphasis on using every available ranking because the final ranking is rarely pivotal. Their argument gives an important reason to distinguish the depth of preferences recorded from whether a ballot influences the decisive round. Completion addresses how fully voters use the opportunity to express preferences, while finalist inclusion addresses whether the ballot includes a candidate who remains in contention.

Completion is useful for studying participation within the ballot because it identifies where voters reach the available ranking limit and where some ranking opportunities remain unused.

Two voters may both rank multiple candidates and include a finalist while differing in how many additional preferences they record. Reaching the completion benchmark captures one difference in the extent of this expression, although individual rankings may never be used during counting. Voters may also have good reasons to leave candidates unranked, so a completed ballot need not represent a more informed or more desirable choice. The institutional question is how often voters use the full ranking opportunity when reaching it requires selecting every candidate, compared with selecting only part of a larger field.

Ballot-use outcomes should also be distinguished from errors that jeopardize otherwise expressed preferences. Pettigrew and Radley (2026) identify overvotes, repeated rankings, and skipped rankings across RCV contests, showing why a count of occupied positions can overstate distinct preference expression. The present analysis counts a candidate once even if the name appears in several ranks. Write-ins and overvote labels are retained as separate features rather than treated as additional named candidates. The measure concerns recorded choices in the export; it does not reconstruct every jurisdiction’s error-handling or tabulation rules.

3.2 Candidate Fields and Ranking Demands

The information required to rank several candidates provides one reason completion may fall as fields expand. Research on ordinary ballots links additional decisions to choice fatigue and reliance on shortcuts (Augenblick & Nicholson, 2016), while candidate information and political sophistication also shape ballot completion (Lamb & Perry, 2020). In primary elections, where candidates share a party label, voters may have fewer readily available cues for distinguishing their options. Goidel and Armstrong (2025) find that larger primary fields lead respondents to spend less time on each candidate and feel more overwhelmed, although their experiment does not find the anticipated increase in abstention.

Under RCV, these information demands concern both identifying acceptable candidates and ordering them. An additional candidate can increase the number of comparisons voters consider even when the number of ranks they may use remains fixed. Larger fields may also offer more

candidates a voter is willing to support, attract more campaign attention, or draw different voters into the election. Consequently, a rank limit fixes the maximum completion requirement without fixing the wider informational environment of the contest.

Below New York City’s five-rank limit, each additional candidate raises the number of distinct named choices required for completion. Above five, the field can continue to grow while the completion requirement remains unchanged. This institutional distinction motivates a directional expectation about the observed relationship:

Hypothesis 1. Ballot completion declines as the candidate field expands below the ranking limit, and this decline attenuates once the limit binds.

The hypothesis concerns the completion measure under a fixed rule. A change in its slope need not establish that voters experience less difficulty, because the threshold defining completion itself stops increasing. Nor does the hypothesis require completion to rise above the cap. A positive post-cap relationship would require further explanation involving how voters, candidates, or campaigns differ across larger fields. The empirical analysis assesses the pattern while leaving these behavioral explanations open.

3.3 Precinct Context and Candidate Availability

Existing research also motivates examining differences across demographic contexts. Cormack (2024) finds higher voidable-overvote rates in New York City Assembly Districts with lower educational attainment. Revisiting the 2025 mayoral primary, Cormack (2026) finds that the overall overvote rate declined while the education gradient persisted in attenuated form. Surveys likewise identify education differences in reported understanding of RCV, although demographic patterns depend on whether the outcome is understanding or ranking multiple candidates (Donovan et al., 2022).

Candidate identity provides another possible source of differences in ballot use. McCarty (2024) examines final-round exhaustion and ranking patterns, finding that aggregate racial and

ethnic relationships vary with the identity of the finalists. Such findings motivate considering the initial field as well: the number of candidates from a particular category may matter differently across precinct contexts. However, familiarity, endorsements, policy agreement, and candidate popularity can accompany the same categories, making the presence of a candidate an incomplete measure of the choices voters perceive.

I therefore ask whether completion and its relationship to field size differ across precinct contexts, and whether adding candidate composition accounts for the average White-share and education associations. These are exploratory questions about aggregate relationships. The demographic composition of residents does not reveal the demographic composition of participating voters, and a candidate’s recorded category does not identify the race or ethnicity of the voters ranking that candidate.

4 Data and Measurement

4.1 Election Records and the Completion Benchmark

The analysis draws on cast vote records (CVRs) from the Harvard Dataverse collection of single-winner RCV elections (Otis, 2022). The New York City sample contains Democratic council primaries in 2021, 2023, and 2025 with at least three named candidates and usable precinct context. Two-candidate contests are excluded because their ranking task offers little scope for the multi-candidate comparison studied here. The main sample includes 1,091,849 engaged ballot records in 6,387 precinct-election cells, spanning 68 district-elections and 48 district labels.

Let R_i denote the number of distinct named candidates ranked on ballot i , K_e the named-candidate field in election e , and C_e the ranking limit. The completion indicator is

$$Y_i = \mathbf{1}\{R_i \geq \min(K_e, C_e)\}. \quad (1)$$

In New York City, $C_e = 5$. A ballot with at least one named candidate enters the denominator. This

restriction measures completion among ballots that participate in the named council contest, rather than combining voluntary stopping with blank or write-in-only records. Repeated appearances of a candidate count once, generic write-ins do not expand the named field, and overvotes are not assigned to a single candidate. These choices make the numerator and denominator comparable across the cases while keeping the outcome distinct from officially tabulated exhaustion.

