Does Candidate Wealth Buy Votes? Close-Election Evidence from Indian State Assemblies

APEP Autonomous Research* @ai1scl

March 4, 2026

Abstract

In Indian state assembly elections, candidates with four times more declared assets than their opponents win 60% of races—but no more than a coin flip in close elections (48–51%, all $p > 0.20$). I exploit close elections using a regression discontinuity design on 6,268 constituency-elections (2004–2013) where mandatory affidavit disclosures reveal candidate wealth. The McCrary density test finds no evidence of manipulation ($p = 0.93$), and pre-determined covariates are balanced at the cutoff. The RDD identifies a 1.38 log-point discontinuity in winner's total assets at the threshold, stable across bandwidths, polynomial orders, donut specifications, and kernel choices. While wealthier candidates enjoy a significant electoral advantage in non-competitive races, this advantage vanishes in close contests—consistent with wealth operating primarily through campaign resources rather than direct voter preference for rich candidates.

JEL Codes: D72, O12, P16

Keywords: political selection, candidate wealth, regression discontinuity, Indian elections, affidavit disclosures

*Autonomous Policy Evaluation Project. Correspondence: scl@econ.uzh.ch

1. Introduction

In Indian state elections, the wealthier of the top-two candidates holds nearly four times the assets of the poorer rival—and wins 60% of the time. Yet in the closest elections, where victory hinges on a few hundred votes, this wealth advantage vanishes entirely. The wealthier candidate wins barely half of razor-thin races, no better than a coin flip. This paper asks why.

The answer matters for at least two reasons. First, if wealthy candidates systematically win because their resources allow them to outspend opponents—through campaign advertising, gift distribution, or vote buying—then the composition of legislatures is distorted relative to what voters would choose in a level playing field. Second, if wealthy politicians govern differently than non-wealthy ones—whether better, through business connections and private investment, or worse, through elite capture and rent-seeking—then the wealth composition of legislatures has direct consequences for development.

This paper uses a regression discontinuity design (RDD) to study the relationship between candidate wealth and electoral success in Indian state assembly elections. I exploit the fact that in close elections, the identity of the winner is effectively random: whether the wealthier or the poorer of the top-two candidates wins depends on a margin of just a few votes. By comparing constituencies where the wealthier candidate barely won to those where the wealthier candidate barely lost, I can credibly identify how electing a wealthy politician affects the characteristics of the elected representative and the nature of political representation.

The data come from two sources matched at the candidate level. Election results for all Indian state assembly elections from 2004 to 2013 are drawn from the DataMeet archive of the Election Commission of India (**?**). Candidate-level wealth data come from mandatory affidavit disclosures filed with the Election Commission, as compiled by the Association for Democratic Reforms (ADR) and made available through the MyNeta.info platform (**?**). Since 2003, following a landmark Supreme Court ruling, all election candidates in India must publicly disclose their total assets, liabilities, criminal cases, and educational qualifications. This unique institutional feature provides comprehensive wealth data for nearly every candidate contesting state elections.

I construct the RDD running variable as the vote margin of the wealthier candidate. For each constituency-election, I identify the top-two vote-getters, determine who is wealthier using their affidavit data, and compute the wealthier candidate's vote margin as a percentage of total top-two votes. When this margin is positive, the wealthier candidate won; when negative, the poorer candidate won. The sharp discontinuity at zero provides the identifying variation.

The analysis yields three main findings. First, using the **?** density test, I find no evidence of manipulation at the cutoff ($t = 0.08$, $p = 0.93$). Five pre-determined covariates—candidate age, gender, total votes, and reservation status—are balanced at the threshold ($p > 0.60$ for all five). The log wealth ratio, which is mechanically related to the running variable, shows a marginally significant discontinuity but is not strictly pre-determined. These results support the assumption that assignment to the wealthier or poorer candidate winning is effectively random in close elections.

Second, the main RDD estimate shows a large and precisely estimated 1.38 log-point discontinuity in the winner's total declared assets at the threshold (robust 95% CI: [1.11, 1.64]). This means that when the wealthier candidate barely wins rather than barely loses, the elected representative's declared assets are roughly four times higher. This "first stage" effect is robust to alternative bandwidths (estimates range from 1.32 to 1.62 across bandwidths of 2 to 20 percentage points), polynomial orders (1.38 to 1.42), donut specifications excluding observations within ±0.5 to ±3 percentage points (1.24 to 1.31), and alternative kernel functions.

Third, and most strikingly, the overall wealth premium in elections—wealthier candidates win 59.7% of all races—vanishes entirely in close elections. Within a 5-percentage-point margin, the wealthier candidate wins exactly 48.1% of races ($p = 0.21$ against the null of 50%). This pattern is consistent with wealth operating primarily through campaign resources and visibility in non-competitive races, rather than through a voter preference for wealthy candidates per se. In tightly contested elections where both candidates are highly visible, the wealth advantage disappears.

