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 contests, candidates who do not match the majority gender of current officeholders are systematically disadvantaged. Near the victory cutoff, these candidates are 1.9 times as likely to lose as win. This gender penalty affects both women and men and is present in elections for city councils, school boards, and county legislatures. Evidence suggests the effect may be mediated by current officeholders steering campaign donations toward candidates who match their own gender. These findings highlight how those in power can shape the compositions of future leadership.

* Email: nduzett@tamu.edu. I am grateful to Andrew Barr, Jonathan Meer, Joanna Lahey, Steve Puller, Melissa Gentry, and participants at the Texas Applied Microeconomics Students Workshop for helpful comments. All errors are my own.
This research did not receive any specific grant from funding agencies in the public, commercial, or not-for-profit sectors.

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 divides in education and labor force participation, gender imbalances persist in many key leadership and decision-making roles. Women hold roughly 10 percent of CEO positions in Fortune 500 companies and under a third of U.S. college presidencies and Federal judgeships.[1] In terms of shaping policy, government holds some of the most impactful positions of power and is no stranger to gender disparities. As of 2025, women comprise 28 percent of the U.S. Congress, a marked increase from the 10 percent recorded in 1992's "Year of the Woman," yet still far from parity.[2]

What explains the gender disparity in leadership? Prior research suggests multiple potential channels: 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 a 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 efforts 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), progress toward parity is slow.

In this paper, I document a new fact regarding the barriers to power faced by underrepresented genders. Specifically, I show that in close elections between candidates of different genders, outcomes systematically favor the candidate whose gender matches 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. Given the widespread overrepresentation of men in positions of power, this has important implications for efforts aimed at increasing the representation of women.

I examine this idea in the context of local government, investigating how the gender composition of U.S. local elected positions influences the likelihood of different-gender

[1] Pew Research Center and the American Bar Association
[2] Pew Research Center

candidates winning a close election for city council, county legislature, or school board. These local governments shape critical 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. Additionally, because these legislative bodies typically rely on majority voting, even slight gender imbalances may play a decisive role in shaping gender-related policies.[3]

Using the largest existing collection of U.S. local government elections,[4] 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. I refer to the underrepresented candidates in such elections as *gender-minority candidates*, and their opponents as *gender-majority candidates*. To isolate the role of gender in determining electoral outcomes, I estimate differences in victory probabilities (i.e., vote share densities) in close elections between these candidates and their opponents.

I find that the gender composition of current officeholders plays a key role in deciding close elections. Gender-minority candidates are 1.9 times as likely to lose a close election as win. This result is robust to bin size, bandwidth size, and density estimation alternatives. Further, the gender-minority penalty in close elections is present in city council, county legislature, and school board elections, demonstrating that this imbalance is not unique to one context. Gender-minority candidates seeking elections to more imbalanced councils face even greater penalties: in particular, gender-minority women are 7.4 times as likely to lose a close election as win when the council is exclusively male. Finally, the penalty is not specific to one gender. Both female and male candidates face penalties in the probability of close election victories when they are the underrepresented gender.[5]

These estimates are meaningful in magnitude. The advantage for gender-majority candidates is comparable in size to the party incumbency advantage estimated by Lee (2008). The corresponding penalty for gender-minority candidates implies long-run consequences for

[3] 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.

[4] The American Local Government Elections Database, Benedictis-Kessner et. al (2023)

[5] While women are more likely to be the gender-minority candidate in city council and county legislature elections (and more likely the gender-majority in school boards), the gender-minority penalty is present for both women and men in elections to all office types, i.e. the penalty for a given gender is not driven by a single office type.

representation as well. A female loss reduces female representation initially, but also lowers the likelihood that future female candidates win seats as the representation of males on a council increases. Governing bodies that become less female due to this penalty may suffer from lower participation among remaining female officials (Rebolledo 2024) and may implement fewer female-friendly policies (Chattopadhyay and Duflo 2004). Female losses may also have negative downstream effects on the pipeline for future leaders (Beaman et al. 2009).

What may be contributing to the gender-minority penalty in close elections? Gaps in financial resources are critical in elections, particularly in close contests where contributions increase and small investments of additional resources could decide the outcome (Gerber 2004, Bouton et al. 2018). I find suggestive evidence that the gender penalty in close elections may be mediated by differences in campaign contributions.[6] I show that gender-minority candidates receive roughly half the contributions received by the average candidate and $16,771 (43 percent) less in contributions than their majority-gender opponents. This gap in contributions nearly triples in close elections. While lacking statistical significance and reliant on a small and non-representative sample, this evidence suggests that 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 current officeholders are steering financial support (e.g., donor networks, endorsements, fundraising advice) toward individuals whose gender matches their own, contributing to systematic losses for gender-minority candidates. This influence channel may play an important role in perpetuating the gender imbalances in positions of power that are prevalent in many sectors.

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

[6] Other explanations, including incumbency advantages, party affiliation, voter preferences, and corruption are not supported by the data (explored in Section IV). Similarly, I see no evidence of differential turnout for majority-gender candidates, the explanation for the increased odds of Black candidate victory in Southern mayoral elections in 1965-2010 (Vogl 2014). Redistricting is another channel by which those in power may influence characteristics of future officials, but there is limited evidence on its effects, focusing on racial minorities (Washington 2012).

information for over 57,000 elections between 1989-2021 in U.S. cities and counties with populations over 50,000 in the 2020 Census. These data cover elections to many local offices, including city councils, county legislatures, and school boards. They also provide some demographic information for candidates, including race, gender, and major party affiliation. While previous work on local government elections in the U.S. is usually restricted to a few hundred elections in one state and only a few years of data, the national nature of this dataset allows me to examine how gender disparities in positions of power fluctuate 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 underrepresented. I examine elections where the gender of both the winner and loser are known and are different.[7] I then attach gender composition data for each institution (i.e. city council, county legislature, or school board) at the time of the election and determine which candidate is the gender-minority.[8] I use the gender-majority status of the council to assess which gender is advantaged—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.[9] I will refer to candidates who do not match the majority gender of the contested multimember office as *gender-minority candidates*, and to their opponents who do match as *gender-majority candidates*.

This process results in 6,564 elections between gender-minority and gender-majority candidates. I observe 352 cities, 421 counties, and 73 school districts in which there has been at least one mixed-gender election between 1989 and 2021. While city councils and county legislatures in these localities are majority male and school boards are more often majority female,[10] both women and men are represented as gender-minority candidates in elections all

[7] 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.
[8] 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).
[9] I later examine heterogeneity in the gender-minority penalty for different gender compositions and find that the penalty is larger with more extreme compositions and more muted around balanced compositions, suggesting that gender-majority status is a reasonable threshold for determining gender disadvantage.
[10] 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.

offices. In the 6,564 elections with both gender-majority and gender-minority candidates, women make up 82.8% of gender-minority candidates, reflecting the widespread male dominance in political institutions.

B. *Detecting gender penalties in close elections*

To formally test for differences in victory probabilities for gender-minority candidates when running against an opposite-gender opponent, I use the McCrary (2008) density estimation method. Conceptually, the test is designed to detect 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 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. Finally, I separately estimate a 4^{th}-order polynomial on each side of the victory cutoff.[11] 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 these underrepresented candidates win close elections at rates equal to their opposite-gender opponents, 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 plots histograms of female candidate margin of victory in mixed-gender elections to city councils, county legislatures, and school boards overall, as well as separately by whether the candidate is seeking election to an institution that is majority-male or majority-female. The figure illustrates a stark contrast in outcomes for female candidates: while close

[11] 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 (0.285), I later show that the main estimates are not sensitive to bandwidth choice.

elections to majority-male institutions overwhelmingly result in female losses, female candidates in close elections to majority-female institutions are more likely to win than lose. These differences in probabilities of close election victory for women are obscured when failing to account for council composition.

Motivated by this descriptive observation, Figure 2 displays the result of a formal density estimate on gender-minority candidate victory probabilities 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 candidates and their opponents 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 close election losers to close election winners. The drop in density implies that gender-minority candidates are 1.86 times more likely to lose than win in close elections at the cutoff.

The estimated gender-minority penalty is robust to choice of bin size, bandwidth size, and density estimation alternatives. First, the estimates are robust to different bin size choices (Figure A.1). Second, Figure A.2 shows that the penalty estimate does not change with choice of bandwidth size between twice and half the bandwidth given by the automated bandwidth selection procedure. Third, the penalty is also robust to choice of density estimator. Cattaneo et al. (2020) introduce a density estimation method similar to McCrary (2008) that does not rely on pre-binning the data and performs better under some additional assumptions. Using this approach in Figure A.3 shows a significant shift in victory probability for gender-minority candidates at the cutoff, identical to the initial result.

