The Composition of Power: Gender Penalties in Close Elections

Neil Duzett*
Texas A&M University

Abstract: Using national data on local government elections in the U.S., I show that in close elections between candidates of different genders, defeat consistently goes to the candidate who does not match current officeholders' most common gender. These gender-minority candidates are 1.9 times as likely to lose as win near the victory cutoff. This penalty exists for both women and men, is present across multiple government offices, and results in long run elected body compositions being 7% less female. I provide evidence that gender disparities in campaign donations contribute to this gender-minority penalty in close elections.

* Email: nduzett@tamu.edu

I. Introduction

Gender disparities are particularly consequential in leadership positions, which shape both contemporary policies and the pipeline for future leaders (Beaman et al. 2009). Despite shrinking gender divides in education and labor force participation, gender gaps persist in many key leadership and decision-making roles. Women hold just 25% of senior management positions in global top companies, 32.8% of U.S. college and university presidencies, and 33.2% of U.S. Federal judgeships. In terms of shaping policy, government holds perhaps the most impactful positions of power, and is no stranger to gender disparities; gender representation in government has long been unequal in many parts of the world. For example, female representation in U.S. congress has risen to 28% in 2025—a marked increase over 10% in 1992's "Year of the Woman," yet still far from parity.

What explains the gender disparity in leadership? A sizable body of research has studied the barriers to gender equity in leadership positions: women avoid competition (Niederle and Vesterlund 2007), receive less promotion-relevant recognition for group work (Sarsons 2017), and have fewer social interactions with their managers that can lead to higher probability of promotion (Cullen and Perez-Truglia 2023). Women running for political office face additional barriers, often running against entrenched incumbent men (Lee 2008) and receiving less or different media coverage during elections (Kahn 1994). Despite work to overcome these barriers, such as gender-blind hiring practices (Goldin and Rouse 2000) and gender quotas in corporate boards and political candidacy (Bertrand et al. 2019, Matsa and Miller 2025, Chattopadhyay and Duflo 2004, Besley et al. 2017, Bagues and Campa 2021)—often seeing only limited success—progress toward parity is slow.

In this paper, I document a new fact regarding the barriers to power faced by minority genders. Specifically, I show that in close elections between candidates of different genders, outcomes systematically favor the candidate whose gender matches that of the current officeholders. This fact suggests that the gender composition of those in power matters for determining who holds power next, even when leadership has no formal channel for picking their successors.

I examine this idea in the local government context, investigating how the gender composition of U.S. city councils, county legislatures, and school boards influences the

likelihood of different-gender candidates winning a close election to those same offices. These local governments shape critical everyday policies such as zoning, policing, education, housing, and health. Furthermore, local elected positions are a key pipeline to higher state and national offices for politicians, making such contests a particularly impactful setting for studying gender barriers to power.[1] Additionally, these legislative bodies often operate under majority voting, making even small gender imbalances potentially pivotal. Using the largest existing collection of U.S. local government elections, I examine mixed-gender elections where one candidate matches the majority gender of the current local institution (i.e. city council, county legislature, or school board), and the opposing candidate does not. Majority gender status is a natural measure of advantage in the local government context, as these legislative bodies normally operate under majority rule. To isolate the role of gender in determining electoral outcomes, I focus on close elections between these candidates, which is an ideal circumstance for examining gender barriers to leadership positions. Candidates close in vote shares are more comparable on the characteristics that determine victory, minimizing gender selection into winning. In these close elections more than any other, gender parity in victory should be expected. I use a formal density test designed by McCrary (2008) to quantify differences in the probability of gender-minority victory and defeat in these close elections.

I find that the gender composition of current officeholders plays a key role in deciding close elections, and gender-minority candidates suffer a victory penalty: at the victory cutoff, gender-minority candidates are 1.9 times as likely to lose a close election as win. This result is robust to bandwidth, bin, and density test alternatives. Further, the gender-minority penalty in close elections is present separately in city council, county legislature, and school board elections, demonstrating that this imbalance is not unique to one context. In addition to its prevalence across local offices, it is also self-perpetuating over time. A female candidate loss results in a long run council composition that is 7% less female than if there were no gender penalty. This creates cascading effects that further penalize future gender-minority candidates, since more imbalanced compositions experience stronger gender-minority penalties. Finally, the penalty is not unique to one gender. Both female candidates and male candidates face penalties in

[1] For example, Chattopadhyay and Duflo (2004) show that local politician gender has important impacts on policy: village councils in India with greater female representation invest in female-friendly public goods such as water access.

the probability of close election victories when they are the gender-minority. I show that this gender heterogeneity can obscure the gender-minority penalty when the council composition is not taken into account.

What may be contributing to the gender-minority penalty in close elections? I find evidence that gender-minority candidates receive fewer campaign contributions than their gender-majority opponents, likely decreasing the probability of a gender-minority victory. Using a difference-in-differences design, I show that gender-minority candidates on average receive $17,553 less in contributions than their gender-majority opponents—half the contributions received by the average candidate. Both gender-minority women and gender-minority men suffer penalties in contributions when facing opposite gender opponents. This descriptive evidence indicates donors provide less financial backing to candidates who do not match the majority gender of the elected body. One explanation consistent with this result is that elected individuals are swaying financial support (e.g. fundraising networks, endorsements, advice) toward individuals whose gender matches their own—particularly in close races where small investments of additional resources could decide the outcome—contributing to systematic losses for gender-minority candidates.

II. Data and Empirical Approach

A. *Local government election data*

I exploit a rich, national dataset of local government elections called the American Local Government Elections Database (Benedictis-Kessner et al. 2023). This dataset features election information for over 57,000 elections between 1989-2021 in U.S. cities and counties over 50,000 population in the 2020 Census. These data provide some demographic information for candidates including race, gender, and major party affiliation, and cover elections to many local offices, including city councils, county legislatures, and school boards. While previous work on local government elections in the U.S. is usually restricted to a handful of elections in one state and restricted to only a few years of data, the national nature of this dataset allows me to examine how the role of gender in positions of political power fluctuates across many different localities in the U.S.

To construct my estimation sample, I first restrict attention to elections for city council, county legislature, or school board. I focus on these multimember offices to provide a convenient setting for establishing which gender is the political minority. I focus on elections where the gender of both the winner and loser are known and are different.[2] I then attach gender composition data for each institution at the time of the election—also provided by Benedictis-Kessner et al. (2023)—and determine which candidate is the gender-minority.[3] I use the gender-majority status of the council to assess disadvantage for the opposite gender—as opposed to another compositional threshold—as it offers a natural and intuitive measure of which group holds greater institutional power, particularly in a legislative body governed by majority rule.[4] I will refer to candidates whose gender does not match the majority gender of the multimember office they are seeking as *gender-minority candidates*, and candidates who do match as *gender-majority candidates*. This process results in 6,564 elections between gender-minority and gender-majority candidates.

Table 1 contains summary statistics for the localities (cities, counties, school districts), offices, and candidates involved in local elections. I observe 352 cities, 421 counties, and 73 school districts in which there has been at least one mixed-gender election between 1989-2021. On average, city councils and county legislatures in these localities are majority male, while school boards are more often majority female.[5] Of the 33,464 candidates in all elections in the data (not only mixed gender), 29% are female. However, of the 6,564 elections involving gender-minority and gender-majority candidates, women make up 82.8% of gender-minority candidates, showing the prevalence of male-majority political candidacy.

[2] In elections where there are multiple winners and losers, I only compare the winner with the lowest vote share and the loser with the highest vote share.

[3] Since nearly all elected bodies have an odd number of seats, there is nearly always a majority gender. However, in some cases in the composition data, candidate gender is not known or ambiguous. When this occurs, I only keep observations for which I can confidently calculate the majority gender (e.g. a council has 9 seats and I know at least 5 are filled by men). I omit councils where male and female representation are precisely equal – this is an exceptionally rare occurrence, as most councils have an odd number of members.

[4] I later examine heterogeneity in the minority gender penalty for different gender compositions and find that the penalty is larger with more extreme compositions and more muted around balanced compositions, suggesting that majority gender status is a reasonable threshold for determining gender disadvantage.

[5] The sample of school boards is relatively small, as the data collection process by Benedictis-Kessner et. al (2023) was primarily focused on other offices.