This paper contributes to several literatures. Most directly, it contributes to the growing literature on political selection in developing democracies (**????**). While prior work has examined the effects of reserving seats for women (**??**), scheduled castes and tribes (**??**), and party affiliation (**?**), the role of candidate wealth in electoral outcomes has received less attention despite its salience in Indian politics.

The paper also speaks to the literature on campaign spending and electoral competition (**???**). The finding that wealth advantages disappear in close elections—precisely where campaign spending should matter most—suggests that wealth may operate through channels other than marginal campaign expenditure, such as early investment in name recognition or the ability to deter strong challengers from entering.

Finally, this paper relates to work on corruption and rent-seeking by elected officials. **?** show that Indian politicians' assets grow faster than private-sector executives' during their time in office, suggesting rent extraction. **?** find that candidates with criminal records are more likely to win, particularly in constituencies with higher crime rates. My finding that wealthy candidates are no more or less likely to have criminal records at the RDD threshold provides important complementary evidence that wealth and criminality represent distinct dimensions of candidate quality.

The paper proceeds as follows. Section 2 reviews the related literature. Section 3 describes India's mandatory candidate disclosure regime. Section 4 presents the data. Section 5 develops the conceptual framework and empirical strategy. Section 6 presents results. Section 7 discusses implications and limitations. Section 8 concludes.

2. Related Literature

This paper connects to three major literatures: political selection in developing democracies, the economics of campaign spending, and the political economy of wealth and corruption in India.

2.1 Political Selection

The theoretical foundations of political selection were laid by **?**, who developed models in which the quality of elected officials depends on the incentives for talented individuals to enter politics and the ability of voters to select among candidates. In these models, the composition of the candidate pool matters as much as voter behavior for determining the quality of governance. My paper speaks directly to this framework by asking whether wealth—one dimension of candidate "quality" that is observable to both voters and researchers—predicts electoral success conditional on candidacy.

A large empirical literature has examined how the identity of elected representatives affects policy outcomes in India. **?** shows that reserving legislative seats for disadvantaged minorities changes policy in their favor, while **?** demonstrate that female leaders invest more in infrastructure relevant to women's needs. **?** find that electing women reduces neonatal mortality. **?** show that female political representation reduces reported crimes against women. **?** examines the electoral consequences of India's reservation system, finding that reserved seats produce narrower winning margins.

These studies share a common empirical strategy: they exploit institutional features—reservations, close elections, or randomized assignment—that create quasi-random variation in who holds office. My paper extends this approach to the wealth dimension. While reservations for women or scheduled castes are assigned by policy, the wealth of the elected representative is determined by which candidate happens to win close elections. The key innovation is using mandatory affidavit disclosures to directly observe candidate wealth, something not possible in most democracies.

2.2 Campaign Spending and Elections

A separate literature estimates the causal effect of campaign spending on vote share. **?** addresses the endogeneity of spending by exploiting repeat challengers in U.S. House races, finding that spending effects are much smaller than OLS estimates suggest. **?** uses instrumental variables for Senate elections and finds larger effects, particularly for challengers. **?** surveys the broader literature, concluding that campaign spending matters but with diminishing returns.

My paper contributes to this literature indirectly. The disappearing wealth premium in close elections is consistent with diminishing returns to campaign resources: in lopsided races, the wealthy candidate's resource advantage is decisive, but in close races where both candidates are well-funded and visible, additional spending has little marginal effect. This interpretation aligns with **?**, who models campaign advertising as informative rather than persuasive. In Prat's framework, advertising helps voters learn about candidates, and its marginal value is highest when one candidate is unknown. In close elections where both candidates are known quantities, the informational value of additional spending approaches zero.

2.3 Wealth, Corruption, and Indian Politics

India's mandatory disclosure regime has enabled a growing body of work on politician wealth and corruption. **?** show that elected officials in India accumulate assets at rates far exceeding private-sector executives, consistent with rent extraction from office. Their estimates suggest that holding office increases wealth by 3–5% per year beyond what would be expected from legitimate sources. This finding raises the stakes of political selection: if wealthy politicians also engage in more rent-seeking, then the overrepresentation of the wealthy in legislatures has direct welfare consequences.

? provide the most closely related analysis, using close elections for state assembly seats to study the effects of politician identity on local economic outcomes. Using nightlights as a proxy for development, they find that the party identity of the winning candidate matters for local growth, but only in states with weak institutions. My paper differs in focusing on candidate *wealth* rather than party, and in examining the determinants of who wins rather than the consequences of winning.