Beyond technical robustness, I show that the gender-minority penalty is not isolated to a particular type of local office: the underrepresented gender is more likely to lose in city council, county legislature, and school board elections. While the main result in Figure 2 pools city council, county legislature, and school board elections to estimate the overall penalty, Figure 3 panels A, B, and C present shifts in win probabilities separately for elections to each of those offices. The figure shows that each of these offices exhibits lower likelihoods of gender-minority candidate victory. The scope of this penalty suggests that gender is a widespread factor in determining the compositions of governments and the identity of those who hold power, and that the penalty may be a more general feature of group-based positions of power in sectors beyond local government.

The gender-minority penalty in close elections is significant in magnitude and has direct implications for local government policy as well as the compositions of future leadership. The penalty is roughly three-quarters the size of the party incumbency advantage estimated by Lee (2008), showing its consequential magnitude for deciding close elections.[12] Furthermore, 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. For example, if a female candidate loses an election in a gender-balanced city council, the female share declines—not only from the immediate loss of the seat, but also from the lower probability of future female victories resulting from the now majority-male composition. The long-run decline in female representation could be even more severe than losses by existing candidates. 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 as well, lowering the potential for future representation. 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 plays a pivotal role in victory, a natural question is whether more gender-imbalanced institutions are associated with stronger gender penalties. I examine how the strength of the gender-minority penalty varies by the gender composition of the office being sought. Figure 4 plots the penalty estimates from separate estimations of gender-minority win probabilities by each percentage of the city council, county

[12] Lee (2008), in work pioneering the vote share regression discontinuity design, estimates party incumbency increases the probability of victory by 40% to 45%. While the effects are of similar magnitude, I show in the next section that the gender-minority penalty is not explained by incumbency advantage.

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—nearly four times the size of the overall estimate. 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 90% female, it is not possible to reliably estimate shifts in win probabilities 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 somewhat true for gender-minority men, but the disparities are more muted and localities where male candidates are underrepresented are much rarer. Extreme gender compositions in 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 penalty and the series of established barriers to political office for women, I examine heterogeneity in the gender-minority penalty by candidate gender. Figure A.4, panels A and B display separate estimations for shifts in win probabilities for female and male gender-minority candidates. These estimates reveal two points of interest. First, the penalty is not unique to one gender: both gender-minority women and gender-minority men face a penalty in the probability of 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 gender has no influence on electoral success conditional on candidacy (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. Figure A.5 estimates changes in victory probabilities 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, while the point estimate suggests that females are somewhat more likely to lose close elections, no statistically significant shift in density is detected when one does not account for institutional gender composition. Candidates whose gender is underrepresented on the local government body are penalized regardless of their own gender, with potentially important implications for representation and the functioning of local governments.

In addition to the policy implications, the existence of the gender-minority penalty in close elections for both women and men also challenges a core aspect of the regression discontinuity (RD) identifying assumption: the similarity of candidates just above and below the victory cutoff. 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. McCrary (2008) suggests that density tests designed to examine this core assumption 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.[13]

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.[14] Finance is a vital component 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

[13] See Marshall (2024) for an overview of challenges in identifying the causal effects of candidate characteristics using a vote share RD method.
[14] I briefly evaluate alternative explanations, such as incumbency, party, and voter preferences, at the end of this section and in the Appendix.

elections, as does the level of contributions (Gerber 2004, Bouton et al. 2018).[15] 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 showing that contributors give more to candidates who match their own gender (Bouton et al. 2022). As elected officials presumably have some influence over the flow of resources to candidates (e.g., via connections with individual donors, political action committees (PACs), political parties, and fundraising networks), their gender preferences could influence the outcomes of elections. Consistent with this idea, Gaudette and Benedictis-Kessner (2024) find evidence supporting the theory that interest groups use campaign contributions to influence the composition of local city councils.[16]

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 elections. 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 for public use. Given limited data availability, I restrict attention to city councils to enable a cleaner comparison between elections, since contributor amounts and financial requirements for running a campaign may vary widely between different local offices. The final merged sample includes 123 city council candidates in mixed-gender elections with complete campaign contribution information.[17] This sample yields an identical—though less precise—estimate of the gender-minority penalty as the main sample (Figure A.6).

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

$$Contributions_i = \beta_0 + \beta_1\, gender\text{-}minority\ candidate_i + \beta_2\, female_i + \mu_t + \gamma_c + \varepsilon_i$$

[15] In particular, Gerber (2004) conducts field experiments in state legislative races, and shows that additional spending in close elections can significantly boost the probability of incumbent victory while having minimal effect on vote share.
[16] Examining real estate industry contributions to city council and mayoral candidates, Gaudette and Benedictis-Kessner (2024) find that donors are more likely to give to winning candidates and mostly donate to non-incumbents.
[17] I discuss the drawbacks of this limited sample size near the end of this section.

$Contributions_i$ 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. μ_t and γ_c are year and city fixed effects, respectively. The primary parameter of interest is β_1 , which describes the difference in contributions between gender-minority and gender-majority candidates in mixed-gender elections.

Table 1 presents the results of this approach. Column 1 reveals that gender-minority candidates receive, on average, $16,771 (43%) less in donations than their gender-majority opponents. The gap in donations widens to $40,073 in elections where the margin of victory is less than 20%, and is comparable in magnitude when controlling for candidate gender. While lacking statistical significance, the signs and magnitudes of these results suggest a relationship between election victory, differences in contributions, and the existing council's influence over the flow of donations.

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

While the evidence suggests the existence of a relationship between campaign finance and the gender-minority penalty, the approach has limitations. First, the sample of elections with contribution information is a small and non-random subset of the main analysis sample. Second, it is possible that gender-minority candidates receive less funding when competing with a gender-majority candidate simply because gender-minority candidates are not very popular. However, I find that the gap in contributions widens substantially in closer elections where candidates are nearly tied in vote shares, indicating that popularity alone does not explain the differences in campaign finance between gender-minority and gender-majority candidates (Table 1). In the face of these limitations, coupled with the identical size of the gender-minority penalty in this smaller sample, I interpret the results on campaign contributions to be suggestive of a core mechanism behind the gender-minority penalty in close elections. Given the prevalence of the gender-minority penalty across multiple offices, this channel of influence over resources

necessary for securing positions of power may play an important role in perpetuating gender inequities across other leadership sectors as well.

Another possibility is that there is an alternative explanation for the gender-minority penalty in close elections (e.g., incumbency, party affiliation, voter preferences, selective turnout, and corruption). These explanations are not supported by the data. 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 A.7). 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 A.8). 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 A.9).

Finally, while corruption could plausibly explain the sharp discontinuity in victory rates by gender-minority status, I find no evidence of ballot manipulation or selective mobilization in close elections involving gender-minority candidates. Given the sharpness of the discontinuity near the victory cutoff, and the low-information nature of many local elections, the possibility for election fraud is a conceivable mechanism. Accusations of voter fraud against public officials often center on "ballot stuffing," i.e. selectively removing ballots for consideration or forging ballots in favor of a certain candidate. If present, this type of fraud would mechanically increase the turnout in manipulated elections relative to fair elections. I test for this form of fraud by examining differences in turnout between elections involving a gender-minority candidate and those without gender-minority candidates. I detect no statistically significant differences in turnout (Table A.2). [18] However, since the gender-minority penalty is detected by using variation in close elections, and local elections can be decided by relatively few votes, it is possible that small-scale ballot manipulation could overturn many local elections while remaining undetected with this method. I take another approach to detecting possible fraud by examining the gender of local election administrators, who oversee election practices in their appointed counties.

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

Conceivably, these election administrators could sway election outcomes toward candidates who match their own gender, either through established channels (i.e., determining polling locations and times, hiring polling workers, policies for dismissing incomplete ballots, etc.), or illicit channels (i.e., directing workers to dismiss legal ballots, improperly purging voter rolls, etc.) To examine this possibility, I link election administrator gender to a subsample of gender-minority elections. I find no apparent difference in the gender-minority penalty whether the reigning election administrator matches the gender-minority candidate's gender or not. Ultimately, while corruption cannot be ruled out as a possible mechanism, I find no evidence in favor of corruption as a driving mechanism behind the widespread gender-minority penalty.