The election outcome is the percentage of these engaged ballots that complete the benchmark. Ballots are pooled within each election, then elections receive equal weight in the primary model. This gives a small and a large district-election the same role in estimating the relationship across fields. The supplement also reports ballot-pooled raw rates, an alternative occupied-position outcome, and a sensitivity retaining additional precincts from the same elections. Candidate fields are reconciled to observed named candidates and the available official election records, with known removals and ambiguous labels documented in the replication materials.

4.2 Demographic and Electoral Context

Precinct demographic measures combine Census population composition and American Community Survey education, income, age, and naturalization estimates, linked through election-year precinct geography. The racial-context variables are mutually exclusive non-Hispanic White, Black, and Asian shares together with Hispanic share; other residents form the remaining category. BA+ share measures four-year college attainment, while the other controls capture median household income, age 65+ share, and naturalized-citizen share. These measures describe precinct residents and their socioeconomic context, rather than the characteristics of individual ballot casters.

The main electoral controls include election year and whether an incumbent is present. I also examine the district-election first-choice margin, calculated as

$$\text{Margin}_e = \frac{V_{(1)e} - V_{(2)e}}{V_e}, \quad (2)$$

where $V_{(1)e}$ and $V_{(2)e}$ are the named first-choice votes for the first- and second-place candidates and V_e is the total valid named first-choice vote. Larger values indicate less competitive first-choice results. Because the candidate field can affect this realized margin, conditioning on it produces a supplementary comparison rather than automatically removing a predetermined confounder.

Candidate-composition measures use existing research workbooks, voting-rights reports, and documented public sources (Handley, 2022; RepresentWomen, 2021, 2023, 2025). I retain fields only when all named candidates can be assigned under the source crosswalk to White, Black, Hispanic, or Asian categories. Broad, ambiguous, conflicting, and unknown classifications remain excluded from the primary composition sample, while an explicit Hispanic identity takes precedence in the operational crosswalk. This produces 39 elections, 31 district labels, and 3,532 precinct-election cells. The source categories are research classifications rather than uniformly verified non-Hispanic self-identification; the supplement documents the crosswalk, provisional person links, and coverage limitations.

5 Empirical Strategy

5.1 Election-Level Completion

The primary model follows the level at which candidate-field size varies:

$$\bar{Y}_e = \alpha + \beta(K_e - 5) + \theta \max(K_e - 5, 0) + \lambda_t + \gamma I_e + \varepsilon_e, \quad (3)$$

where $\bar{Y}_e$ is completion in percentage points, λ_t denotes year indicators, and I_e indicates an incumbent. The below-cap slope is β , the change in slope at five is θ , and the above-cap slope is $\beta + \theta$. I estimate the model by ordinary least squares with equal election weights. Linear percentage-point coefficients make the substantive change directly interpretable; fitted percentages remain within the outcome's bounds on the estimation sample.

The main election-level results and the precinct-context comparisons use district-clustered

CR2 uncertainty with small-sample degrees of freedom (Pustejovsky & Tipton, 2018). This calculation allows observations from the same district to be related without treating them as independent evidence. Precincts improve the measurement of each election’s completion rate, but they do not supply additional independent candidate fields. The supplement reports a district-label fixed-effects specification, while treating geographic continuity cautiously because the labels span the 2023 redistricting. Splitting district labels into boundary eras leaves too few repeated elections to sustain a separate within-era model.

Candidate entry is not randomly assigned around five. Larger fields can arise in open seats, more salient contests, or places with different political networks and electorates. Consequently, the hinge estimates describe a conditional association around a fixed institutional threshold. They do not identify the effect of raising or removing the cap. The changing completion benchmark is part of that association and must remain visible when interpreting the result.

5.2 Joint Precinct-Context Models

The demographic analysis compares precincts within the same election, using an election-specific intercept and all eight demographic and resource measures jointly. Each context gradient may vary with candidate-field size below and above five:

$$y_{pe} = \alpha_e + X'_{pe}b + (K_e - 5)X'_{pe}g + \max(K_e - 5, 0)X'_{pe}h + u_{pe}. \quad (4)$$

Here y_{pe} is precinct completion in percentage points, and X_{pe} contains the demographic measures. Each election receives equal total weight across its precincts. The election intercept absorbs conditions shared throughout the contest, including the main effects of year, incumbency, margin, candidate count, and candidate composition. Unmeasured conditions that affect precincts differently can still shape the gradients.

I summarize the exploratory White-share and education associations through ten-percentage-point context contrasts. These comparisons describe average completion differences

across precinct contexts, rather than effects of the cap on demographic groups. Each precinct's lower and higher profiles lie five points below and above its observed share, and both must fall inside the election's central observed range. For White share, other racial shares are reallocated proportionally so the composition remains coherent; resources remain fixed. For BA+ share, the other measures remain fixed. Contrasts are averaged within elections and then equally across supported elections, with district CR2 uncertainty calculated for the aggregate comparison.

The candidate-composition models add interactions between each precinct racial share and the available number of candidates coded in the corresponding category. A more flexible specification allows all four precinct racial shares to interact with candidate racial shares. The before-and-after comparisons use identical observations, weights, outcomes, and prediction populations. Five exploratory tests evaluate changes at the cap for White composition, education, age, income, and naturalization, with Holm adjustment within each specification (Holm, 1979). Average context differences and changes in those differences at the cap remain separate quantities.

6 Results

6.1 The Completion Pattern Around Five Candidates

The raw evidence in Figure 1 establishes the pattern that the election model summarizes. Completion declines toward five candidates and rises in larger fields, while differences in the calendar and individual elections leave room for observed electoral context to account for part of the relationship.