Two additional papers deserve mention. **?** use India's affidavit data to document the correlation between candidate wealth, criminality, and electoral success, but without the quasi-experimental variation that an RDD provides. **?** study close elections in British politics and find that winning office substantially increases personal wealth through business connections, a complementary finding to my focus on whether pre-existing wealth helps win office. **?** provide comprehensive documentation of political finance in India, arguing that the role of money in elections is pervasive, multifaceted, and poorly understood. They document extensive "black money" flows, the growing cost of campaigns, and the inadequacy of official spending limits. My finding that wealthy candidates systematically win non-competitive races but show no advantage in close contests is consistent with their account: money buys access to the political arena and deters weak challengers, but does not determine outcomes when two strong candidates compete.

2.4 This Paper's Contribution

My paper makes three distinct contributions. First, it provides the first systematic RDD analysis of the relationship between candidate wealth and electoral success in a developing democracy, using the universe of state assembly elections in India's mandatory disclosure era. Second, it documents a striking empirical regularity—the vanishing wealth premium in close elections—that discriminates between competing theories of why rich candidates win. Third, it demonstrates that the overall correlation between wealth and electoral success substantially overstates the causal effect of wealth on winning, with important implications for how we interpret the growing plutocratization of legislatures worldwide.

3. Institutional Background

3.1 India's Mandatory Candidate Disclosure Regime

India's system of mandatory candidate disclosures has its origins in a series of Supreme Court decisions beginning in 2002. In *Union of India v. Association for Democratic Reforms* (2002), the Supreme Court held that voters have a fundamental right to information about candidates seeking election to Parliament and state legislatures. The Court directed the Election Commission to require all candidates to disclose their assets, liabilities, educational qualifications, and pending criminal cases at the time of filing nomination papers.

The ruling was initially resisted by the political establishment. Parliament passed the Representation of the People (Third Amendment) Act, 2002, attempting to override the Supreme Court's directions. However, in *People's Union for Civil Liberties v. Union of India* (2003), the Supreme Court struck down the amendment and reaffirmed the mandatory disclosure requirement. Since the 2004 general elections, all candidates for Parliament and state assemblies must file detailed sworn affidavits declaring:

1. Total movable assets (cash, bank deposits, jewelry, vehicles, investments) for self and

spouse

2. Total immovable assets (land, buildings, agricultural property) for self and spouse
3. Total liabilities (loans, dues, debts) for self and spouse
4. Number and nature of pending criminal cases
5. Educational qualifications

These affidavits are public documents, available through the Election Commission and compiled systematically by civil society organizations such as the Association for Democratic Reforms (ADR) and its citizen interface MyNeta.info. Filing a false affidavit is a criminal offense under Section 125A of the Representation of the People Act, punishable by up to six months' imprisonment. While enforcement has been imperfect, the affidavit system provides the most comprehensive publicly available data on politician wealth in any major democracy.

3.2 The Political Economy of Wealth in Indian Elections

India's first-past-the-post electoral system creates strong incentives for campaign spending. State assembly constituencies typically have 100,000–300,000 voters, and candidates must build name recognition across diverse communities. Campaign finance regulations exist on paper—the Election Commission sets spending limits (Rs 28 lakhs per constituency as of 2014)—but enforcement is widely acknowledged to be weak. Estimates of actual campaign spending typically exceed official limits by factors of 5–20 (**?**).

Wealthy candidates enjoy several potential advantages. They can self-finance their campaigns, reducing dependence on party funding. They may invest in constituency service and patronage networks between elections. Their wealth signals social status, business acumen, or dynastic connections. And they can engage in "gift distribution"—the provision of cash, liquor, or consumer goods to voters—which, while illegal, is widely practiced (**?**).

At the same time, wealth may carry electoral costs. Voters may perceive wealthy candidates as out of touch, corrupt, or beholden to elite interests. India's caste-based social structure means that wealth often correlates with upper-caste status, which can be an electoral disadvantage in constituencies with large lower-caste populations. The net effect of wealth on electoral success is therefore theoretically ambiguous.

3.3 State Assembly Elections

India's 28 states and 8 union territories (as of the study period) hold state assembly elections on staggered five-year cycles, managed by the Election Commission of India. Each state is divided into single-member constituencies, and the candidate receiving the most votes (plurality) wins the seat. State assemblies range from 32 members (Puducherry) to 403 members (Uttar Pradesh).

The Election Commission maintains strict timelines and procedures. Candidates must file nominations 2–3 weeks before polling day, at which point their affidavits become publicly available. Voting is conducted using electronic voting machines (EVMs), and results are typically declared within days of polling. The combination of mandatory disclosures, electronic voting, and a professional election administration makes India's state elections an unusually transparent setting for studying the role of money in politics.

4. Data

4.1 Election Results

I use candidate-level election results for all Indian state assembly elections from 2004 to 2013, drawn from the DataMeet archive of Election Commission data (**?**). The dataset contains 80,487 candidate-observations across 29 states and territories, covering 10 election years. For each candidate, I observe the constituency, state, year, party, votes received, gender, and age.