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 supported by institutional composition: candidates whose gender does not match the most common gender of current officeholders are more likely to lose, even in elections decided by narrow margins. Notably, the result is not unique to one political office, suggesting it is a general result and thus may apply to gender imbalances in positions of power beyond government, particularly in settings in which those currently in power have some influence over the flow of resources to those seeking the same positions.

The scope of the gender-minority penalty has broad implications for understanding why gender parity remains elusive in positions of power, despite gains in other domains. The ability of current powerholders to shape the characteristics of their successors—potentially through access to critical resources—can help explain the widespread gender imbalances observed in many powerful positions. As men are often overrepresented in such positions, this has particularly sharp implications for women seeking leadership roles and policies aimed at increasing female representation. Despite gains in female representation across some domains, gender imbalances in government and other high-level positions may self-perpetuate even when those in power have no formal channel for picking their successors.

Figure 1. Female Candidate Margins of Victory in Mixed-Gender Elections

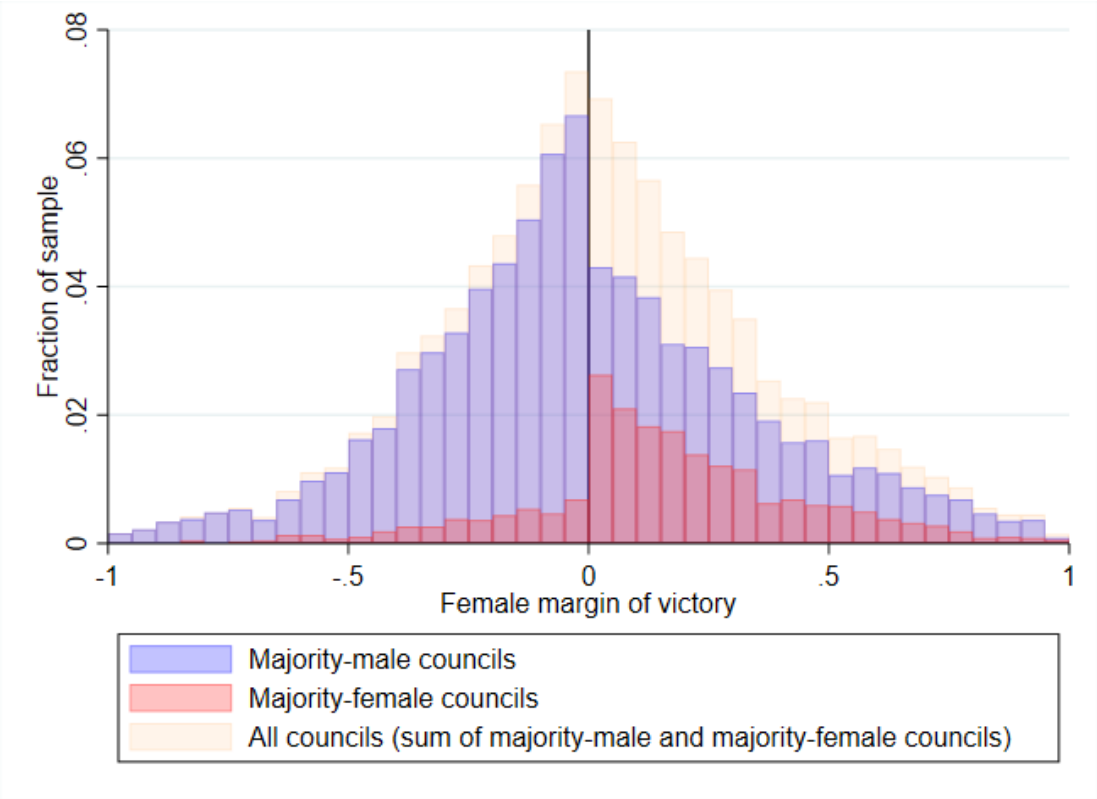

Notes: This figure presents histograms for female candidates' margin of victory in mixed-gender elections. Margin of victory is the difference between the female candidate's vote share and their male opponent's vote share (positive values indicate female victory while negative values indicate female loss). Total vote share is calculated only using votes received by the female and male candidates when one wins and the other loses. Margin of victory is plotted within 5% bins. The y-axis is the fraction of the total sample, meaning the fraction of sample in both council types sums to the corresponding value of the orange bars (the fraction at 10% margin of victory in the orange bars is equal to the sum of majority-male and majority-female samples at 10% margin of victory). While the legend uses the term "councils" for brevity, elections to county legislatures and school boards are also included.

Figure 2. The Gender-Minority Penalty: Differences in Victory Probabilities in Close Mixed-Gender Elections

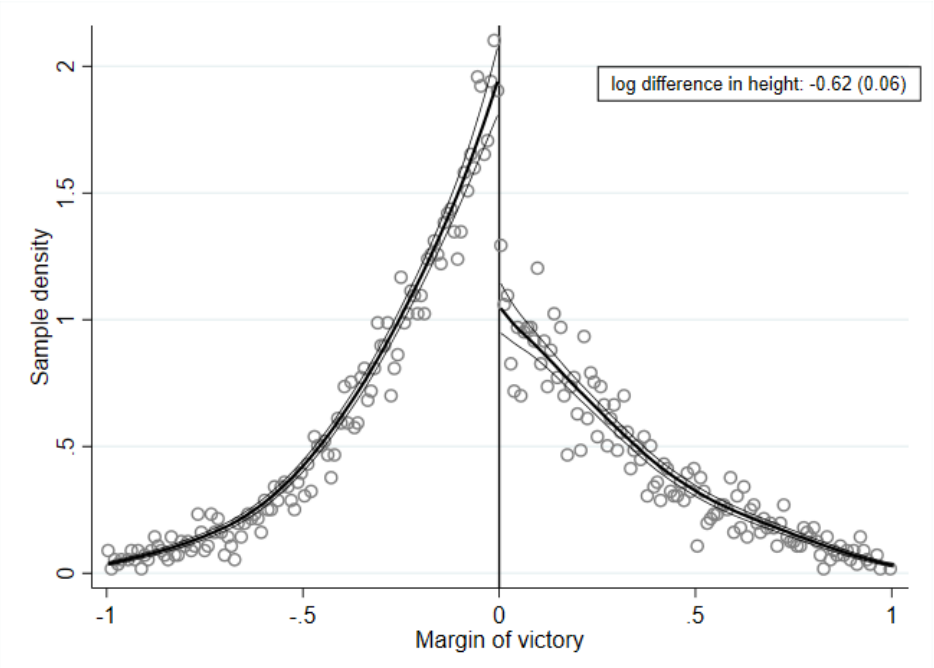

Note: N= 6564. . This figure presents estimates of shifts in win probabilities for gender-minority candidates in mixed-gender elections. Margin of victory is the difference between the gender-minority candidate's vote share and their gender-majority opponent's vote share (positive values indicate gender-minority victory while negative values indicate gender-minority loss). Total vote share is calculated only using votes received by the gender-minority and gender-majority candidates when one wins and the other loses. The sample size will be smaller than the sample in Table A.1, which also includes contests between two gender-minority candidates.