Table 1 shows that the relationship survives adjustment for year and incumbency. The below-cap slope is -7.25 percentage points per additional candidate, with a 95% confidence interval of [-10.53, -3.98]. The slope change is 9.66 points ([5.99, 13.33]), yielding a positive post-cap slope of 2.41 points ([1.69, 3.12]). Adding first-choice margin produces a below-cap slope of -7.68 and an above-cap slope of 2.23. These specifications support attenuation of the preceding decline, with an observed reversal rather than a flat post-cap relationship.

Figure 2 translates the election model into adjusted completion levels. Averaging over the observed distribution of election years and incumbency, predicted completion is 38.0% with three candidates, 23.5% with five, and 28.3% with seven. The figure shows the fitted relationship under New York City’s existing rule and carries its district-clustered uncertainty. It does not extend the below-cap slope into a hypothetical policy setting without a ranking limit.

The positive post-cap slope does not follow from cognitive load alone. Reaching the five-rank cap fixes the number of choices required for completion, while leaving the number of candidates to assess free to grow. A larger field may also increase the supply of acceptable candidates or accompany different campaign and turnout conditions. The result therefore establishes a structured completion pattern around the benchmark, while leaving these accounts observationally compatible with the evidence.

The local-window checks qualify the precision of this pattern. Restricting the analysis to four through six candidates retains a positive slope-change estimate, but its confidence interval includes zero. Wider four-to-seven and three-to-eight windows retain positive changes with intervals above zero, as the supplementary local estimates show. The evidence is therefore stronger for the broader completion pattern than for locating a distinct change precisely at five. The model places the hinge at the institutional threshold, while the narrowest comparison leaves considerable uncertainty about the size of the local change.

6.2 Exploratory Completion Differences Across Precinct Contexts

Table 2 reports the average precinct-context differences. In the full sample, a ten-point increase in White share, accompanied by proportional reductions in the other racial shares, corresponds to -1.75 completion points, with a 95% interval of [-2.22, -1.28]. A ten-point increase in BA+ share corresponds to 0.67 points ([0.19, 1.16]). Both comparisons hold the remaining included context measures fixed and use within-election variation, so differences shared by entire elections cannot alone explain them.

Restricting the comparison to fields with complete candidate categories leaves the White-

share association at -1.58 points $([-2.03, -1.12])$. Adding own-category candidate availability changes it to -1.55 $([-1.99, -1.11])$, and the full candidate-share interaction model produces -1.49 $([-1.92, -1.06])$. The corresponding BA+ differences remain positive, ranging from 0.50 before candidate adjustment to 0.39 under the full interaction specification. Candidate composition therefore accounts for little of these average contextual differences under the models examined.

The change in the White-share difference after adding own-category availability is only 0.029 points, with a paired district-deletion interval of $[-0.33, 0.39]$. Negative White-share differences also persist when provisional person links are excluded and when completion is defined using occupied positions. These sensitivities strengthen the descriptive finding, while leaving measurement limits and unobserved precinct characteristics relevant to its explanation.

Figure 3 summarizes the average White-share and BA+ completion differences across the four specifications, with district-clustered confidence intervals. The candidate-adjustment comparisons use the same observations and supported precinct profiles. These average associations provide no clear evidence that the cap closes a racial completion gap or widens an education-based gap. None of the five exploratory demographic tests of a slope change at five survives Holm adjustment; the supplement reports the full family.

6.3 Minneapolis and San Francisco as Descriptive Context

To supplement this analysis, I take a look at both Minneapolis and San Francisco for further descriptive context. Minneapolis's three-rank limit is already reached in every included field, while San Francisco's ten-rank limit exceeds the observed field sizes. The supplement presents raw election-level outcomes and sample counts rather than fitted curves. The raw patterns vary across candidate counts and years, including thin and nonmonotonic portions that a straight fitted slope would conceal. Differences in election type, candidate supply, electorate, ballot design, and administrative history prevent these cities from serving as controls for New York City.

7 Discussion and Conclusion

Completion in New York City’s council primaries declines as candidate fields expand toward five, while the election-level models and raw summaries show an increase above five. The reversal occurs across fields where the number of rankings required for completion has stopped increasing. The association persists when candidate-field variation is represented by elections rather than thousands of precinct observations. Because all included elections operate under the same five-rank rule, these comparisons leave the effect of changing the cap unresolved.

The results also clarify what completion contributes to debates about electoral participation. Electoral rules offer opportunities to record preferences beyond a first choice, and completion identifies how often voters use the full allowance for named candidates. In the NYC sample, the frequency of doing so differs substantially between fields that require ranking everyone and fields that allow some candidates to remain unranked. Interpreting these differences as levels of engagement without considering the ranking requirement would overlook a feature of the ballot itself. Assessments of ranking limits should consider this use of available choices alongside the candidates ranked, continuing-ballot participation, and the representation achieved through the count.

The exploratory precinct findings raise a further question about the contexts in which ranking opportunities are used. Completion is lower in precincts with larger White shares and somewhat higher where college attainment is greater, conditional on the joint controls and election intercepts. Available candidate-category counts do little to account for these average relationships. The associations concern precinct composition, with no robust evidence that either relationship changes specifically at the cap. Candidate familiarity, turnout composition, endorsements, information, and other neighborhood differences remain unresolved. A planned survey experiment will examine individual ranking choices under controlled ballot conditions in future work.

The main limits concern candidate-entry endogeneity, the changing completion benchmark, and the ecological nature of the demographic evidence. Cross-city patterns provide context but do not identify the consequences of changing ranking rules. Within these boundaries, the paper doc-

uments how use of named-candidate ranking capacity varies around a consequential institutional threshold and shows why completion levels, demographic context, and behavioral explanations should be evaluated separately.