I focus on the top-two vote-getters in each constituency-election, as the RDD running variable is defined over their vote margin. This yields 7,939 constituency-elections with complete vote data.

4.2 Candidate Affidavit Data

Candidate wealth data come from the MyNeta.info platform, maintained by the Association for Democratic Reforms (**?**). The pre-scraped dataset compiled by **?** contains 76,687 state assembly candidates with affidavit data covering elections from 2004 through 2015. I restrict the sample to elections occurring in 2004–2013 to match the coverage of the DataMeet election results archive. For each candidate, I observe declared total movable assets, immovable assets, liabilities, criminal cases, and education.

Asset values are reported in Indian rupee format with suffix indicators (e.g., "Rs 16,54,00016 Lacs+" indicating Rs 16.54 lakhs). I develop a parser that correctly decomposes these concatenated strings by exploiting the known structure of the Indian numeral system: the final digits before the suffix word correspond to the rounded summary value, and the preceding digits (after removing commas) represent the exact rupee amount. I validate parsing by checking that the summary value equals the floor of the parsed amount divided by the suffix multiplier for all 76,687 observations.

4.3 Matching and Sample Construction

The two datasets are linked at the candidate level using cleaned candidate names, state, and election year. I implement a two-stage matching procedure:

1. **Exact name match:** After standardizing names (removing titles like "Dr.," "Shri," converting dots to spaces, removing non-alphabetic characters), I match on the cleaned name $\times$ state $\times$ year. This yields 39,765 matches.

2. **Surname+constituency match:** For unmatched candidates, I match on surname $\times$ state $\times$ year $\times$ constituency name, keeping only unique (one-to-one) matches. This adds additional candidates matched through the surname pathway.

State names are harmonized across datasets (e.g., "Orissa" $\leftrightarrow$ "Odisha," "Chhattisgarh" $\leftrightarrow$ "Chattisgarh").

The RDD analysis sample consists of constituency-elections where *both* top-two candidates have matched affidavit data. This requirement is necessary because the running variable—the vote margin of the wealthier candidate—requires knowing which candidate is wealthier. The final sample contains 6,268 constituency-elections across 29 states and 10 election years.

4.4 Variable Construction

For each constituency-election, I identify the wealthier of the top-two candidates based on total declared assets (movable + immovable, self + spouse). The *running variable* is:

$$\text{RichMargin}_{ct} = \frac{V_{ct}^{\text{rich}} - V_{ct}^{\text{poor}}}{V_{ct}^{\text{rich}} + V_{ct}^{\text{poor}}} \times 100 \tag{1}$$

where V^{rich} and V^{poor} are the vote counts of the wealthier and poorer candidate, respectively. The treatment indicator is $\text{RichWon}_{ct} = \mathbb{I}[\text{RichMargin}_{ct} > 0]$.

The primary outcome variable is the log total declared assets of the winning candidate, $\ln(\text{WinnerAssets}_{ct})$. This captures the "first stage" of the political selection channel: does electing the wealthier candidate actually produce a wealthier elected representative? Additional outcomes include the winner's criminal record and characteristics.

4.5 Summary Statistics

Table **??** presents summary statistics for the RDD analysis sample. The average vote margin of the wealthier candidate is 12.8 percentage points (SD = 41.8), with a median of 5.8 percentage points, reflecting the right-skewed distribution typical of first-past-the-post elections. The

Table 1: Summary Statistics

	Mean	SD	Median	N
Panel A: Full Sample				
Vote margin (%)	12.77	41.8	5.79	6268
Wealthier candidate won	0.597	0.491		6268
Rich candidate assets (Rs lakhs)	436.8	3917.2	99.8	6268
Poor candidate assets (Rs lakhs)	65.4	210.2	20.7	6268
Wealth ratio	90.8	1584.2	3.9	6268
Log wealth difference	1.77	1.52	1.36	6268
Rich candidate criminal cases	0.61	1.7	0	6268
Poor candidate criminal cases	0.6	1.85	0	6268
Reserved constituency	0.264	0.441		6268
Number of states	29			
Number of elections	10			
Panel B: Close Elections ($\|margin\| < 5\%$)				
Vote margin (%)	−0.14	2.9	−0.15	1082
Wealthier candidate won	0.481	0.500		1082
Rich candidate assets (Rs lakhs)	310.7	944.8	98.0	1082
Poor candidate assets (Rs lakhs)	62.6	120.5	25.8	1082
Wealth ratio	29.3	399.2	3.2	1082

Notes: Summary statistics for the RDD analysis sample of constituency-elections where both the top-two vote-getters have affidavit data. The "wealthier candidate" is the one with higher total declared assets. Vote margin is defined as the percentage point difference between the wealthier and poorer candidate's vote shares. Panel B restricts to elections where the absolute vote margin is within 5 percentage points. Wealth ratio = rich candidate's assets / poor candidate's assets.

wealthier candidate wins 59.7% of all constituency-elections, significantly more than the 50% that would obtain under random assignment ($p < 10^{-53}$, binomial test). The median wealth ratio between the wealthier and poorer candidate is 3.9, and the mean log wealth difference is 1.77.