Figure 3. Heterogeneity in the Gender-Minority Penalty by Office Type

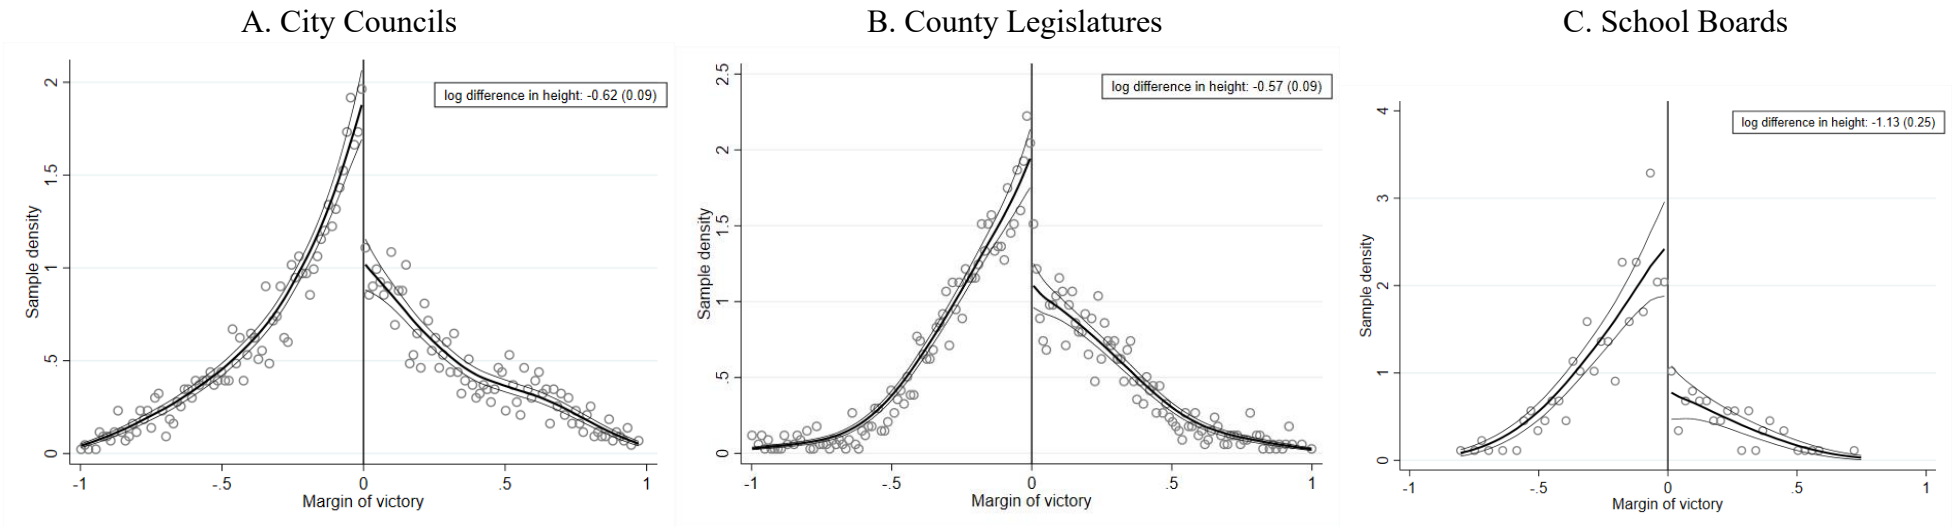

Note: Panel A: N = 3,319, Panel B: N = 2,921, Panel C: N = 324. This figure presents estimates of shifts in win probabilities for gender-minority candidates in mixed-gender elections separately by the office type being contested. Margin of victory is the difference between the gender-minority candidate's vote share and their gender-majority opponent's vote share (positive values indicate gender-minority victory while negative values indicate gender-minority loss). Total vote share is calculated only using votes received by the gender-minority and gender-majority candidates when one wins and the other loses.

Figure 4. Heterogeneity in the Gender-Minority Penalty by Institutional Gender Composition

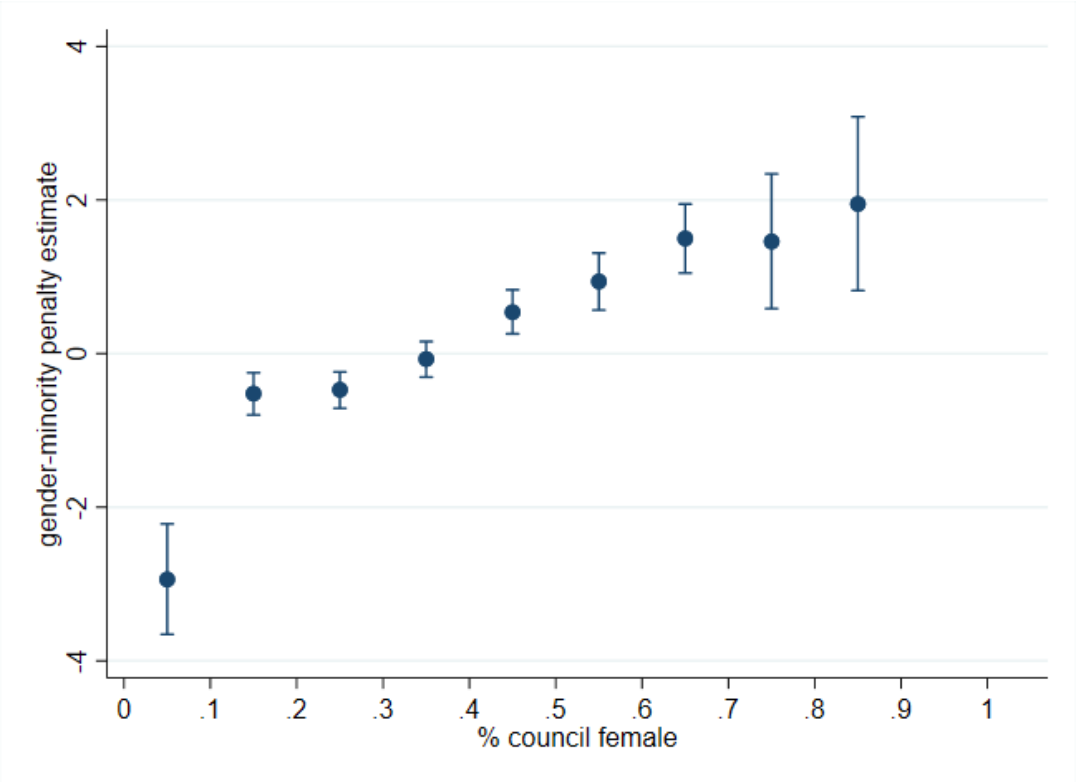

Notes: This figure plots estimates of shifts in win probabilities for female candidates in institutions with different gender compositions. Institutions are binned in 10% intervals by what percentage of their members are female (e.g. the first bin includes all institutions that are between 0%-9.9% female, inclusive). The y-axis is the estimated gender penalty (or advantage) for female candidates running against a male opponent. While the x-axis uses the term "council" for brevity, county legislatures and school boards are also included.

Table 1. Imbalanced Campaign Contributions in Mixed-Gender Elections

	Contributions ($)		
	(1)	(2)	(3)
		$\|MOV\| \leq 0.28$	$\|MOV\| \leq 0.28$
Gender-Minority Candidate (MGC)	-16,771	-40,073	-26,742
	(15,649)	(25,068)	(16,799)
Female			-33,199
			(21,636)
Mean	36,086	35,841	35,841
N	123	76	76

Notes: This table present regression results from the specification given in Section IV. The sample is a subsample of the main data for which campaign contribution data are available. In this table, all candidates are seeking election to city council. While the sample is small, the gender-minority penalty estimate in this smaller sample is identical to the main estimate shown in Figure 2, though less precise. Gender-Minority Candidate (GMC) = 1 if the candidate's gender does not match the majority gender of the current city council. MOV is the margin of victory for the candidate, i.e. the difference between their vote share and their opponent's vote share. Total vote share is calculated only using votes received by the candidate and their opponent when one wins and the other loses. Columns 2 and 3 restrict the sample further to close elections where the margin of victory is less than 0.28, selected according to the bandwidth selection procedure outline in McCrary (2008), and is the same bandwidth used in Figure 2. All columns include city and year fixed effects.

APPENDIX

Table A1. Summary Statistics for Election Data

Panel A: Localities	Cities	Counties	School districts
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 number of 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. Turnout in Mixed-Gender Elections

	Turnout rate
Mixed-gender election (MGE)	-0.0005
	(0.0025)
Gender-minority election (GME)	-0.0008
	(0.0012)
MGE * GME	0.0019
	(0.0028)
Mean	0.0368
N	8,039

Notes: This table presents regression results from regressing election gender attributes on the election turnout rate. The sample includes all city council elections that match to city population in 2023. Turnout rate is the sum of votes received by top two candidates divided by city population in 2023; intuitively, turnout rate is the share of the city's population that voted for one of the two candidates. In elections with multiple winners, only the winner with the lowest vote total is included. Mixed-gender election (MGE) = 1 if the two competing candidates differ by gender. Gender-minority election (GME) = 1 if one or both of the candidates' gender does not match the majority gender of the city council. This approach is similar to one used by Vogl (2014) to study differences in turnout in U.S. Black-White mayoral elections.