References

- Augenblick, N., & Nicholson, S. (2016). Ballot position, choice fatigue, and voter behaviour. *The Review of Economic Studies*, 83(2), 460–480. <https://doi.org/10.1093/restud/rdv047>
- Boudreau, C., Colner, J., & MacKenzie, S. A. (2026). Start spreading the views: Improving political expression and representation in local elections using ranked choice voting. *Political Research Quarterly*, 79(3), 1082–1098. <https://doi.org/10.1177/10659129261454685>
- Burnett, C. M., & Kogan, V. (2015). Ballot (and voter) "exhaustion" under instant runoff voting: An examination of four ranked-choice elections. *Electoral Studies*, 37, 41–49. <https://doi.org/10.1016/j.electstud.2014.11.006>
- Cormack, L. (2024). More choices, more problems? Ranked choice voting errors in New York City. *American Politics Research*, 52(3), 306–319. <https://doi.org/10.1177/1532673X231220640>
- Cormack, L. (2026). More chances, fewer problems? Revisiting ranked choice voting errors in New York City. *American Politics Research*, 54(3), 340–346. <https://doi.org/10.1177/1532673X261420766>
- Donovan, T., Tolbert, C., & Harper, S. (2022). Demographic differences in understanding and utilization of ranked choice voting. *Social Science Quarterly*, 103(7), 1539–1550. <https://doi.org/10.1111/ssqu.13215>
- Goidel, S., & Armstrong, B. (2025). Choice overload in crowded primary elections. *American Politics Research*, 53(1), 3–16. <https://doi.org/10.1177/1532673X241281246>
- Graham-Squire, A., & McCune, D. (2025). An examination of ranked-choice voting in the United States, 2004–2022. *Representation*, 61(1), 1–19. <https://doi.org/10.1080/00344893.2023.2221689>
- Handley, L. (2022). *Report to the New York City Districting Commission: Analysis of voting patterns by race and assessment of compliance of proposed New York City Council district plan with the Voting Rights Act*. New York City Districting Commission. <https://www.nyc.gov/assets/districting/downloads/pdf/RBV-Report.pdf>

- Helgaas, T. (2025). The instant runoff voting patchwork: Handling voter errors and ballot length. *University of Pennsylvania Law Review*, 173(4), 1217–1267. <https://doi.org/10.58112/UPLR.173-4.5>
- Holm, S. (1979). A simple sequentially rejective multiple test procedure. *Scandinavian Journal of Statistics*, 6(2), 65–70. <https://www.jstor.org/stable/4615733>
- Kilgour, D. M., Grégoire, J.-C., & Foley, A. M. (2020). The prevalence and consequences of ballot truncation in ranked-choice elections. *Public Choice*, 184(1-2), 197–218. <https://doi.org/10.1007/s11127-019-00723-2>
- Lamb, M., & Perry, S. (2020). Knowing what you don't know: The role of information and sophistication in ballot completion. *Social Science Quarterly*, 101(3), 1132–1149. <https://doi.org/10.1111/ssqu.12788>
- McCarty, N. (2024). *Minority electorates and ranked choice voting*. <https://electionconfidence.org/wp-content/uploads/2024/01/FINAL-RCV-study-1-10-24.pdf>
- Otis, D. (2022). *Single winner ranked choice voting CVRs*. Data set; Harvard Dataverse. <https://doi.org/10.7910/DVN/AMK8PJ>
- Pettigrew, S., & Radley, D. (2026). Overvotes, overranks, and skips: Mismarked and rejected votes in ranked choice voting. *Political Behavior*, 48(1), 229–251. <https://doi.org/10.1007/s11109-025-10028-4>
- Pustejovsky, J. E., & Tipton, E. (2018). Small-sample methods for cluster-robust variance estimation and hypothesis testing in fixed effects models. *Journal of Business & Economic Statistics*, 36(4), 672–683. <https://doi.org/10.1080/07350015.2016.1247004>
- RepresentWomen. (2021). *New York City Council candidate research workbook, 2021*. Research data. <https://docs.google.com/spreadsheets/d/19j8B5KCsar6jW3FqAQ27Q9o5Tok4v9qb5HzvmUspxCE/edit>
- RepresentWomen. (2023). *New York City Council candidate research workbook, 2023*. Research data. <https://docs.google.com/spreadsheets/d/1mNjHWCjQqTbj3ixWOfMnexJJJ4FQh-3cSZZSxhpcKrQ/edit>

RepresentWomen. (2025). *New York City Council candidate research workbook*, 2025. Research data. https://docs.google.com/spreadsheets/d/1ZMsL_53kcBsfGEfu5YsDwSz3ulFxkiHd2shbyJatMkE/edit

Table 1: Completion Declines Below Five Candidates and Increases Above Five

	Election Controls	Also Margin
Below-cap slope	-7.25*** [-10.53, -3.98]	-7.68*** [-10.92, -4.45]
Slope change	9.66*** [5.99, 13.33]	9.92*** [6.23, 13.60]
Above-cap slope	2.41*** [1.69, 3.12]	2.23*** [1.42, 3.04]
Incumbent present	-0.66 [-3.90, 2.59]	0.16 [-3.28, 3.60]
First-choice margin (10 pp)		-0.69* [-1.32, -0.07]
Elections	68	68
District clusters	48	48
Year indicators	Included	Included
Adjusted R^2	0.68	0.69


Percentage-point coefficients with district CR2 95% confidence intervals. * $p < .05$; ** $p < .01$; *** $p < .001$.