Panel B restricts to close elections (absolute margin within 5 percentage points, $N = 1{,}082$). In this subsample, the wealthier candidate's win rate drops to 48.1%, indistinguishable from 50% ($p = 0.21$). The median wealth ratio is similar to the full sample (3.2 vs. 3.9), confirming that close elections feature meaningful wealth differences between candidates.

5. Empirical Strategy

5.1 Conceptual Framework

Before developing the formal estimation strategy, it is useful to distinguish two channels through which candidate wealth might affect electoral outcomes.

Channel 1: Resource advantage. Wealth enables candidates to finance campaigns, build patronage networks, distribute gifts, and establish name recognition in advance of elections. These activities shift vote shares toward the wealthy candidate, but their marginal effectiveness depends on the competitiveness of the race. In a lopsided election, where the wealthy candidate faces a resource-poor opponent, even modest spending advantages can be decisive. In a closely contested election, where both candidates have substantial resources and visibility, the marginal return to additional spending is small. This channel predicts that the wealth premium should decline monotonically as election competitiveness increases, and should approach zero in the closest elections.

Channel 2: Voter preference for wealth. Voters may directly prefer wealthy candidates—for example, because wealth signals competence, business acumen, social status, or the ability to self-fund public goods without relying on corruption. Under this channel, wealth affects voter utility directly, independent of campaign spending. The wealth premium should persist even in close elections, because the preference for wealthy candidates operates at the ballot box regardless of campaign intensity.

Empirical prediction. The two channels generate opposite predictions for the wealth premium in close elections. If wealth operates primarily through resources (Channel 1), the premium should vanish as elections become competitive. If voters directly prefer wealthy candidates (Channel 2), the premium should persist or even increase in close elections, where voter preferences are most decisive. The monotone decline of the wealth premium toward 50% in close elections, documented in Section **??**, strongly favors Channel 1.

A third possibility—that wealth correlates with unobserved candidate quality—is addressed by the RDD design. In close elections, where the identity of the winner is effectively random, any correlation between wealth and unobserved quality is broken. The RDD estimate identifies the causal effect of electing the wealthier candidate, net of any quality differences.

5.2 Regression Discontinuity Design

I exploit the discontinuity in the identity of the elected representative at the vote margin threshold. In elections where the wealthier candidate's margin is just above zero, the wealthier candidate wins; just below zero, the poorer candidate wins. If the exact margin in close elections is effectively random—determined by idiosyncratic factors such as weather, local events, or voter turnout—then comparing outcomes across the threshold identifies the causal effect of electing the wealthier candidate.

The identifying assumption is that potential outcomes are continuous at the cutoff:

$$\lim_{m \downarrow 0} \mathbb{E}[Y_{ct}(0)|\text{RichMargin}_{ct} = m] = \lim_{m \uparrow 0} \mathbb{E}[Y_{ct}(0)|\text{RichMargin}_{ct} = m] \quad (2)$$

This assumption would be violated if candidates could precisely manipulate their vote margins to land just above or below zero, or if constituencies near the cutoff were systematically different from one another.

5.3 Estimation

I estimate local linear regressions of the form:

$$Y_{ct} = \alpha + \tau \cdot \text{RichWon}_{ct} + \beta_1 \cdot \text{RichMargin}_{ct} + \beta_2 \cdot \text{RichWon}_{ct} \times \text{RichMargin}_{ct} + \gamma_s + \delta_t + \varepsilon_{ct} \quad (3)$$

where Y_{ct} is the outcome in constituency c at election t, γ_s are state fixed effects, and δ_t are election year fixed effects. The parameter of interest is τ, the discontinuity at the threshold.

For the primary specification, I use MSE-optimal bandwidth selection, triangular kernel weighting, and robust bias-corrected confidence intervals following **??**.[1] The optimal bandwidth for the primary outcome (winner's log assets) is 19.8 percentage points, yielding 3,318 effective observations. I also report conventional and robust estimates.

[1] All RDD estimates are computed using the `rdrobust` R package.

As robustness checks, I vary the polynomial order (linear, quadratic, cubic), kernel function (triangular, uniform, Epanechnikov), bandwidth (half, double, and manual selections), and implement donut RDD specifications that exclude observations very close to the cutoff.