Figure A1. Gender-Minority Penalty Robustness to Bin Size

a. Half optimal bin size

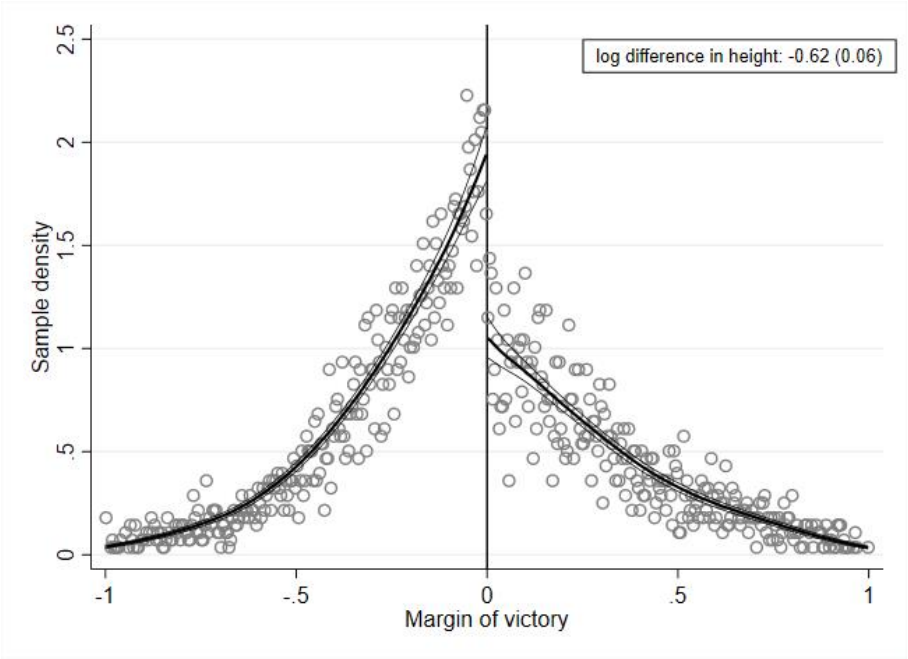

b. Twice optimal bin size

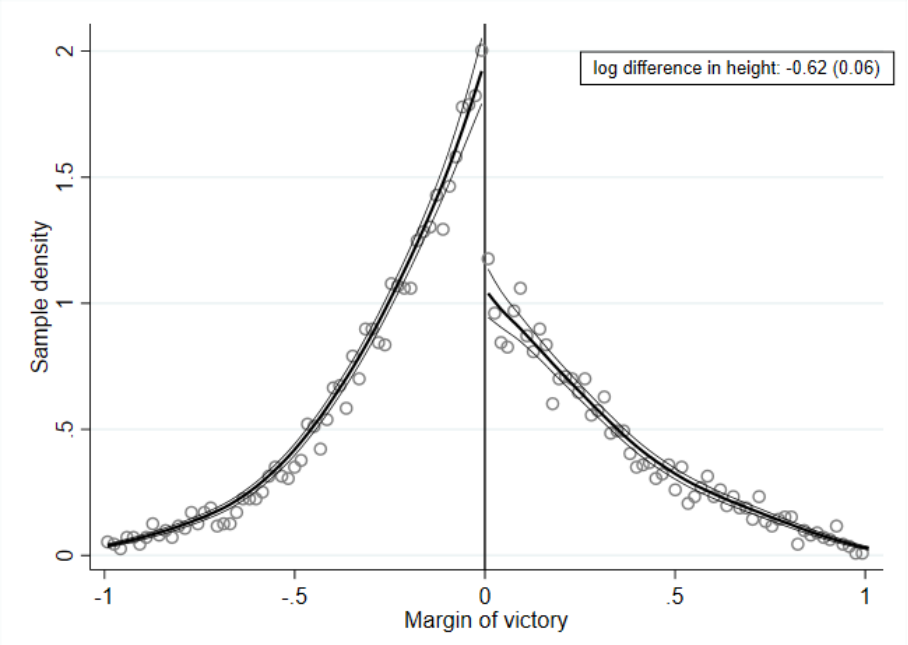

Notes: N= 6564 for both panels. This figure presents estimates of shifts in win probabilities for gender-minority candidates in mixed-gender elections for different bin sizes. Margin of victory is the difference between the gender-minority candidate's vote share and their gender-majority opponent's vote share (positive values indicate gender-minority victory while negative values indicate gender-minority loss). Total vote share is calculated only using votes received by the gender-minority and gender-majority candidates when one wins and the other loses. Default bin size calculation is b = 0.0084.

Figure A2. Gender-Minority Penalty Robustness to Bandwidth Choice

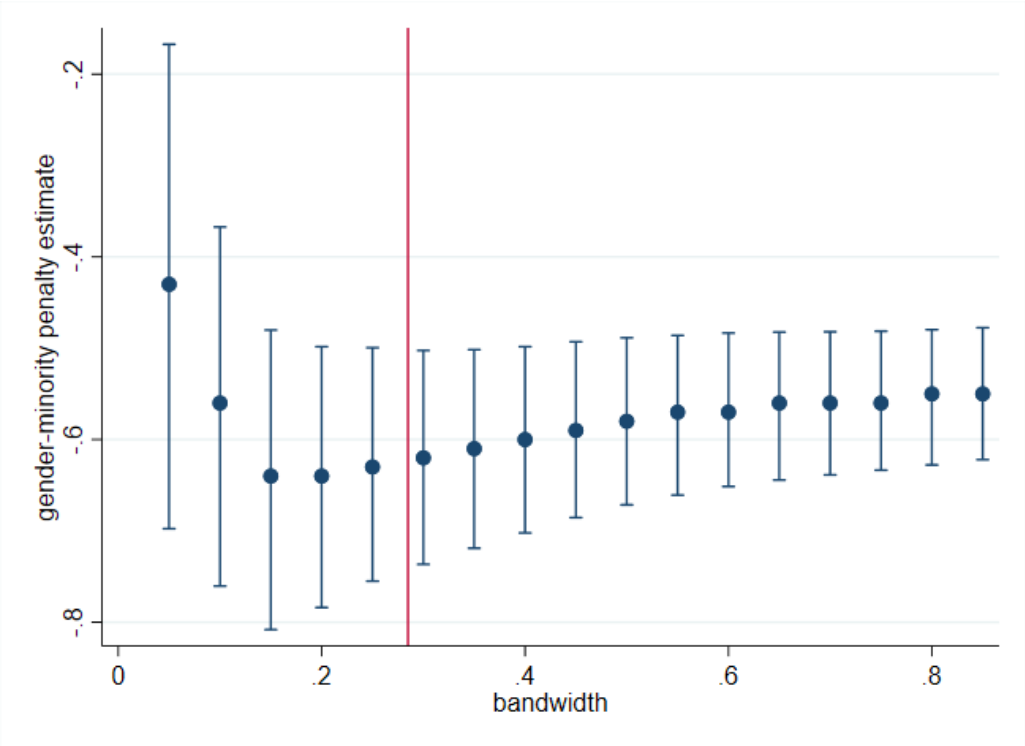

Notes: This figure plots estimates of the difference in win probabilities between gender-minority and gender-majority candidates with varying bandwidths. The red line at x = 0.28 denotes the value calculated from the bandwidth selection method described in McCrary (2008), which uses the average mean squared error of the estimated polynomials on each side of the victory cutoff to choose an appropriate bandwidth.

Figure A.3 The Gender-Minority Penalty – Estimated Using Cattaneo et al. (2020) Density Test