Table 2: The White-Share Completion Difference Persists After Candidate Adjustment

	All Fields	Common Fields	Own-Group	Full Composition
White difference	-1.75*** [-2.22, -1.28]	-1.58*** [-2.03, -1.12]	-1.55*** [-1.99, -1.11]	-1.49*** [-1.92, -1.06]
Education difference	0.67** [0.19, 1.16]	0.50* [0.11, 0.88]	0.47* [0.12, 0.82]	0.39* [0.05, 0.73]
Precinct-election cells	6,387	3,532	3,532	3,532
Elections	68	39	39	39
District clusters	48	31	31	31


Exploratory average differences per ten-point context shift; district CR2 95% intervals. * $p < .05$; ** $p < .01$; *** $p < .001$.

Figure 1: Completion Falls Toward Five Candidates Before Rising in Larger Fields




Points represent elections, black ticks show equal-election means, and counts identify the supporting elections. These raw summaries carry no inferential intervals.

Figure 2: The Adjusted Completion Relationship Changes at Five Candidates



Predictions average observed year and incumbency profiles; shading shows district CR2 95% confidence intervals.

Figure 3: Exploratory Completion Differences Across Precinct Contexts



Average differences per ten-point context shift; bars show pointwise district CR2 95% confidence intervals. Candidate-adjustment comparisons use identical samples. These are contextual associations, not cap effects.

Supplement to Ballot Caps and Candidate Fields

1 Measurement, Coverage, and Inference

1.1 Distinct Named Choices and the Engaged Denominator

The completion numerator counts distinct named candidates selected anywhere in the exported ranking positions. Repeated names count once, and generic write-in, overvote, skipped, undervote, and unknown markers do not become named candidates. A record enters the engaged denominator when at least one named choice appears. The measure retains later recorded marks rather than applying a full jurisdiction-specific tabulation algorithm, so it does not establish whether every recorded preference would remain usable after an earlier error.

The New York City sample uses the same 68 elections in each electoral specification, with 6,387 context-complete precinct-election cells. Table A.1 distinguishes equal-election means from pooled ballot rates. Candidate count is reconciled to named choices and election records; the 2021 District 16 field contains four candidates following the documented removal of Eric Stevenson. A zero-vote entry in a results file is not sufficient by itself to establish that a candidate remained on the ballot.

Retaining additional precincts from these same elections produces a below-cap slope of -7.27 points (95% CI $[-10.54, -3.99]$), a slope change of 9.67 ($[6.01, 13.34]$), and an above-cap slope of 2.41 ($[1.69, 3.12]$). These estimates are close to the primary context-complete specification. An alternative completion outcome counts occupied positions, allowing repeated choices and special labels to contribute differently. In the candidate-complete sample, its White-share con-

trast is -1.68 points ($[-2.17, -1.19]$), compared with -1.58 under the distinct-name measure. The measures remain substantively different despite the similar contextual association.

1.2 Election-Level and Local Specifications

Table A.2 reports the year-and-incumbency model, the margin-adjusted specification, and the district-label fixed-effects model. The last comparison uses 38 elections in 18 repeated district labels. Because these labels span redistricting, the model should not be interpreted as a verified constant-geography comparison. Separating pre-2023 and 2023-onward labels leaves six repeated elections in three districts, which is insufficient for a credible separate specification.

Table A.3 reports models restricted to four through six, four through seven, and three through eight candidates, retaining year and incumbency controls. In the narrowest window, the estimated slope change remains positive but its interval crosses zero. The wider windows provide more precise evidence of attenuation. These windows are checks of an observational relationship, rather than local randomization around a candidate-entry threshold.

Election-level and precinct-context regression intervals use district-clustered CR2 and contrast-specific small-sample degrees of freedom. A post-cap slope is calculated from the sum of the candidate-count and hinge coefficients, using their covariance. The variance of a context contrast is calculated after averaging its design differences, rather than by averaging precinct intervals. These choices follow the small-sample inference framework described by Pustejovsky and Tipton (2018). The earlier mixed-logit sensitivity uses a district-deletion jackknife, and matched comparisons use the district-endpoint calculation described below. All reported model uncertainty is calculated across districts; raw counts and observed means have no inferential intervals. These calculations allow dependence within district labels but do not remove spatial dependence between different labels or uncertainty in demographic interpolation.

2 Demographic and Candidate-Composition Comparisons

2.1 Context Profiles and Multiple Testing

Resident racial shares come from mutually exclusive 2020 Census population categories, spatially linked to the precinct boundaries for each election year. The operational categories are non-Hispanic White alone, Black alone, Asian alone, and Hispanic of any race; other and multiracial residents remain in a residual category. Small source-rounding totals above 100% are normalized only where the overflow is less than 0.25 points. Education, income, age, and naturalization use the existing ACS-derived precinct measures.

Each White-share comparison shifts the observed share five points down and five points up. The other racial shares move proportionally, preserving the composition, and resources remain fixed. The BA+ comparison changes education by the same ten-point amount while holding the other measures fixed. Both endpoints must lie within the election's 5th-to-95th-percentile range for the moderator. This is a marginal support guard; it does not establish complete overlap for every combination of race, resources, and candidate field.

The full-sample White contrast uses 3,237 supported cells in 57 elections; the education contrast uses 4,586 cells in all 68 elections. In the candidate-complete sample, the corresponding populations contain 1,742 cells in 31 elections and 2,544 cells in 39 elections. Every before-and-after candidate comparison uses the same supported population for its moderator. Each election receives equal weight after its supported precincts are averaged.