B. *Detecting potential gender penalties in close elections*

To formally test for a gender-minority penalty when running against a gender-majority candidate, I use the McCrary (2008) density test. Conceptually, the test is designed to detect statistically significant shifts in density in a continuous variable near a treatment cutoff. Implementing this test proceeds in three steps. First, I construct each gender-minority candidate's margin of victory by taking the difference between their vote share and their gender-majority opponent's vote share. This results in a continuous running variable ranging from -1 to 1, for which positive values indicate gender-minority candidate victory and negative values indicate gender-minority candidate loss. Next, I pool observations into bins of gender-minority candidate margin of victory and plot a histogram. Finally, I separately estimate a 4^{th} order polynomial on each side of the victory cutoff.[6] The parameter of interest is the difference in log height between these polynomials at the victory cutoff (where margin of victory is nearly zero), representing shifts in the relative probability of gender-minority candidate victory in the closest elections. If gender-minority candidates win close elections at rates equal to gender-majority candidates, this test would fail to reject the null hypothesis of no shift in density at the victory cutoff.

III. Evidence of a Gender-Minority Penalty

A. *Gender-minority candidates are less likely to win close elections*

Figure 1 displays the result of a density test on gender-minority candidate margins of victory when their opponent is the majority gender. The figure provides strong evidence of a discontinuity in density at the victory cutoff, rejecting the null hypothesis of gender-minority and gender-majority candidates winning elections at equal rates. At the victory cutoff, there is a log difference in density of -0.62 (0.06), indicating a 46.2% drop in the density of gender-minority candidates from candidates who barely lose to those who barely win. This drop in density translates to gender-minority candidates being 1.9 times more likely to lose than to win in close elections.

[6] This estimation gives greater weight to bins nearer the cutoff, effectively estimating within a bandwidth. While I use an automated bandwidth selection procedure and report results primarily for that bandwidth, I later show that the main estimates are not sensitive to bandwidth choice.

The estimated gender-minority penalty is robust to choice of bin size, bandwidth, and density test alternatives. First, McCrary (2008) shows that choice of running variable bin size does not change the estimate of the parameter of interest and verifies this with simulations. Second, Figure A.1 shows that the gender-minority penalty estimate does not change with choice of bandwidth up to twice and half the bandwidth given by the automated bandwidth selection procedure. Third, the gender-minority penalty is also robust to choice of density test. Cattaneo et al. (2020) introduce a density estimator similar to McCrary (2008) that does not rely on pre-binning the data and performs better under some additional assumptions. Using this method in Figure A.2 shows a significant shift in density at the victory cutoff for gender-minority candidates, identical to the initial result.

Beyond technical robustness, I show that the gender-minority penalty is not unique to a particular local office. While the main result in Figure 1 pools city council, county legislature, and school board elections to estimate the overall gender-minority penalty, Figure 5 panels A, B, and C present density tests separately for elections to each of those offices. The figure shows that the scope of the gender-minority penalty is not isolated to one kind of public office: city councils, county legislatures, and school board elections all exhibit lower likelihoods of gender-minority candidate victory. The scope of this gender-minority penalty suggests that gender is a widespread factor in determining the compositions of governments and the identity of those who hold power.

The gender-minority penalty in close elections has direct implications for local government policy as well as the compositions of future leadership. Governments that are less female due to this penalty may employ fewer female-friendly policies (Chattopadhyay and Duflo 2004) and remaining female policymakers may participate less in the legislative process when they have fewer female colleagues (Rebolledo et al. 2024). In addition to the initial minority loss, the gender-minority penalty implies long run consequences for gender representation in government as well. The penalty is self-perpetuating: gender-minority losses shift the council composition even closer toward the majority gender, further disadvantaging future gender-minority candidates. Assuming all mixed-gender elections are close, a single female loss in a gender-even composition results in future councils that are 7% less female than if there were no

gender-minority penalty.[7] The long run drop in female compositions could be even more drastic than estimated here, as previous work has shown that female victories can inspire other women to run in the future (Beaman et al. 2009), implying that additional female losses remove potential female candidates from the election pool. Further, the cascading impact of gender-minority losses is potentially augmented if more imbalanced compositions experience greater penalties.

B. *Heterogeneity in the gender-minority penalty*

Institutional gender composition

As the gender-minority status of the candidate is shown to play a pivotal role in victory, a natural question is whether more gender imbalanced institutions are associated with stronger gender penalties. I examine how the gender-minority penalty varies in strength by the composition of the office being sought. Figure 4 plots the penalty estimates from separate density tests by each percentage of the city council, county legislature, or school board that is female, binned in 10% intervals. For elections to the left of 50% female, all gender-minority candidates are female, while those to the right are all male. In institutions with no female members, the penalty against female candidate victory is exceptionally large. The penalty becomes significantly more muted as soon as the institution's gender composition reaches 10% female, and begins to favor female candidates upon reaching 40% to 50%. The advantage for female candidates continues to grow with more female-dominant compositions but never reaches the magnitude of the penalty at 0% to 10% female. As so few institutions feature gender compositions above 80% female, it is not possible to reliably estimate a density test for mixed-gender elections to those institutions.

These results show that the gender-minority penalty in close elections is stronger when gender compositions are more imbalanced. In places where few or no women hold office, women face a near insurmountable probability of close election loss. The same is largely true for gender-minority men, but the disparities are more muted and localities where male candidates are

[7] This is estimated using 5000 simulations of two consecutive election cycles, assuming half the council is up for reelection at a time and there is a 41% chance of an election to be mixed-gender (based on the in-sample share of mixed-gender elections). Simulations with a minority gender penalty of 34% (as estimated in Figure 1) result in an average female composition of 39% after two full cycles, while simulations without the penalty result in an average female composition of 46% after two full cycles.

the gender-minority are much rarer. Extreme gender compositions in elected government are more resistant to disruption, which speaks to the entrenched nature of the characteristics of those who hold positions of power.

Candidate gender

Given the gendered nature of the disparity and the series of established barriers to political office for women, I examine heterogeneity in the gender-minority penalty by candidate gender. Figure 2 panels A and B display separate density tests for female and male gender-minority candidates. These tests reveal two points of interest. First, the penalty is not unique to one gender: both gender-minority women and gender-minority men face a disparity in victories when running against the opposite gender. Second, while women are more commonly disadvantaged, men face a larger penalty when they are the gender-minority: -0.43 (0.06) for women and -1.55 (0.16) for men, which correspond to those candidates being 1.54 and 4.72 times as likely to lose as to win, respectively.

The gender heterogeneity in the gender-minority penalty contrasts with the prevailing view in the election literature that women and men win elections at similar rates (Seltzer et al. 1997, Lawless 2015). Rather, the perceived gender equity in victory rates could be a function of hidden heterogeneity. While women face significant barriers and are more commonly the gender-minority, both gender-minority women and gender-minority men face lower probability of victory in close elections. Without accounting for gender-minority status, this heterogeneity is obscured.[8] Figure A4 shows the result of a density test using the more common approach of calculating margins of victory for female candidates versus male candidates, effectively pooling gender-majority and gender-minority female candidates. Indeed, no shift in density is detected—the opposing shifts in density between gender-majority and gender-minority candidates effectively cancel out, showing how the gender-minority penalty is obscured when the approach fails to account for institutional gender composition.

The existence of the gender-minority penalty in close elections for both women and men also challenges an important implication of the common regression discontinuity (RD)

[8] Figure A3 shows histograms for female candidate margin of victory, further displaying how the heterogeneity in the gender-minority penalty obscures the penalty itself.

identifying assumption: candidates just above and below the victory cutoff should be similar in all relevant respects except for the election outcome. If the probability of victory differs systematically by gender-minority candidate status in close races, vote share-based RD estimates of the causal effects of candidate characteristics may be biased.[9]

IV. Why Is There a Gender-Minority Penalty in Close Elections?

Why do gender-minority candidates lose close elections at disproportionate rates? One plausible explanation is differential access to campaign resources.[10] Finance is a vital feature of all political campaigns. Candidate spending, particularly in advertising and events, has been shown to increase voter support and the probability of candidate victory (Stratmann 2009, Schuster 2020, Le et al. 2024). The effectiveness of additional spending increases in close elections, as does the amount of contributions (Gerber 2004, Bouton et al. 2018). It is possible that the gender of those on an elected body influences the flow of campaign resources. This would be consistent with recent work by Bouton et al. 2022, which shows that contributors give more to candidates who match their own gender. As elected officials presumably have some influence over the flow of resources to candidates (e.g., via connections with donors including political action committees (PACs), political parties, and fundraising networks), their gender preferences could influence the outcomes of elections.

To examine the role campaign finance plays in determining the gender-minority penalty in close elections, I attach local campaign contribution data to a subsample of the election data. These contribution data are collected from individual campaign finance reports and collated by state agencies. The Accountability Project, a nonprofit initiative for transparency in campaign finance, collects and standardizes much of the data across states, and offers a subset of the data available for download from their site. I restrict attention to city council elections for a clean comparison between elections, since contributor amounts and financial requirements for running

[9] McCrary (2008) suggests that density tests may fail to identify selection across the victory cutoff if such selection is nonmonotonic, i.e. observations sort in both directions across the cutoff, as this paper demonstrates is often the case in local mixed gender elections. See Marshall (2024) for an overview of challenges in identifying the causal effects of candidate characteristics using a vote share RD method.