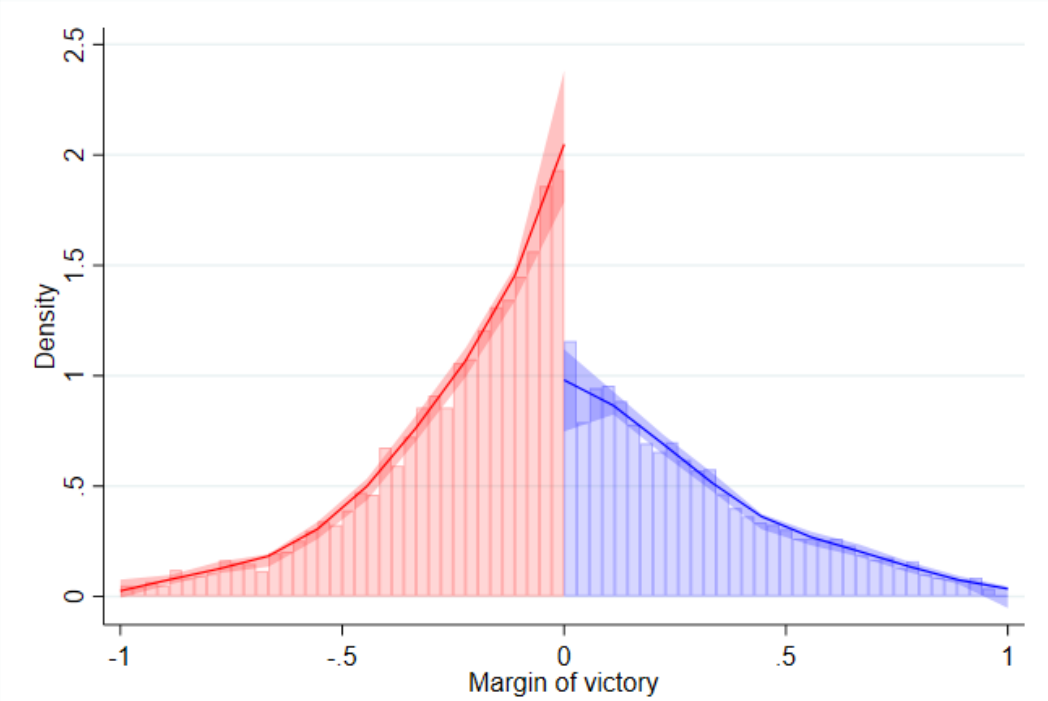

Notes: N = 6564. This figure presents estimates of shifts in win probabilities for gender-minority candidates in mixed-gender elections. Margin of victory is the difference between the gender-minority candidate's vote share and their gender-majority opponent's vote share (positive values indicate gender-minority victory while negative values indicate gender-minority loss). Total vote share is calculated only using votes received by the gender-minority and gender-majority candidates when one wins and the other loses. Calculated using the *rddensity* Stata package provided by Cattaneo et al. (2020).

Figure A.4 Heterogeneity in the Gender-Minority Penalty by Candidate Gender

a. Gender-Minority Female Candidates

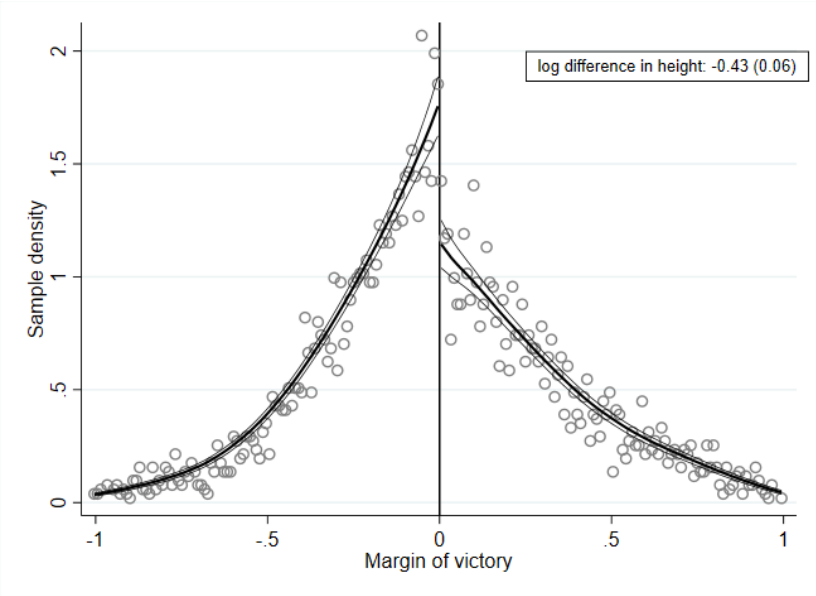

b. Gender-Minority Male Candidates

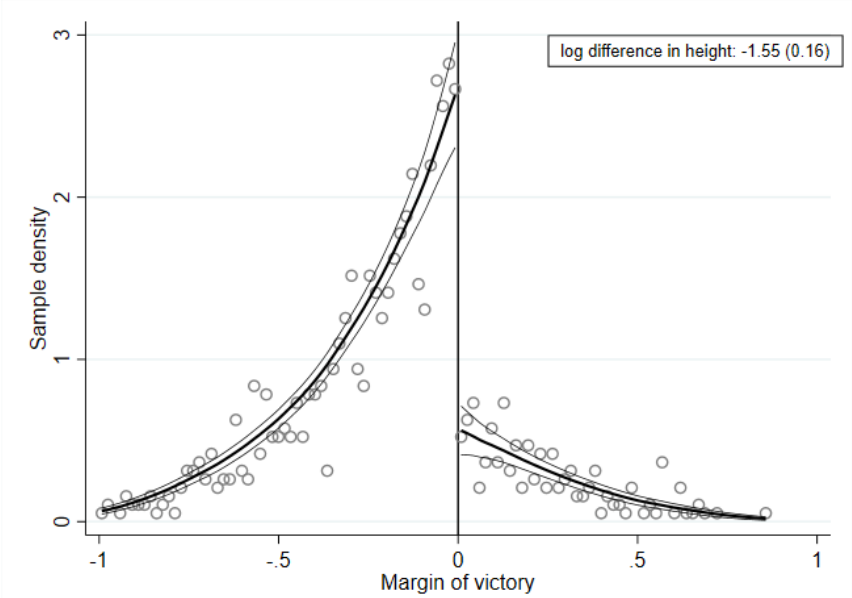

Note: Panel A: N = 5437, Panel B: N = 1127. This figure presents estimates of shifts in win probabilities for gender-minority candidates in mixed-gender elections separately by gender. Margin of victory is the difference between the gender-minority candidate's vote share and their gender-majority opponent's vote share (positive values indicate gender-minority victory while negative values indicate gender-minority loss). Total vote share is calculated only using votes received by the gender-minority and gender-majority candidates when one wins and the other loses.

Figure A.5 Female Candidate Victories in Mixed Gender Elections

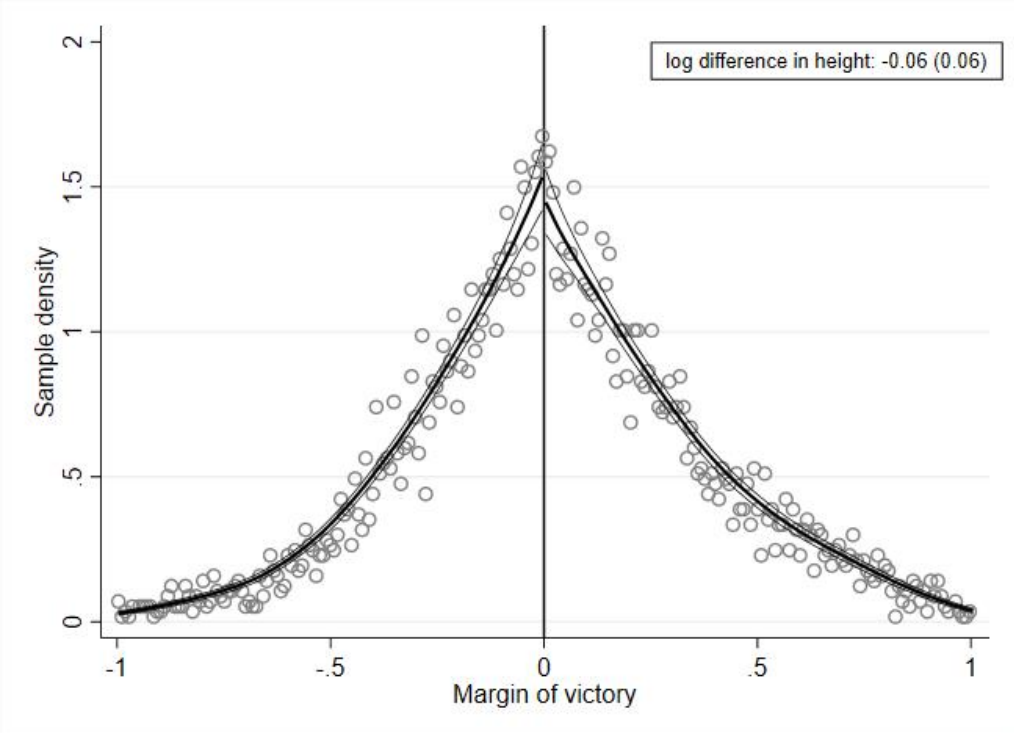

Note: N = 6,855. This figure presents estimates of shifts in win probabilities for female candidates running against a male candidate. Margin of victory is the difference between the female candidate's vote share and their male opponent's vote share (positive values indicate female victory while negative values indicate female loss). Total vote share is calculated only using votes received by the female and male candidates when one wins and the other loses.