Table A.4 separates these average differences from changes in demographic gradients at the cap. The five-test family includes White composition, BA+ share, age 65+ share, median income, and naturalized-citizen share. Holm's sequential adjustment orders the five raw p-values and evaluates each against a threshold that accounts for the remaining family (Holm, 1979). No cap-specific interaction survives this correction in the main specifications. The intervals are pointwise 95% intervals, while stars use the family-adjusted p-values.

The White contrast is measured per ten-point composition shift; BA+, age, and naturaliza-

tion use ten-point changes, and income uses a doubling of median household income. Each table entry is the change in that context gradient per additional candidate after five. The joint model includes the other racial shares and resources throughout. These tests therefore differ from separate one-moderator models and do not support selecting only the strongest demographic interaction.

2.2 Candidate Categories and Sensitivity

Candidate categories draw primarily on RepresentWomen’s council research workbooks and the candidate tables in the New York City Districting Commission’s voting-rights report (Handley, 2022; RepresentWomen, 2021, 2023, 2025). Additional public reports and biographies provide documented supplements, with the source, location, person link, and original category retained in the coding file. Names or images alone are not used to infer race, and a source classification is not relabeled as independently verified self-identification.

The crosswalk groups Black and Black/Haitian labels as Black, Asian and South Asian labels as Asian, and explicit Hispanic, Afro-Latino, Afro-Latina, and Hispanic-Asian labels as Hispanic. The original mixed labels remain recorded. Broad AAPI and people-of-color labels, conflicting classifications, and other unresolved identities do not enter the strict four-category analysis. A field is retained only if every named candidate can be classified under these rules; an unknown candidate never becomes a zero count for a category.

The own-category availability model adds

$$\sum_{g \in G} \rho_g S_{peg} A_{eg}, \quad A_{eg} = \min\{K_{eg}, \min(K_e, C_e)\}, \quad (\text{A.1})$$

where S_{peg} is precinct share in ten-point units and K_{eg} counts candidates in category g . Capping availability recognizes that candidates beyond the completion requirement cannot add another required rank. The more flexible model replaces these terms with all four precinct shares interacted with White, Black, and Asian candidate shares, leaving Hispanic candidate share as the reference composition. It allows cross-category associations rather than assuming that only same-category

supply is relevant.

Neither model treats its change in the White-share contrast as causal mediation. Candidate composition may accompany differences in ideology, familiarity, endorsements, and campaign resources. Before-and-after estimates nevertheless use the same sample, so their small difference is not generated by dropping additional observations. The paired deletion check removes the same district from both models while holding the prediction population fixed. The White-share contrast remains negative in every district-deletion fit.

Excluding fields involving provisional name aliases leaves 2,973 cells in 34 elections and 27 districts. The average White-share contrast is -1.64 before own-category adjustment and -1.63 after it, with both district CR2 intervals below zero. These checks preserve the aggregate finding without identifying the demographic composition of turnout or individual coethnic voting.

2.3 Comparison with the Earlier Joint Model

This comparison documents how demographic-interaction inference depends on the model specification and uncertainty calculation. The mixed-logit results are retained as a sensitivity audit, while the manuscript reports the current election-level and exploratory precinct-context models.

Table A.5 compares the earlier grouped-binomial mixed logit with estimates on a fixed common sample. The full diagnostic sequence first reproduces the submitted-era model, retains the documented District 16 candidate-field correction, and restricts the sample to the same 6,387 precinct-election cells. It then changes the random-effects structure, fixed intercepts, model scale, weights, outcome construction, racial measures, income transformation, and demographic interactions separately. The outcome comparison reapplies the old recorded-position rule to the rebuilt exports before excluding overvote marks, counting distinct named candidates, or changing to the named-engagement denominator. Separate one-change comparisons start again from the original common-sample specification, since a sequential comparison need not give a unique explanation independent of its ordering.

The earlier joint mixed-logit table used model-based Wald standard errors from a model

with district random effects. Those random effects did not supply a district-clustered robust covariance calculation. Table A.5 reports only district-based uncertainty: a jackknife that leaves both common-sample point estimates unchanged and uses their variation across 48 district-deletion refits, with a t_{47} reference distribution. The resulting raw p -values are 0.374 for White share and 0.572 for education, and neither passes Holm. Every deletion fit converges. Adding demographic pre-cap interactions also leaves both interactions uncertain under the district jackknife. The complete diagnostic record preserves the earlier model-based calculations for auditing, rather than presenting them as the paper’s inferential results.

This comparison concerns the earlier demographic interactions. The current main election-level estimates already use district-clustered CR2 uncertainty and retain a positive slope change at five; the narrow four-to-six-candidate comparison remains less precise. The exploratory race and education findings concern average contextual completion differences, not demonstrated effects of the cap. Non-rejection of their cap interactions does not establish that the underlying differences are zero.

2.4 Matching and Its Support Limits

The matching analysis examines whether candidate-identity comparisons retain sufficient support after balancing precinct composition and electoral context. These diagnostics supplement the regression adjustments without establishing an individual coethnic mechanism.

Table A.6 summarizes outcome-blind matching within the same year, candidate count, and rank limit. Pairs must involve different districts and different White-candidate counts, with no precinct reused within a specification. The tight resource limits are five points for White share and BA+ share, three points for age, and an income ratio no greater than 1.25. The all-race design applies the five-point limit to Black, Hispanic, and Asian shares as well and requires identical incumbency status. Mean absolute standardized differences must not exceed 0.10 within each year-by-field-size stratum.