[10] I briefly evaluate alternative explanations, such as incumbency, party, and voter preferences, at the end of this section and in the Appendix.

a campaign may vary widely between different local offices. The final merged sample includes 227 city council candidates with complete campaign contribution information.[11] I will refer to this sample as the *contribution sample*.

Using these data, I employ a difference-in-differences design to describe potential differences in contributions between gender-minority and gender-majority candidates while controlling for time and locality specific effects:

$$Contributions_i = \beta_0 + \beta_1\, gender\text{-}minority\ candidate_i + \beta_2\, mixed\text{-}gender\ election_i + \beta_3\, gender\text{-}minority\ candidate_i * mixed\text{-}gender\ election_i + \mu_t + \gamma_c + \varepsilon_i$$

$Contributions_i$ is the dollar total of contributions received by the city council candidate. Gender-minority candidate (GMC) = 1 if the candidate does not match the majority gender of the council they are contesting. Mixed-gender election (MGE) = 1 if the candidate's opponent is the opposite gender. μ_t and γ_c are year and city fixed effects respectively. The primary parameter of interest is the sum of β_1 and β_3, which represents the difference in contributions between gender-minority and gender-majority candidates in mixed-gender elections.

Table 2 presents the results of this approach. Column 1 reveals that gender-minority candidates suffer a net loss of $17,553 compared to their gender-majority opponents. This loss is roughly one half of the average contributions received by candidates in this sample. These results suggest a relationship between election victory, differences in contributions, and the existing council's influence over the flow of donations. Contributions are critical for campaign success, particularly in close elections when candidates must expend more resources to differentiate themselves from their opponent, and even small differences in this resource can tip the scales away from victory. As the council's gender composition defines which candidate is financially disadvantaged as well as which candidate is less likely to win a close election, the sum of the evidence is consistent with council members wielding influence over the flow of donations to favor candidates who match their own gender, resulting in disproportionate close election losses for gender-minority candidates. This suggests that those who hold positions of power, even when

[11] I discuss the drawbacks of this limited sample size near the end of this section.

they possess no formal channel for determining their successor, can still influence the characteristics of future leadership.

While the evidence supports the existence of a relationship between campaign finance and the gender-minority penalty, the approach in this paper is not without drawbacks. First, the sample of elections with contribution information is admittedly small, covering a fraction of the original sample, forfeiting some statistical power and rendering the approach more susceptible to potential selection. Second, the approach does not take advantage of election closeness, which would provide more explanatory power for the gender-minority penalty in close elections, but also further severely limit the sample size. Third, any interpretation of gaps in contributions between minority and majority candidates as affecting their probability of victory is confounded by the possibility that more popular candidates – who are already more likely to win – tend to attract greater contributions. However, I find that gender-minority candidates running against opposite gender opponents do not receive significantly higher or lower vote shares overall, in either the campaign contribution sample or the full sample, even while gender-minority candidates lose close elections more often (Table 1). In the face of these limitations, I interpret the results on campaign contributions to be suggestive of a core mechanism behind the gender-minority penalty in close elections.

Another possibility is that there is an alternative explanation for the gender-minority penalty in close elections. Several plausible explanations for the observed victory gap (e.g., incumbency, party affiliation, selective turnout, and corruption) do not receive empirical support. First, the gender-minority penalty might reflect incumbency advantage, since gender-minority candidates are more likely to be challengers. However, the gap persists in open-seat elections without incumbents, suggesting incumbency is not the primary driver (Figure A5). Second, if gender correlates with party, the effect might stem from partisan preferences. But the penalty remains even when both candidates share the same party affiliation, ruling out party composition as the explanation (Figure A6). Third, voter preferences for candidate gender could explain the pattern if such preferences are stable across time. Yet when institutions flip majority gender—e.g., from male-dominated to female-dominated—the direction of the penalty also reverses, suggesting that voter preferences do not drive the gap (Figure A7). Finally, I find no evidence of

ballot manipulation or selective mobilization in close elections involving gender-minority candidates (Table A1, Figure A8).[12]

V. Conclusion

This paper provides new evidence on the self-reinforcing nature of gender in positions of power. I show that gender disparities in political representation are reinforced by institutional composition: candidates whose gender does not match that of current officeholders are more likely to lose, even in elections decided by narrow margins. The result is not unique to one gender or one political office, and is associated with strong gender imbalances in campaign finance linked to institutional gender composition.

The scope of the gender-minority penalty has broad implications for understanding why gender parity in government remains elusive, despite gains in other domains. This penalty raises questions about the feasibility of gender parity in political office; even when major barriers such as incumbency advantage and lower candidacy rates are overcome, gender-minority candidates may still struggle to win office against otherwise comparable opponents. In particular, the existence of the penalty across several different local offices suggests it is not an isolated phenomenon. Rather, gender in government and other positions of power may self-perpetuate even when those in power have no formal channel for picking their successors.

[12] See Vogl (2014), who uses a similar approach to examine changes in turnout in Black-White U.S. mayoral elections.

REFERENCES

Bagues, M., & Campa, P. (2021). Can gender quotas in candidate lists empower women? Evidence from a regression discontinuity design. *Journal of Public Economics*, *194*, 104315.

Bagues, M., Sylos-Labini, M., & Zinovyeva, N. (2017). Does the gender composition of scientific committees matter?. *American Economic Review*, *107*(4), 1207-1238.

Barber, M., Butler, D. M., & Preece, J. (2016). Gender inequalities in campaign finance. *Quarterly Journal of Political Science*, *11*(2), 219-48.

Beaman, L., Chattopadhyay, R., Duflo, E., Pande, R., & Topalova, P. (2009). Powerful women: does exposure reduce bias?. *The Quarterly Journal of Economics*, *124*(4), 1497-1540.

de Benedictis-Kessner, J., Lee, D. D. I., Velez, Y. R., & Warshaw, C. (2023). American local government elections database. *Scientific Data*, *10*(1), 912.

Bertrand, M., Black, S. E., Jensen, S., & Lleras-Muney, A. (2019). Breaking the glass ceiling? The effect of board quotas on female labour market outcomes in Norway. *The Review of Economic Studies*, *86*(1), 191-239.

Besley, T., Folke, O., Persson, T., & Rickne, J. (2017). Gender quotas and the crisis of the mediocre man: Theory and evidence from Sweden. *American Economic Review*, *107*(8), 2204-2242.

Bouton, L., Castanheira, M., & Drazen, A. (2018). A theory of small campaign contributions (No. w24413). National Bureau of Economic Research.

Bouton, L., Cagé, J., Dewitte, E., & Pons, V. (2022). Small campaign donors (No. w30050). National Bureau of Economic Research.

Cattaneo, M. D., Jansson, M., & Ma, X. (2020). Simple local polynomial density estimators. *Journal of the American Statistical Association*, *115*(531), 1449-1455.

Chattopadhyay, R., & Duflo, E. (2004). Women as policy makers: Evidence from a randomized policy experiment in India. *Econometrica*, *72*(5), 1409-1443.

Cordero, A. A., Bandiera, A., Luna, E. A. G., & Meriläinen, J. (2024). When fortune favors women: Do marginal increases in female representation persist?. *Journal of Economic Behavior & Organization*, *226*, 106680.

Cullen, Z., & Perez-Truglia, R. (2023). The old boys' club: Schmoozing and the gender gap. *American Economic Review*, 113(7), 1703-1740.

Ferreira, F., & Gyourko, J. (2014). Does gender matter for political leadership? The case of US mayors. *Journal of Public Economics*, *112*, 24-39.

Ferrer, J., & Thompson, D. (2025). Does Leader Turnover Degrade Local Government Performance? Evidence from Local Election Officials. Working paper.

Gerber, A., Hoffman, M., Morgan, J., & Raymond, C. (2020). One in a million: Field experiments on perceived closeness of the election and voter turnout. *American Economic Journal: Applied Economics*, *12*(3), 287-325.

Goldin, C., & Rouse, C. (2000). Orchestrating impartiality: The impact of "blind" auditions on female musicians. *American Economic Review*, 90(4), 715-741.

Kahn, K. F. (1994). Does gender make a difference? An experimental examination of sex stereotypes and press patterns in statewide campaigns. *American Journal of Political Science*, 162-195.

Lawless, J. L. (2015). Female candidates and legislators. *Annual Review of Political Science*, *18*(1), 349-366.