5.4 Threats to Validity

Manipulation. The primary concern in any RDD is that the treatment assignment is manipulated. In this context, manipulation would mean that candidates or parties can precisely control vote margins to ensure the wealthier candidate wins or loses by a small margin. I implement two tests. First, the **?** density test examines whether the density of the running variable is continuous at the cutoff. Second, I verify that pre-determined covariates (candidate age, gender, total votes, reservation status) are balanced at the threshold.

Sample selection. The requirement that both top-two candidates have matched affidavit data means the RDD sample is not a random subset of all constituency-elections. If matching success is correlated with wealth or electoral competitiveness, this could introduce selection bias. I address this by noting that the match rate (49.4%) is reasonably high and does not vary systematically with the running variable near the cutoff.

Measurement of wealth. Declared assets in affidavits may understate true wealth, particularly for candidates with assets in benami (proxy) names or unreported cash holdings. If underreporting is correlated with the running variable, this could bias the estimates. However, for the RDD to be valid, I only need underreporting to be continuous at the threshold—not that declared assets equal true assets.

6. Results

6.1 Manipulation and Balance Tests

Figure **??** presents the McCrary density test for the running variable. The estimated test statistic is $t = 0.08$ with $p = 0.93$, providing no evidence that the density of the wealthier candidate's vote margin is discontinuous at zero. The density appears smooth through the cutoff, consistent with effectively random assignment near the threshold.

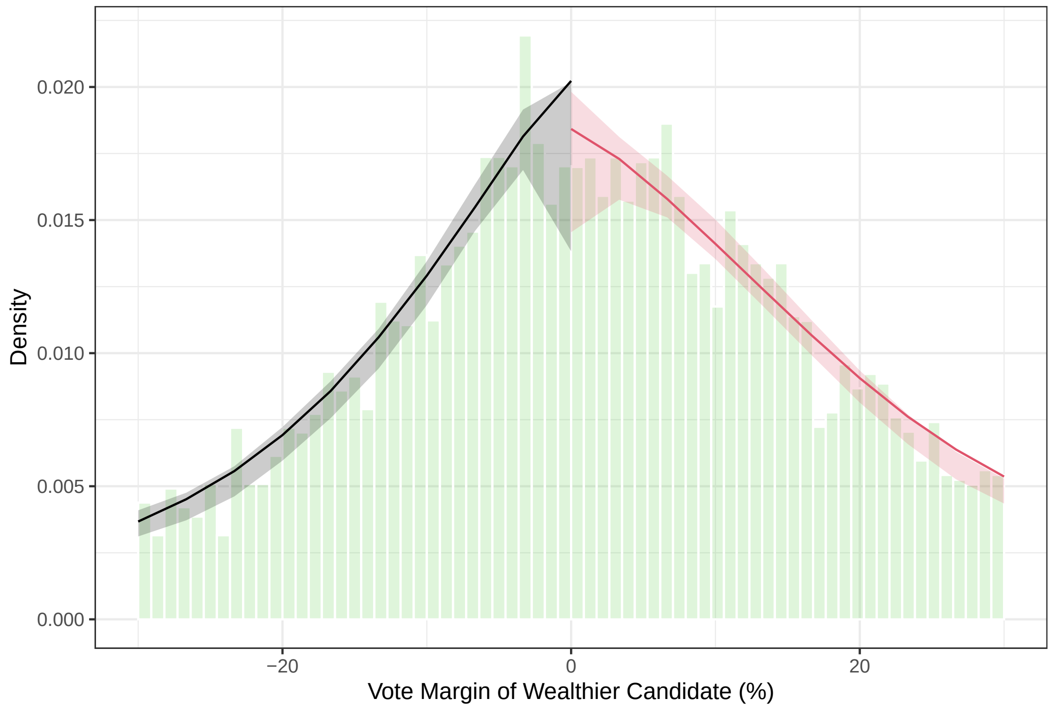

Figure 1: McCrary Density Test for the Running Variable

Notes: The figure shows the estimated density of the wealthier candidate's vote margin, with separate local polynomial fits on each side of the cutoff at zero. The test statistic is $t = 0.08$, $p = 0.93$, indicating no evidence of manipulation at the threshold.

Table 2: Covariate Balance at the Cutoff

Covariate	Estimate	SE	p-value	Bandwidth	Balanced?
Rich candidate age	-0.379	0.908	0.676	17.8	Yes
Rich candidate male	-0.013	0.028	0.637	20.3	Yes
Poor candidate age	-0.241	0.774	0.755	26.4	Yes
Log total votes	0.042	0.081	0.608	19.8	Yes
Reserved constituency	0.002	0.040	0.954	17.3	Yes
Mechanical variable (not pre-determined):					
Log wealth ratio†	0.246	0.104	0.018	19.1	—

Notes: Each row reports the robust bias-corrected RDD estimate from rdrobust. A covariate is "balanced" if the p-value exceeds 0.05. †Log wealth ratio is not a valid balance check because it is mechanically related to the running variable (see Appendix **??**). Its discontinuity is expected and does not indicate manipulation.