Matching on White share and resources yields 129 pairs in 20 elections. The adjusted

gap between more- and fewer-White-candidate fields increases by 1.12 completion points per ten-point increase in the pair’s White-share midpoint, with an approximate interval $[0.03, 2.21]$. The gradient is suggestive, but the matched sample retains large absolute standardized differences in Black, Hispanic, and Asian shares, approximately 0.51, 0.40, and 1.16. Regression adjustment for those residual differences cannot substitute for strong overlap.

Requiring all-race balance leaves seven pairs in three 2023 three-candidate elections, comparing District 19 with Districts 23 and 43. Exact incumbency does not further reduce this sample. Only two pair midpoints exceed 50% White, and the average margin difference between paired elections remains about 29 points. Broader tolerances recover 80 pairs in eight elections, all with three or six candidates, but the average absolute within-pair margin difference remains about 20 points. A five-point margin restriction and mean-balance requirement leave no feasible pairs under either set of demographic tolerances.

Matched-model uncertainty permits comparisons to share either district through a positive-semidefinite district-endpoint sandwich. The 129-pair interval is approximate and unadjusted for the exploratory specification comparisons. Intervals for the tighter and broader all-race designs are withheld because only three and eight district endpoints remain. The numerical support guard requires at least ten district endpoints and four effective supporting districts; exceeding it is not a guarantee of reliable inference. These diagnostics support treating the candidate-identity mechanism as unresolved.

3 Minneapolis and San Francisco

3.1 What the Cross-City Figures Describe

Table A.7 records the case coverage. Minneapolis includes 45 regular council elections in 2009, 2013, 2017, 2021, and 2023 under a three-rank limit. San Francisco includes 12 district-supervisor elections in 2020, 2022, and 2024 under a ten-rank limit. Two-candidate contests and the 2020 Minneapolis special election are excluded. Candidate fields refer to distinct named candidates

observed in the export, with source-specific labels and write-ins checked in the common data construction.

Both figures use the same distinct-name completion indicator and engaged-ballot denominator as New York City. Each point represents an election, and a black horizontal tick shows the equal-election mean for a candidate count. Counts identify where a mean depends on only one or two elections. The figures contain no fitted regression, parametric confidence ribbon, or extrapolated no-cap line. Their purpose is to expose the observed support and completion requirement in each setting.

Figure A.1 shows that Minneapolis completion increases from 27.5% at three candidates to 51.6% at six, before falling to 43.4% and 36.3% at seven and eight. Each of the latter counts contains one election. In San Francisco, the single three-candidate election has 66.5% completion, while the mean at four candidates is 39.5%. The means at five, six, and seven are 14.8%, 15.3%, and 23.9%, so a description of continuously falling completion would be inaccurate.

These summaries pool different election years and districts. For example, San Francisco's five- and six-candidate observations occur in 2024, while seven-candidate fields occur in both 2020 and 2024. A straight fitted slope can hide this calendar composition and produce an apparently smooth pattern unsupported at some counts. It also adds dependence on the chosen controls, reference values, functional form, and uncertainty calculation. The raw figure keeps those features visible rather than treating a fitted curve as independent confirmation of the New York City threshold.

Figure A.2 restricts New York City and Minneapolis to their common three-to-five-candidate range. The required number of named rankings increases across this range in New York City and stays at three in Minneapolis. The comparison illustrates why the same candidate count can imply different completion benchmarks. Differences in primaries versus general elections, electorates, campaign information, candidate selection, and ballot administration remain bundled with the city comparison, precluding a causal interpretation of its contrasting means.

3.2 Source and Replication Boundaries

Cross-city coverage is restricted to council and supervisorial elections from the years listed in Table A.7. The six 2024 San Francisco exports are identified as district-supervisor contests using the source names and contest roster, with their record counts reconciled to the source inventory. Named-candidate counts exclude generic write-ins and special labels, placing San Francisco’s 2020 District 11 field at three candidates. An occupied-position definition would answer a different completion question if these labels entered the benchmark.

The complete replication record retains original file identifiers, checksums, eligibility rules, source labels, and selected aggregate inputs. Independent Python calculations reproduce election coefficients, raw summaries, candidate-count support, education contrasts, and table quantities. The existing audit independently reproduces the White-share models and matched calculations. CR2 intervals use the documented R implementation, with additional exported estimating objects available for examination; computational agreement does not validate candidate identities, geographic continuity, or a causal interpretation.

References

- Handley, L. (2022). *Report to the New York City Districting Commission: Analysis of voting patterns by race and assessment of compliance of proposed New York City Council district plan with the Voting Rights Act*. New York City Districting Commission. <https://www.nyc.gov/assets/districting/downloads/pdf/RBV-Report.pdf>
- Holm, S. (1979). A simple sequentially rejective multiple test procedure. *Scandinavian Journal of Statistics*, 6(2), 65–70. <https://www.jstor.org/stable/4615733>
- Pustejovsky, J. E., & Tipton, E. (2018). Small-sample methods for cluster-robust variance estimation and hypothesis testing in fixed effects models. *Journal of Business & Economic Statistics*, 36(4), 672–683. <https://doi.org/10.1080/07350015.2016.1247004>
- RepresentWomen. (2021). *New York City Council candidate research workbook, 2021*. Research data. <https://docs.google.com/spreadsheets/d/19j8B5KCsar6jW3FqAQ27Q9o5Tok4v9qb5HzvmUspxCE/edit>
- RepresentWomen. (2023). *New York City Council candidate research workbook, 2023*. Research data. <https://docs.google.com/spreadsheets/d/1mNjHWCjQqTbj3ixWOfMnexJJJ4FQh-3cSZZSxhpcKrQ/edit>
- RepresentWomen. (2025). *New York City Council candidate research workbook, 2025*. Research data. https://docs.google.com/spreadsheets/d/1ZMsL_53kcBsfGEfu5YsDwSz3ulFykiHd2shbyJatMkE/edit