Le, T., Onur, I., Sarwar, R., & Yalcin, E. (2024). Money in politics: How does it affect election outcomes?. *SAGE Open*, *14*(4), 21582440241279659.

Lee, D. S. (2008). Randomized experiments from non-random selection in US House elections. *Journal of Econometrics*, *142*(2), 675-697.

Marshall, J. (2024). Can close election regression discontinuity designs identify effects of winning politician characteristics?. *American Journal of Political Science*, *68*(2), 494-510.

Matsa, D. A., & Miller, A. R. (2025, May). Did Board Gender Quotas Break the Glass Ceiling in Europe?. In AEA Papers and Proceedings (Vol. 115, pp. 529-535). 2014 Broadway, Suite 305, Nashville, TN 37203: American Economic Association.

McCrary, J. (2008). Manipulation of the running variable in the regression discontinuity design: A density test. *Journal of Econometrics*, *142*(2), 698-714.

Niederle, M., & Vesterlund, L. (2007). Do women shy away from competition? Do men compete too much?. *The Quarterly Journal of Economics*, *122*(3), 1067-1101.

Rebolledo, E. B., Bruhn, J., Choon, T. H., & Weber, E. A. (2024). Gender Composition and Group Behavior: Evidence from US City Councils (No. w33223). National Bureau of Economic Research.

Sarsons, H. (2017). Recognition for group work: Gender differences in academia. *American Economic Review*, *107*(5), 141-145.

Schuster, S. S. (2020). Does campaign spending affect election outcomes? New evidence from transaction-level disbursement data. *The Journal of Politics*, *82*(4), 1502-1515.

Seltzer, R. A., Newman, J., & Leighton, M. V. (1997). *Sex as a political variable: Women as candidates and voters in US elections*. Lynne Rienner Publishers.

Stratmann, T. (2009). How prices matter in politics: The returns to campaign advertising. *Public Choice*, *140*, 357-377.

Sudulich, L., Trumm, S., & Makropoulos, I. (2024). Running uphill: A comparative analysis of the gender gap in campaign financing. *European Journal of Political Research*.

Vogl, T. S. (2014). Race and the politics of close elections. *Journal of Public Economics*, *109*, 101-113.

Table 1. Campaign Contributions in Elections by Minority Gender Candidate Status

	Contributions ($)			Vote share	
	(1)	(2)	(3)	(4)	(5)
	Contribution sample	Women	Men	Contribution sample	Full sample
Minority gender candidate (MGC)	39,458*	36,631	12,039	-0.069	-0.014*
	(22,646)	(30,169)	(13,345)	(0.053)	(0.008)
MGC * Mixed gender election	-57,011**	-36,753	-28,485*	0.009	-0.009
	(27,458)	(25,184)	(16,981)	(0.069)	(0.011)
Mean	37,126	34,664	38,269	0.524	0.501
N	227	72	155	227	8,396

Notes: Each column is difference-in-differences specification, all of which include city and year fixed effects. The specifications in columns 1-4 include only a subset of city council elections that match to campaign finance data provided by the Accountability Project. Columns 2 and 3 include only candidates who match the gender noted in the column. Column 5 is the full sample of all city council, county legislature, and school boards elections.

Figure 1. Density Test for Gender-Minority Candidate Victories in Mixed-Gender Elections

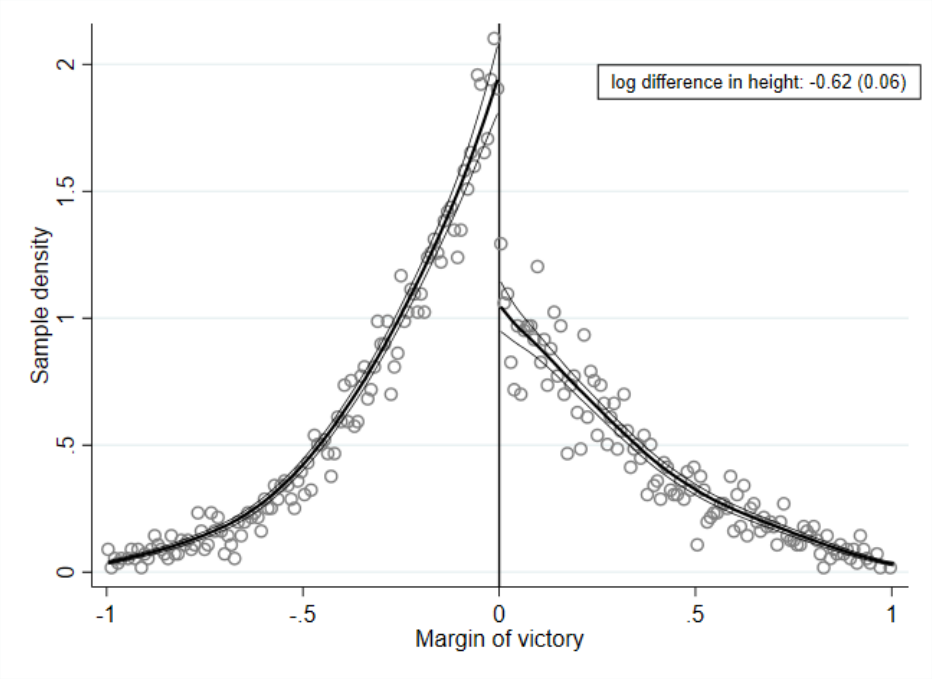

Note: N= 6564. This figure presents a density test for minority gender candidates' margin of victory, which is the difference between the minority gender candidate's vote share and their majority gender opponent's vote share. Vote shares are normalized between both candidates. The sample size will be smaller than the sample in Table 1, which also includes contests between two minority gender candidates.

Figure 2. Density Test for Gender-Minority Candidate Victories - Heterogeneity by Office Type

A. City Councils

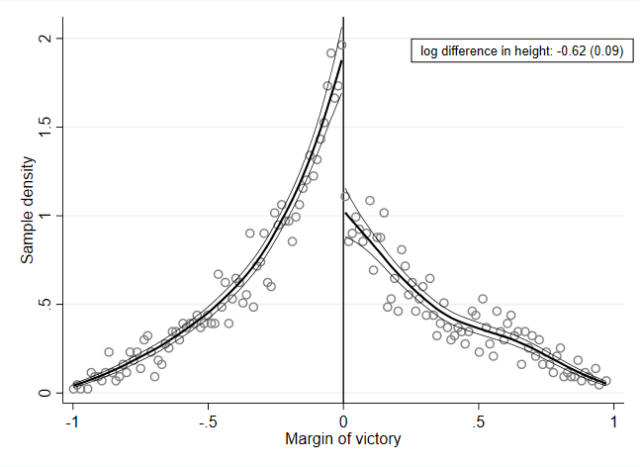

B. County Legislatures

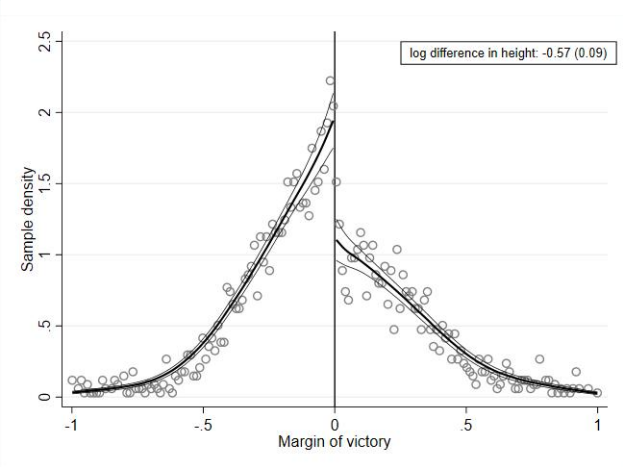

C. School Boards

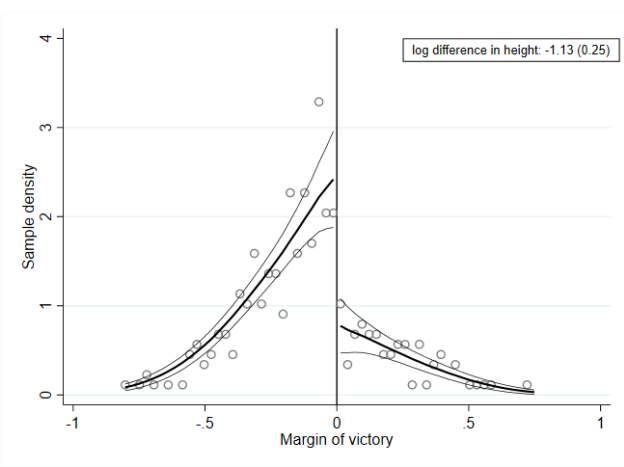

Note: Panel A: N = 3,319, Panel B: N = 2,921, Panel C: N = 324. This figure presents a density test for minority gender candidates' margin of victory separately by gender, which is the difference between the minority gender candidate's vote share and their majority gender opponent's vote share. Vote shares are normalized between both candidates.