Figure A.6 The Gender Minority Penalty in the Contribution Sample

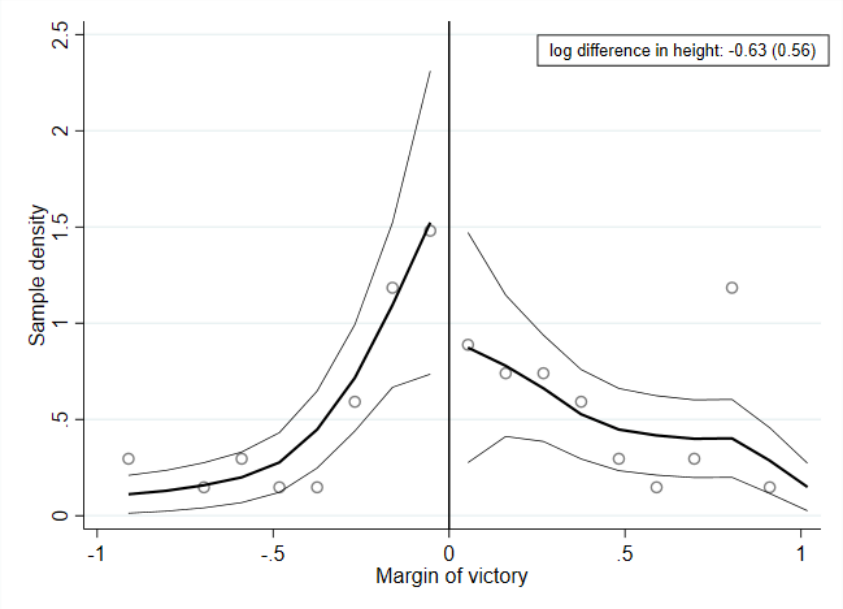

Notes: N = 63. This figure presents an estimate of the shift in win probabilities for gender-minority candidates in mixed-gender elections for candidates who match to campaign contribution data. Margin of victory is the difference between the gender-minority candidate's vote share and their gender-majority opponent's vote share (positive values indicate gender-minority victory while negative values indicate gender-minority loss). Total vote share is calculated only using votes received by the gender-minority and gender-majority candidates when one wins and the other loses.

Figure A.7 The Gender Minority Penalty in Open Elections

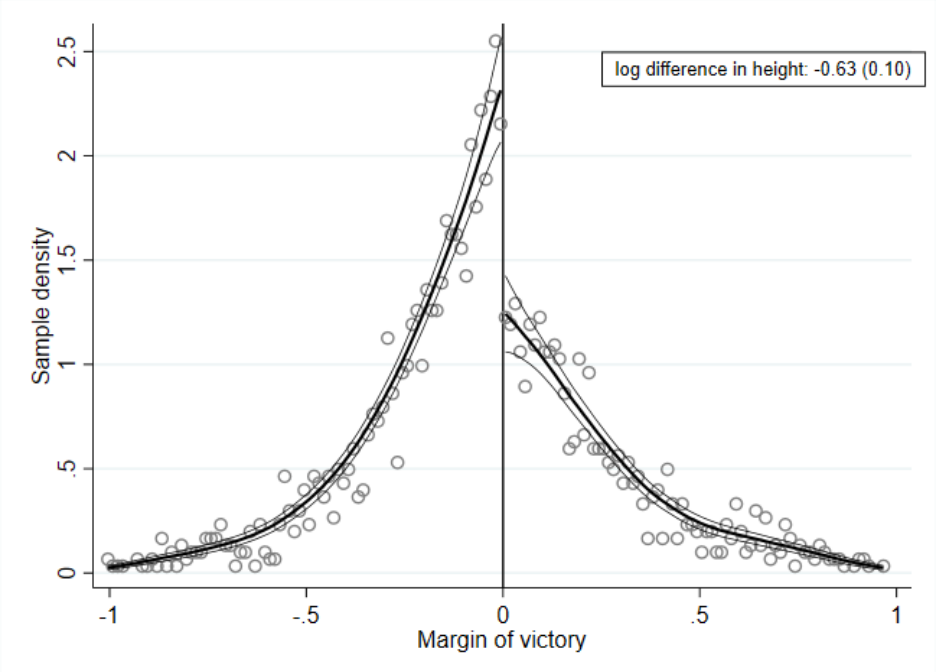

Notes: N = 2,420. This figure presents an estimate of the shift in win probabilities for gender-minority candidates in mixed-gender elections when neither candidate is an incumbent. Margin of victory is the difference between the gender-minority candidate's vote share and their gender-majority opponent's vote share (positive values indicate gender-minority victory while negative values indicate gender-minority loss). Total vote share is calculated only using votes received by the gender-minority and gender-majority candidates when one wins and the other loses.

Figure A.8 Heterogeneity in the Gender-Minority Penalty by Party Affiliation

A. GMC is Majority Party, OPP is Majority Party

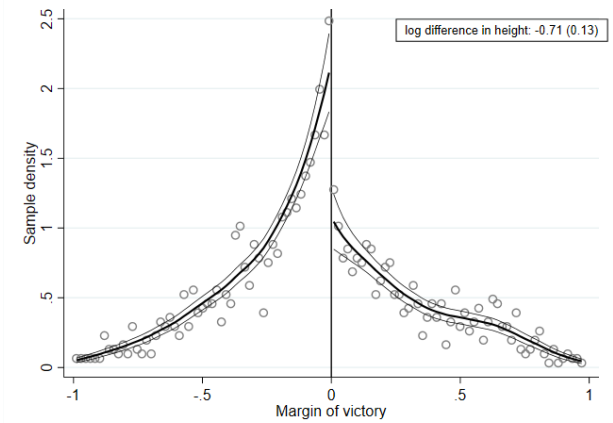

B. GMC is Minority Party, OPP is Minority Party

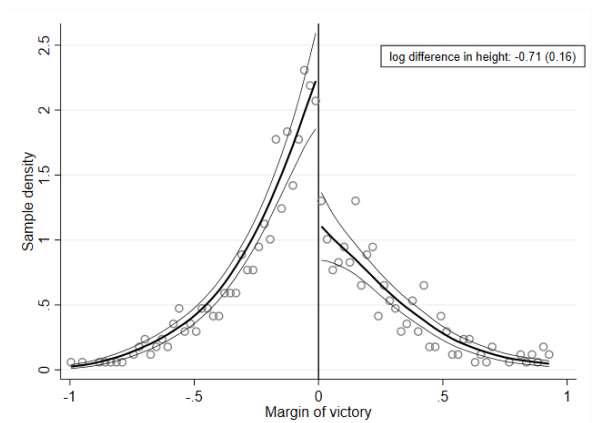

C. GMC is Minority Party, OPP is Majority Party

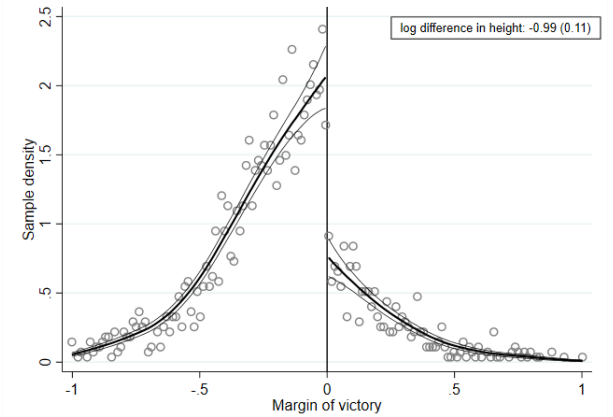

D. GMC is Majority Party, OPP is Minority Party

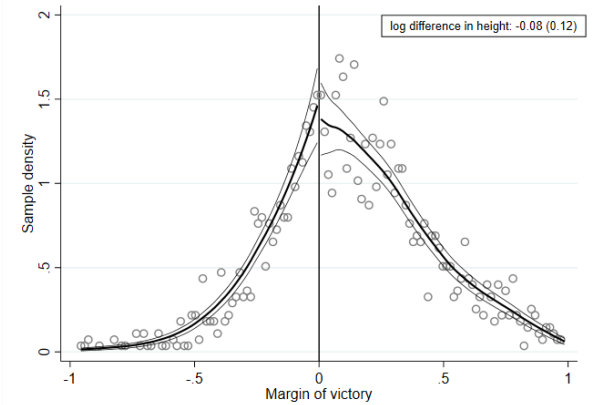

Notes: Panel A: N = 1,685, Panel B: N = 738, Panel C: N = 2,283, Panel D: 1,858. This figure presents estimates of shifts in win probabilities for gender-minority candidates in mixed-gender elections. Margin of victory is the difference between the gender-minority candidate's vote share and their gender-majority opponent's vote share (positive values indicate gender-minority victory while negative values indicate gender-minority loss). Total vote share is calculated only using votes received by the gender-minority and gender-majority candidates when one wins and the other loses. 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.9 The Gender-Minority Penalty After the Majority Gender Switches