Table A.1: Observed Completion and Candidate-Field Support in New York City

Candidates	Elections	Engaged ballots	Election mean (%)	Pooled rate (%)
3	15	208,175	40.4	35.7
4	15	194,242	31.5	28.5
5	6	91,944	22.1	20.6
6	10	171,004	24.4	22.8
7	6	139,928	26.6	25.8
8	8	127,972	30.2	28.8
9	3	51,610	31.8	30.8
11	1	22,361	34.0	34.0
12	2	41,895	41.0	40.8
13	1	25,251	41.6	41.6
15	1	17,467	47.6	47.6

Election means weight elections equally; pooled rates weight engaged ballots.

Table A.2: Election-Level Estimates Across Alternative Specifications

	Election Controls	Also Margin	District Labels
Below-cap slope	-7.25*** [-10.53, -3.98]	-7.68*** [-10.92, -4.45]	-6.68** [-10.97, -2.40]
Slope change	9.66*** [5.99, 13.33]	9.92*** [6.23, 13.60]	8.26** [3.04, 13.47]
Above-cap slope	2.41*** [1.69, 3.12]	2.23*** [1.42, 3.04]	1.57* [0.24, 2.90]
Incumbent present	-0.66 [-3.90, 2.59]	0.16 [-3.28, 3.60]	-1.91 [-5.76, 1.93]
First-choice margin (10 pp)		-0.69* [-1.32, -0.07]	
Elections	68	68	38
District clusters	48	48	18

District CR2 95% intervals; district labels span redistricting. * $p < .05$; ** $p < .01$; *** $p < .001$.

Table A.3: Local Candidate-Field Windows Test the Same Slope Change

	4–6 Candidates	4–7 Candidates	3–8 Candidates
Below-cap slope	-6.56*	-6.56*	-7.23***
	[-12.30, -0.82]	[-12.01, -1.11]	[-10.78, -3.67]
Slope change	8.46	8.66*	9.61***
	[-0.71, 17.62]	[1.24, 16.08]	[4.96, 14.27]
Above-cap slope	1.90	2.10	2.38**
	[-2.68, 6.48]	[-0.41, 4.60]	[0.71, 4.06]
Elections	31	37	60
District clusters	27	31	45

Year and incumbency controls; district CR2 95% confidence intervals. * $p < .05$; ** $p < .01$; *** $p < .001$.

Table A.4: Changes in Demographic Gradients at Five Candidates Are Uncertain

	All Fields	Common Fields	Own-Group	Full Composition
White share	0.04 [-0.82, 0.89]	-0.22 [-1.11, 0.66]	-0.44 [-1.32, 0.44]	-0.85 [-2.01, 0.31]
BA+ share	0.21 [-0.35, 0.76]	0.13 [-0.48, 0.74]	0.08 [-0.56, 0.72]	0.08 [-0.57, 0.73]
Age 65+ share	0.13 [-0.66, 0.92]	0.10 [-0.78, 0.97]	0.03 [-0.83, 0.89]	0.03 [-0.83, 0.88]
Median income	0.63 [-0.21, 1.46]	0.81 [-0.19, 1.80]	0.88 [-0.13, 1.89]	0.89 [-0.16, 1.94]
Naturalized share	-0.44 [-1.33, 0.44]	-0.69 [-1.78, 0.41]	-0.69 [-1.80, 0.42]	-0.79 [-1.75, 0.16]
Elections	68	39	39	39
District clusters	48	31	31	31

Pointwise district CR2 95% intervals; stars use Holm adjustment across five tests per model. * $p < .05$; ** $p < .01$; *** $p < .001$.

Table A.5: Earlier Demographic Interactions with District-Based Uncertainty

	Original Common Model	Also Pre-Cap Interactions
White share	0.0039 [-0.0048, 0.0126]	0.0026 [-0.0319, 0.0371]
BA+ share	0.0017 [-0.0043, 0.0077]	0.0107 [-0.0201, 0.0415]
Precinct-election cells	6387	6387
Elections	68	68
District clusters	48	48

Log-odds interactions per ten-point context shift; district jackknife 95% intervals (47 df). No interaction survives five-test Holm adjustment.

Table A.6: Stronger Matching Leaves Few Comparable Elections

Rule	Pairs	Elections	Districts	White midpoint $\geq 50\%$
White share and resources	129	20	18	16
All racial shares and incumbency	7	3	3	2
Broader demographic tolerances	80	8	8	14
Tight matching plus margin	0	0	0	0
Broader matching plus margin	0	0	0	0


Resources include BA+, income and age; margin sensitivities require margins within five points.

Table A.7: Different Cities Provide Different Ranking Contexts

City	Years	Limit	Elections	Districts	Engaged
New York City	2021, 2023, 2025	5	68	48	1,091,849
Minneapolis	2009, 2013, 2017, 2021, 2023	3	45	13	321,682
San Francisco	2020, 2022, 2024	10	12	7	386,114


Council or supervisorial elections with at least three named candidates; cities are descriptive comparisons.

Figure A.1: Raw Cross-City Completion Patterns Have Uneven Support



Points represent elections; ticks show equal-election means and labels show election counts.

Figure A.2: The Same Candidate Counts Imply Different Completion Requirements



New York City requires three, four, or five rankings; Minneapolis requires three throughout.