Figure 3. Density Test for Gender-Minority Candidate Victories - Heterogeneity by Candidate Gender

a. Gender-Minority Female Candidates

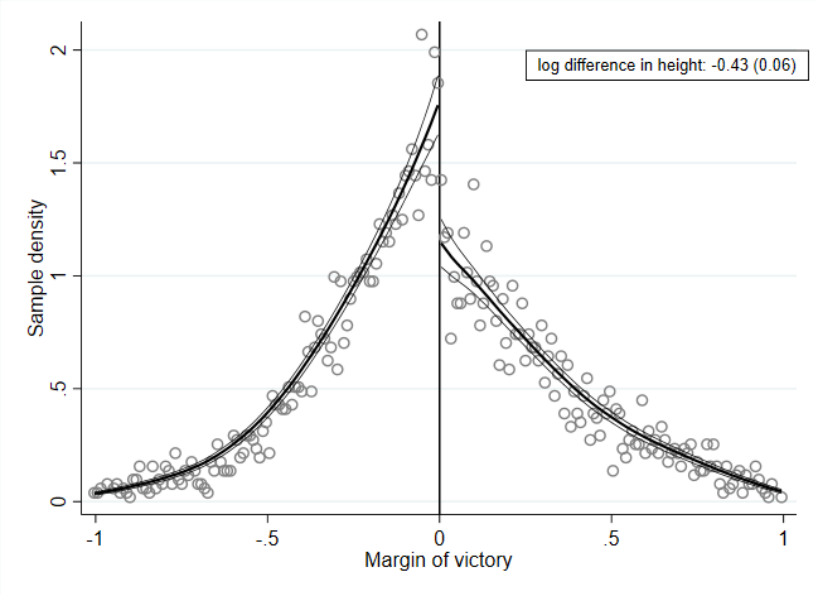

b. Gender-Minority Male Candidates

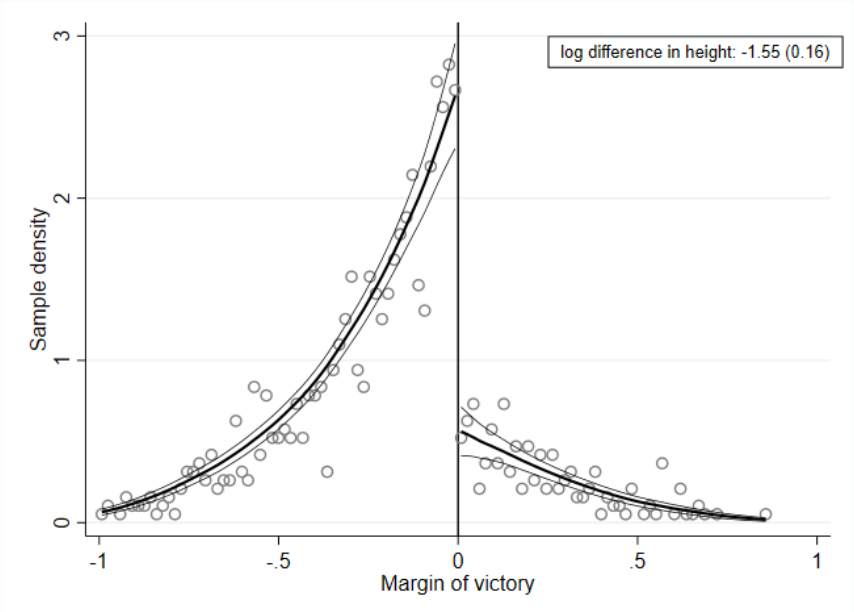

Note: Panel A: N = 5437, Panel B: N = 1127. This figure presents a density test for minority gender candidates' margin of victory separately by gender, which is the difference between the minority gender candidate's vote share and their majority gender opponent's vote share. Vote shares are normalized between both candidates.

Figure 4. Density Test for Gender-Minority Candidate Victories - Heterogeneity by Institutional Gender Composition

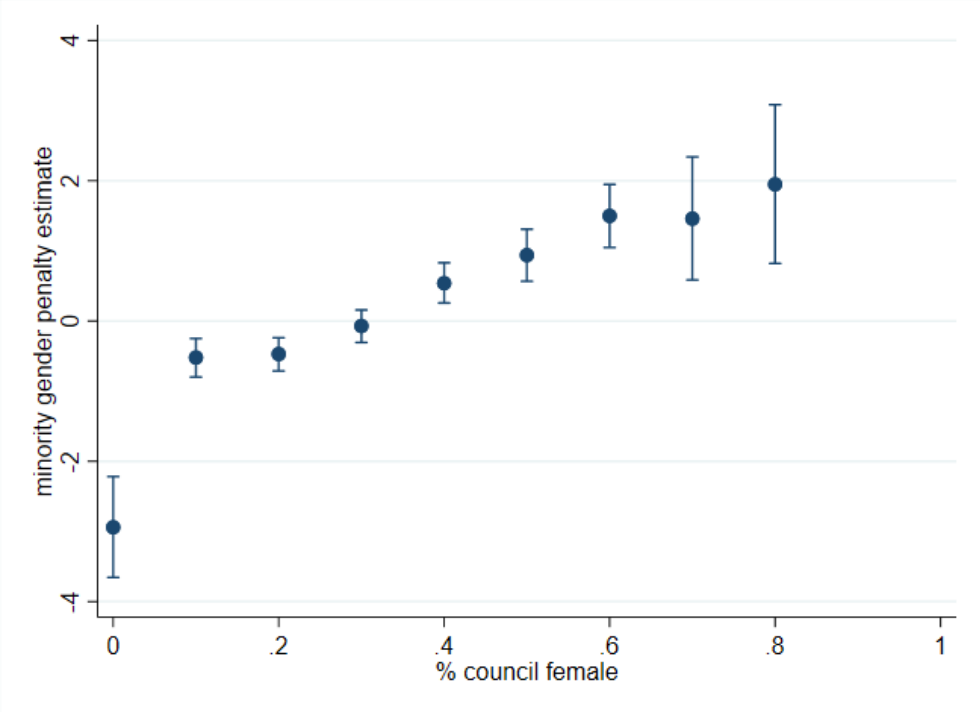

Notes: This figure plots density tests for institutions with different gender compositions. Institutions are binned in 10% intervals by what percentage of their members are female (e.g. the 10% bin includes all institutions that are between 10%-19% female, inclusive). The y-axis is the estimated gender penalty (or advantage) for female candidates running against a male opponent.

APPENDIX

Table A1. Summary Statistics for Election Data

Panel A: Localities	Cities	Counties	School districts
Population	594,451	-	-
	[1,543,800]		
N	352	421	73
Panel B: Institutions	City councils	County legislatures	School boards
Num male members	5.99	5.70	2.33
Num female members	2.73	1.84	3.28
Num seats	9.27	7.97	6.32
N	2,627	2,676	302
Panel C: Candidates	All	Majority gender	Minority gender
Male	23,643	20,979	2,664
Female	9,821	2,514	7,307
N	33,464		

Notes: Panel A includes the populations for all cities, counties and school districts that are represented by at least one election in the data. Panel B includes a breakdown of the gender composition of each institution (i.e. city council, county legislature, and school board) for which the majority gender can be determined. Each institution is considered a unique observation after an election. Panel C includes candidates involved in elections to an institution included in panel B, broken down by whether that candidate matches the majority gender of the institution they are contesting.

Table A2. TWFE estimates on turnout rate and margin of victory

	Outcome	
	Turnout rate	Margin of victory
Mixed gender election	0.000682	-0.00831*
	(0.000637)	(0.004659)
Minority gender election	0.000412	-0.01224**
	(0.000647)	(0.004687)
Mixed * Minority	0.000763	-0.00747*
	(0.000649)	(0.004717)
Mean	0.037087	0.293369

Notes: Each cell is a separate regression, all including city and year FE's. Turnout rate is the sum of votes received by top two candidates divided by city population in 2023. Margin of victory is the absolute value of the difference between the top two candidates' normalized vote shares.