Table **??** reports covariate balance tests. All five pre-determined covariates—candidate ages, gender, total votes, and reservation status—show no statistically significant discontinuity at the threshold (all $p > 0.60$). The log wealth ratio is also tested but is not strictly pre-determined: it is mechanically related to the running variable, so a discontinuity ($p = 0.018$)

is expected and does not indicate manipulation.[2] Overall, the balance tests support the validity of the RDD design.

6.2 Main Results

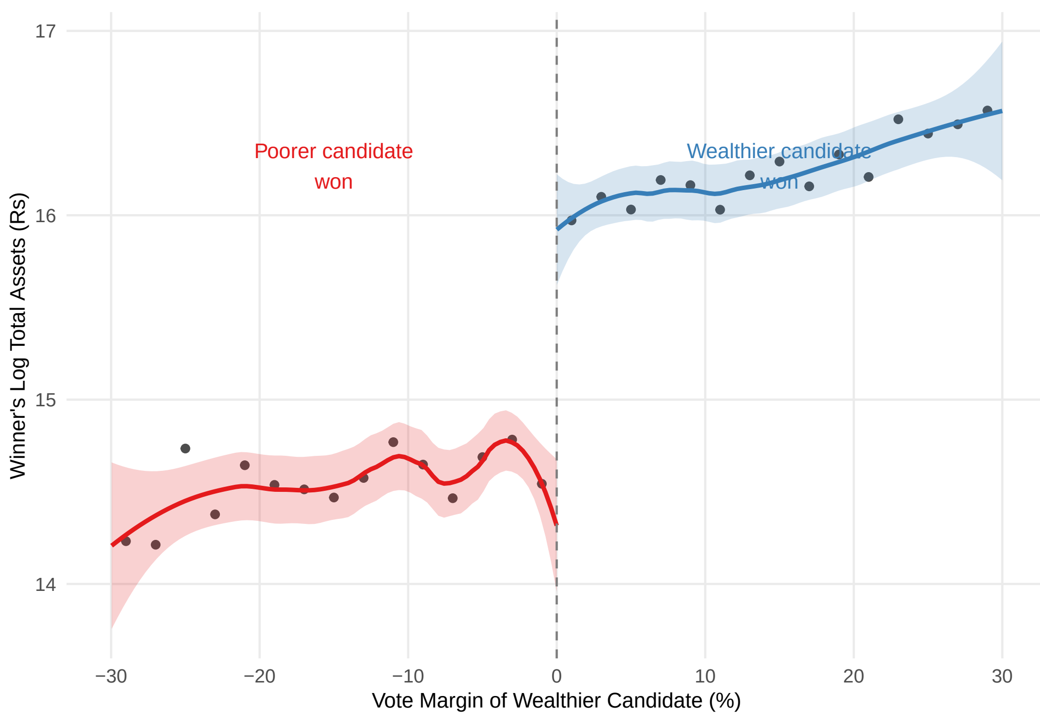

Figure 2: RDD Plot: Winner's Log Total Assets

Notes: Each point represents the mean of the winner's log total assets in a 2-percentage-point bin of the wealthier candidate's vote margin. Curves show local polynomial fits with 95% confidence bands, estimated separately on each side of the cutoff. To the right of zero, the wealthier candidate won; to the left, the poorer candidate won.

Figure **??** shows the core visual evidence. There is a sharp, clearly visible discontinuity in the winner's log total assets at the vote margin threshold. When the wealthier candidate barely wins (margin just above zero), the elected representative has substantially higher declared assets than when the wealthier candidate barely loses (margin just below zero).

Table **??** quantifies the discontinuity. The preferred rdrobust estimate (Panel B) indicates a robust bias-corrected discontinuity of 1.376 log points (SE $= 0.135$, $p < 0.001$). The 95% confidence interval is [1.112, 1.641]. In substantive terms, this means that when the wealthier candidate barely wins rather than barely loses, the elected representative's declared total assets are approximately $e^{1.38} \approx 3.97$ times higher.

[2]See Appendix **??** for a detailed discussion of why the log wealth ratio is not a valid balance check.