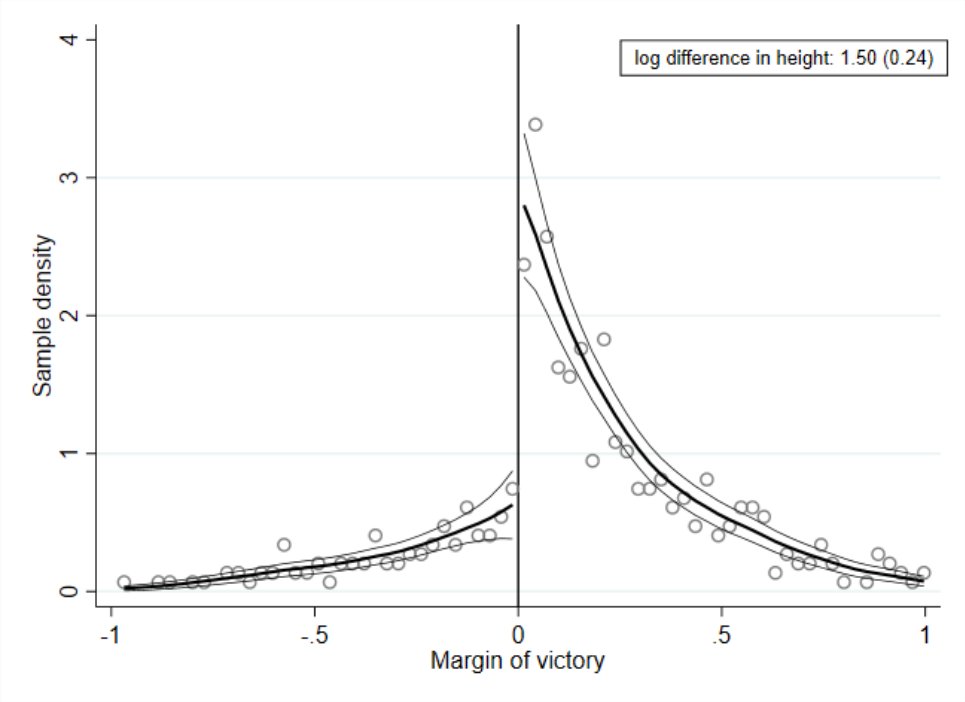

Notes: N = 526. This figure presents an estimate of the shift in win probabilities for gender-minority candidates in mixed-gender elections. Margin of victory is the difference between the gender-minority candidate's vote share and their gender-majority opponent's vote share (positive values indicate gender-minority victory while negative values indicate gender-minority loss). Total vote share is calculated only using votes received by the gender-minority and gender-majority candidates when one wins and the other loses. 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.10 Role of Election Administrator Gender in the Gender-Minority Penalty

A. Gender-Minority Candidate and Administrator Share Gender

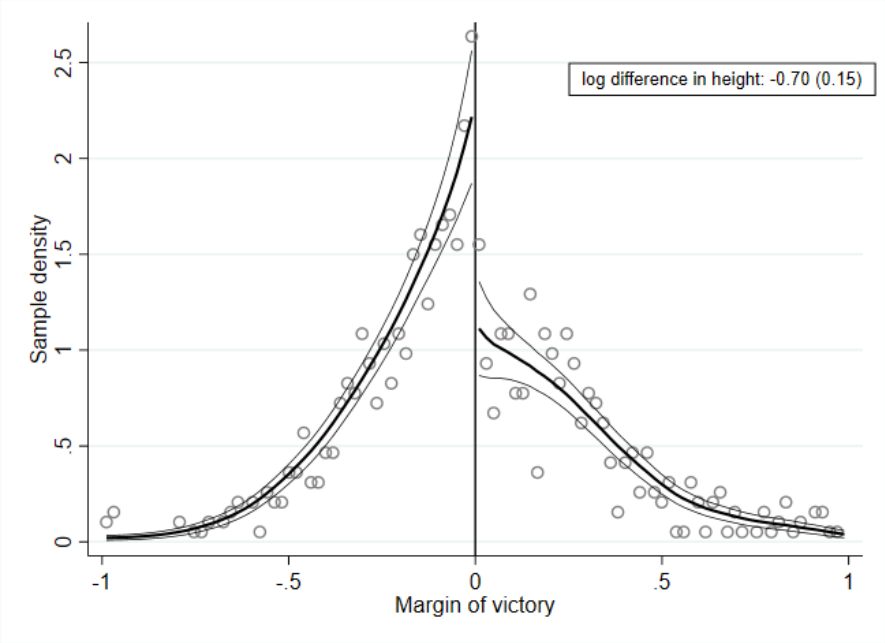

B. Gender-Minority Candidate and Administrator Are Different Genders

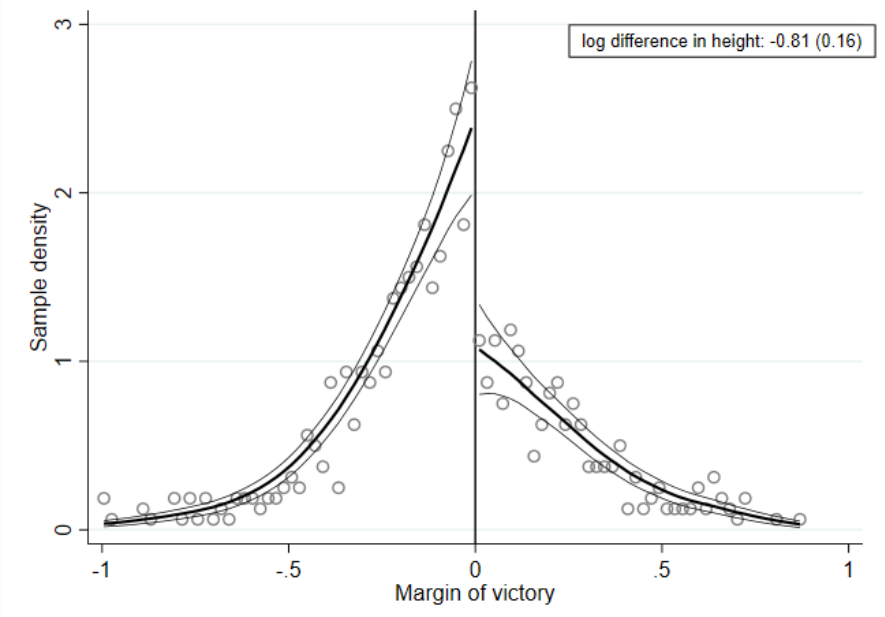

Notes: Panel A: N = 988, Panel B: N = 764. This figure presents estimates of shifts in win probabilities for gender-minority candidates in mixed-gender elections. Margin of victory is the difference between the gender-minority candidate's vote share and their gender-majority opponent's vote share (positive values indicate gender-minority victory while negative values indicate gender-minority loss). Total vote share is calculated only using votes received by the gender-minority and gender-majority candidates when one wins and the other loses. 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).

DECLARATION OF GENERATIVE AI AND AI-ASSISTED TECHNOLOGIES IN THE WRITING PROCESS

During the preparation of this work the author used Chat GPT in order to make grammatical corrections and suggested edits for clarity. After using this tool/service, the author reviewed and edited the content as needed and takes full responsibility for the content of the published article.

REFERENCES

Adams, B. (2018). Campaigning in Lilliput: Money's Influence in Small and Mid-Sized City Elections. California Journal of Politics and Policy, 10(2).

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.

Washington, E. (2012). Do majority-black districts limit blacks' representation? The case of the 1990 redistricting. *The Journal of Law and Economics*, *55*(2), 251-274.

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.

Gaudette, J., & de Benedictis-Kessner, J. (2024). Local Money: Evaluating the Effects of Municipal Campaign Contributions on Housing Policy Outcomes. 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). Did Board Gender Quotas Break the Glass Ceiling in Europe?. *AEA Papers and Proceedings*, *115*, 529-535.

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.