Figure A1. Gender-Minority Penalty Robustness to Bandwidth Choice

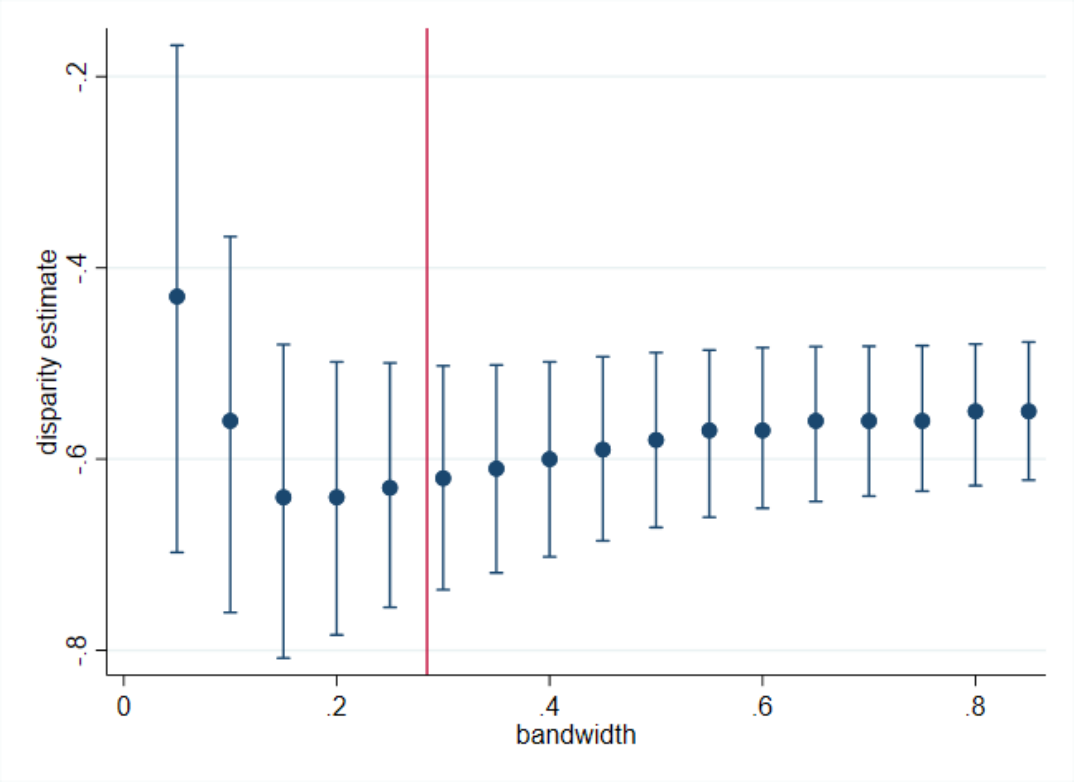

Notes: Plots penalty estimates from Figure 1 with varying bandwidths. The red line at x = 0.28 denotes the bandwidth calculated from McCrary's default bandwidth calculation.

Figure A.2 Density Test for Gender-Minority Candidate Victories – Using Cattaneo et al. (2020) Density Test

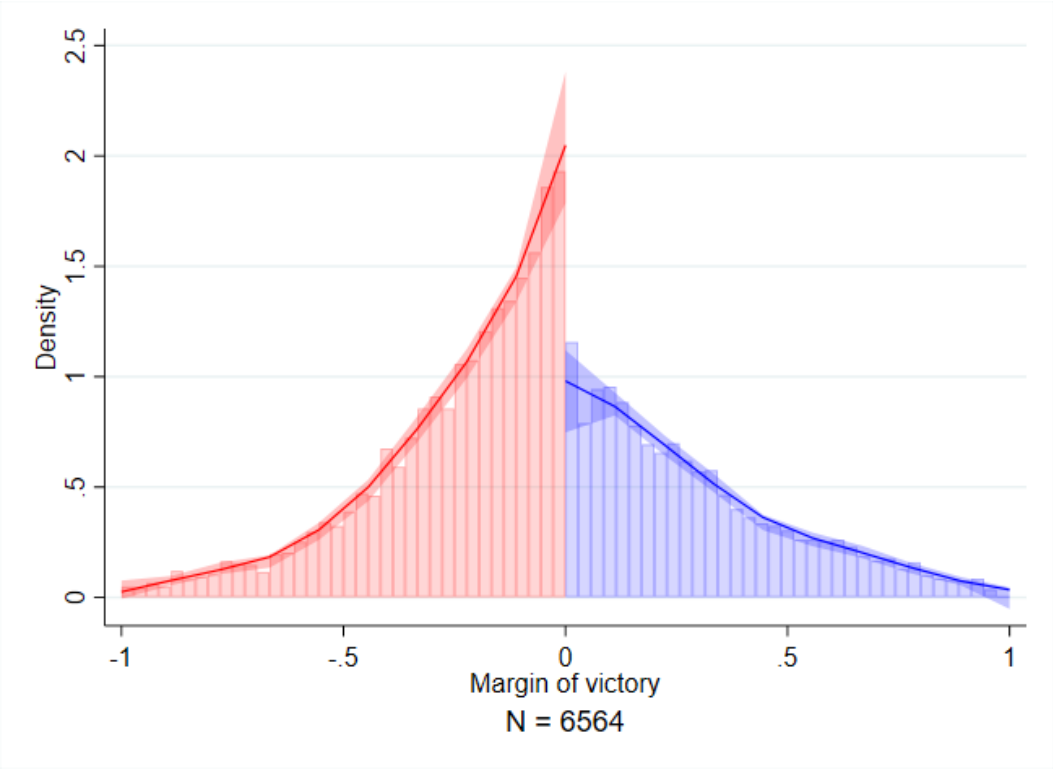

Notes: This figure presents a density test for gender-minority candidates' margin of victory, which is the difference between the gender-minority candidate's vote share and their gender-majority opponent's vote share. Vote shares are normalized between both candidates. Calculated using the *rddensity* Stata package provided by Cattaneo et al. (2020).

Figure A3. Female Candidate Victories in Mixed Gender Elections

a. Female candidates

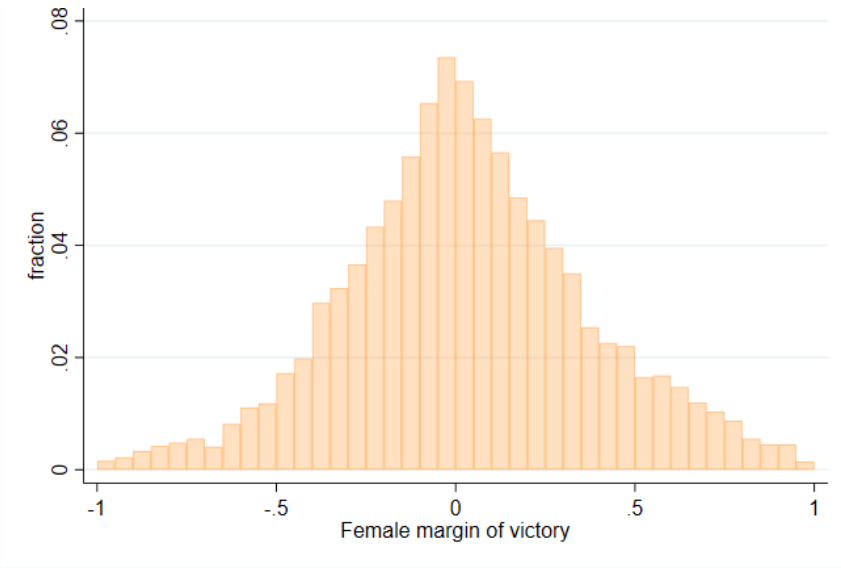

b. Female candidates by institution gender composition

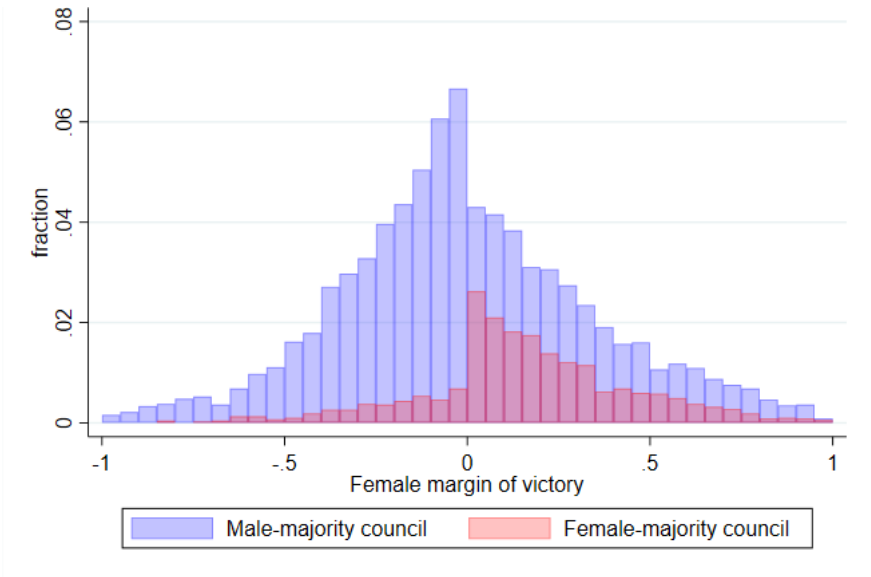

Notes: This figure presents histograms for female candidates' margin of victory, which is the difference between the female candidate's vote share and their male opponent's vote share. Vote shares are normalized between both candidates. Margin of victory is plotted within 5% bins. In panel B, the y-axis is fraction of the total sample, meaning the fraction of sample in both council types sums to the corresponding value in panel A (the fraction at 10% margin of victory in panel A is equal to the sum of male-majority and female-majority samples at 10% margin of victory in panel B).

Figure A4. Density Test for Female Candidate Victories in Mixed Gender Elections

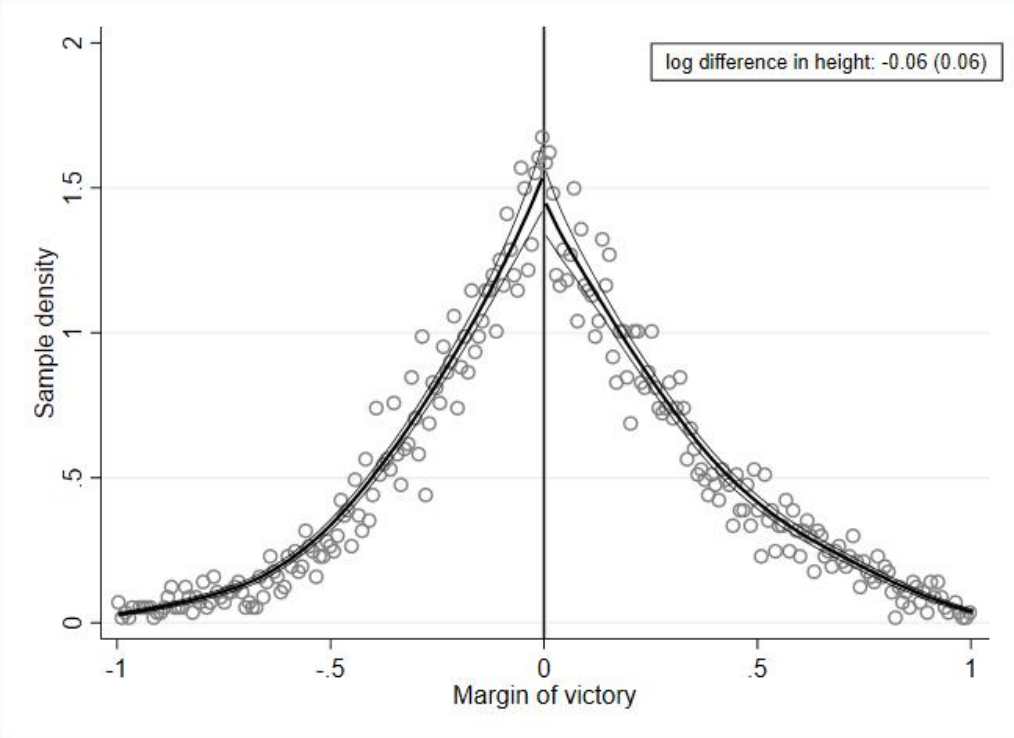

Note: N = 6,855. This figure presents a density test for female candidates' margin of victory, which is the difference between the female candidate's vote share and their male opponent's vote share. Vote shares are normalized between both candidates.

Figure A.5 Density Test for Gender-Minority Candidate Victories – Open Elections

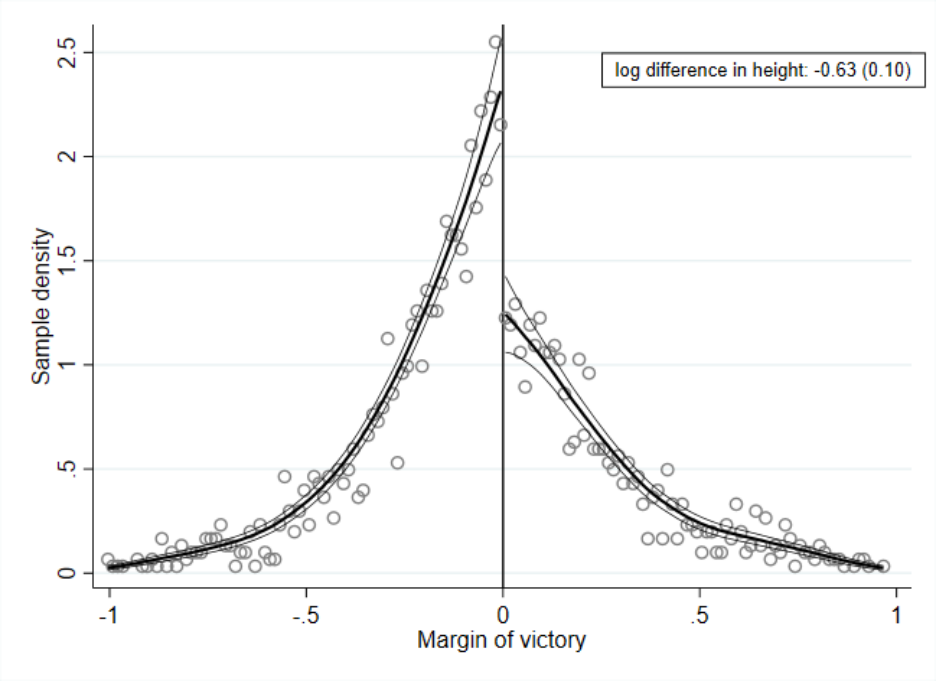

Notes: N = 2,420. This figure presents a density test for gender-minority candidates' margin of victory when neither candidate is an incumbent, which is the difference between the gender-minority candidate's vote share and their gender-majority opponent's vote share. Vote shares are normalized between both candidates.

Figure A.6 Density Test for Gender-Minority Candidate Victories – Heterogeneity by Party Affiliation

A. GMC is Majority Party, OPP is Majority Party

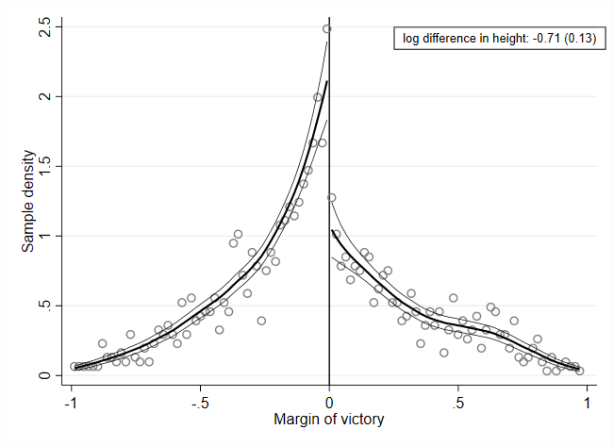

B. GMC is Minority Party, OPP is Minority Party

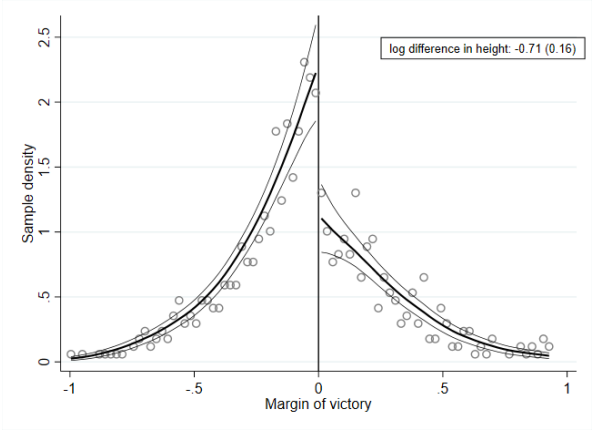

C. GMC is Minority Party, OPP is Majority Party

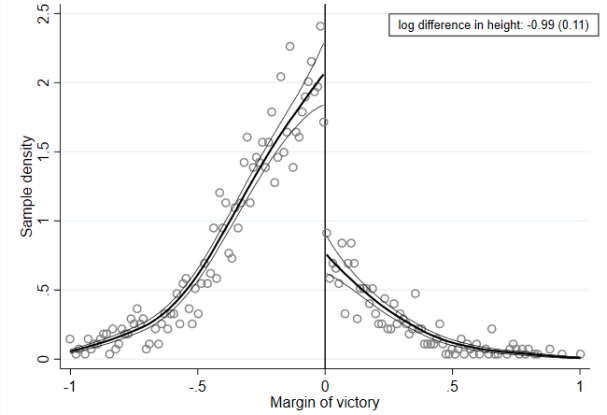

D. GMC is Majority Party, OPP is Minority Party

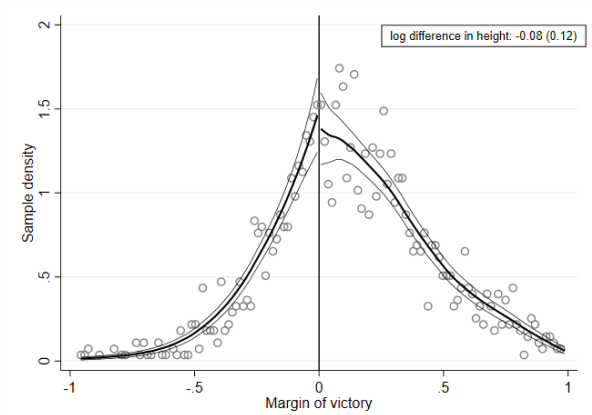

Notes: Panel A: N = 1,685, Panel B: N = 738, Panel C: N = 2,283, Panel D: 1,858. This figure presents a density test for gender-minority candidates' margin of victory, which is the difference between the gender-minority candidate's vote share and their gender-majority opponent's vote share. Vote shares are normalized between both candidates. Each figure is estimated separately by whether the gender-minority candidate (GMC) or their opponent (OPP) match the institution's (i.e. city council, county legislature, or school board) majority political party (democrat or republican)

Figure A.7 Density Test for Historically Gender-Minority Candidate Victories - After Majority Gender Switch

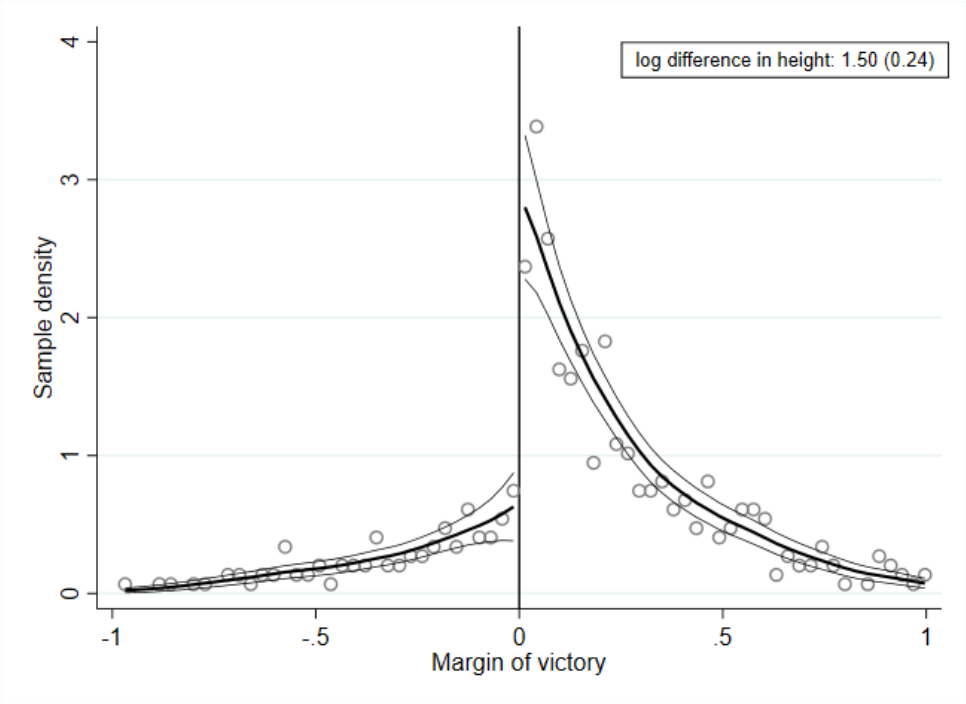

Notes: N = 526. This figure presents a density test for gender-minority candidates' margin of victory, which is the difference between the gender-minority candidate's vote share and their gender-majority opponent's vote share. Vote shares are normalized between both candidates. The estimation only includes candidates whose gender would previously have been the institution's (i.e. city council, county legislature, or school board) minority gender, but now matches the current institution's majority gender (e.g. the majority gender has switched from male to female).

Figure A.8 Density Test for Gender-Minority Candidate Victories – By Matching Gender of Election Administrator

A. Gender-Minority Candidate and Administrator Share Gender

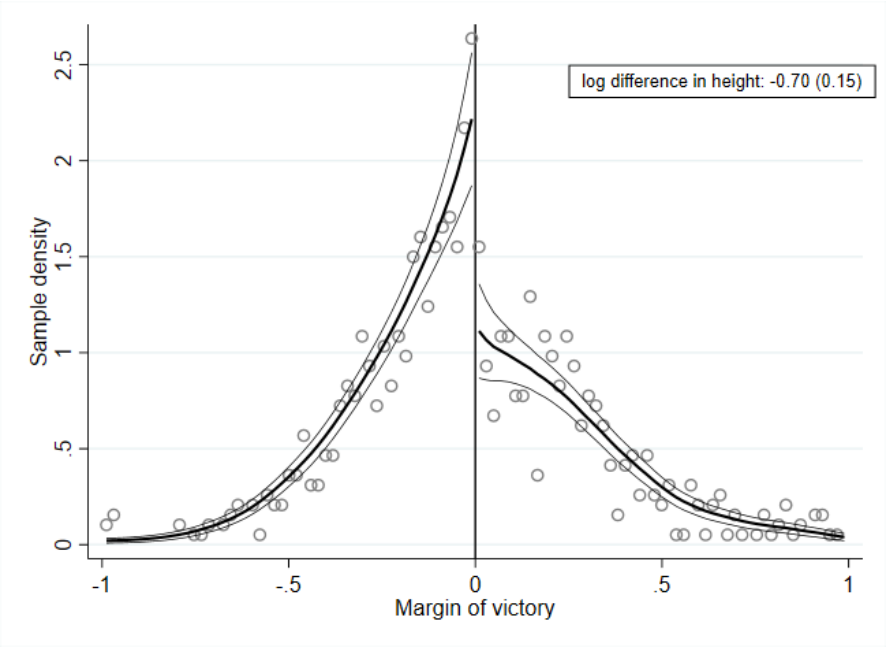

B. Gender-Minority Candidate and Administrator Are Different Genders

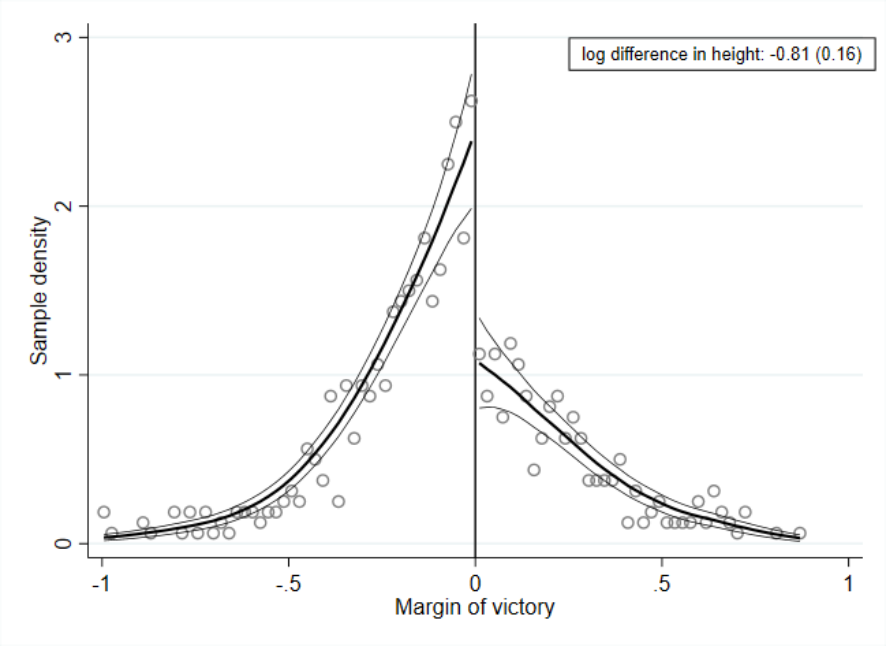

Notes: Panel A: N = 988, Panel B: N = 764. This figure presents a density test for gender-minority candidates' margin of victory, which is the difference between the gender-minority candidate's vote share and their gender-majority opponent's vote share. Vote shares are normalized between both candidates. The estimate only includes candidates for whom the gender of the head local election administrator can be determined, using data generously provided by the authors of Ferrer and Thompson (2025).