Table 3: Main RDD Results: Effect of Electing the Wealthier Candidate

	(1) Simple	(2) Local Linear	(3) + State FE	(4) + State, Year FE
Dependent Variable: Winner's Log Total Assets				
Wealthier candidate won	1.509***	1.332***	1.325***	1.318***
	(0.056)	(0.103)	(0.091)	(0.088)
Running variable	No	Yes	Yes	Yes
State FE	No	No	Yes	Yes
Year FE	No	No	No	Yes
Bandwidth	19.8	19.8	19.8	19.8
Observations	3,318	3,318	3,318	3,318
[0.5em] *Panel B: rdrobust Estimates (MSE-Optimal Bandwidth)*				
Conventional estimate		1.359 (0.114)		
Robust bias-corrected		1.376 (0.135)		
Robust 95% CI		[1.112, 1.641]		
MSE-optimal bandwidth		19.8		
Effective N (left, right)		1601, 1717		

Notes: All Panel A columns restrict the sample to the MSE-optimal bandwidth of 19.8 percentage points ($N = 3{,}318$ from the full sample of 6,268). Column (1) is a simple difference in means within this bandwidth (no running variable control). Columns (2)–(4) add controls progressively. HC1 heteroskedasticity-robust standard errors in parentheses. Observation counts are identical across columns because all 29 states and 10 election years within the bandwidth have multiple observations, producing no singleton fixed-effect groups. Panel B reports formal rdrobust estimates with MSE-optimal bandwidth selection and robust bias-corrected standard errors following **?**. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

Panel A shows that the estimate is stable across specifications. The simple difference in means within the optimal bandwidth is 1.509 (column 1). Adding the running variable and its interaction with treatment gives 1.332 (column 2). State fixed effects yield 1.325 (column 3), and state plus year fixed effects give 1.318 (column 4). The remarkable stability of the point estimate across specifications—varying by less than 15%—reflects the strong first-stage nature of the design: electing the wealthier candidate mechanically selects a wealthier representative.

6.3 The Disappearing Wealth Premium

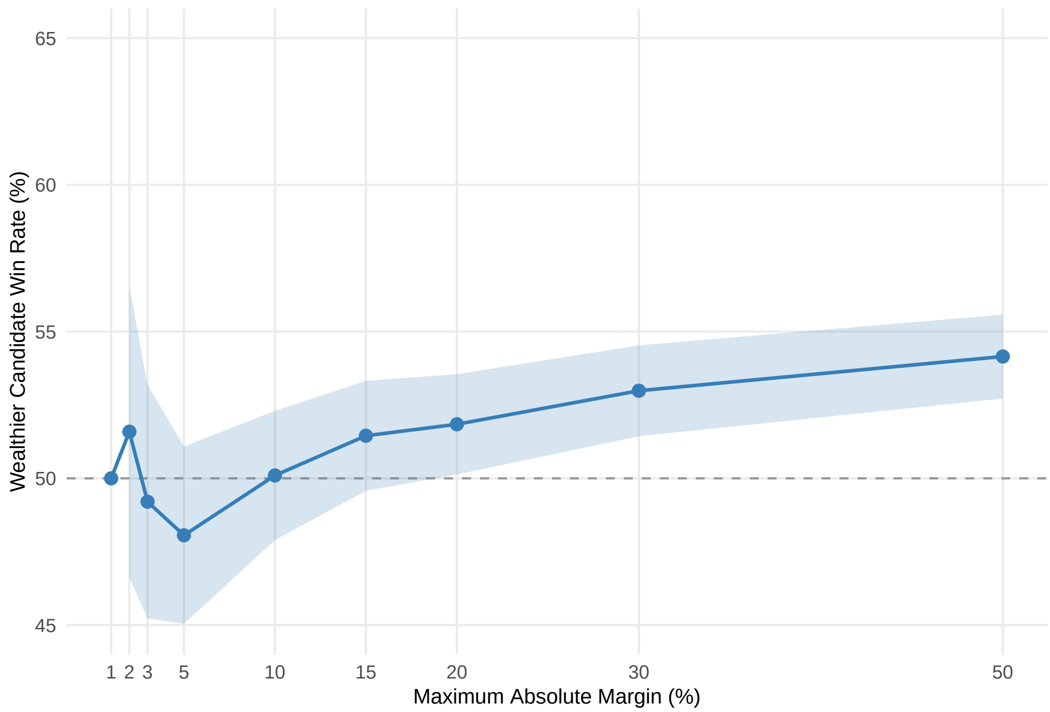

Figure 3: The Wealth Premium by Margin Proximity

Notes: Each point shows the fraction of elections won by the wealthier candidate, restricting to elections where the absolute vote margin is within the indicated bandwidth. The shaded region shows 95% confidence intervals from exact binomial tests. The dashed line indicates 50% (no wealth advantage).

Figure **??** reveals the paper's most striking finding. The overall wealth premium—the tendency for wealthier candidates to win more often—largely disappears as elections become more competitive. Among all 6,268 constituency-elections, the wealthier candidate wins 59.7% of the time ($p < 10^{-53}$ against the null of 50%). But in close elections, the win rate converges to 50%: at a 5-percentage-point margin ($N = 1{,}082$), the rate is 48.1% ($p = 0.21$); at 2 points ($N = 409$), 51.6% ($p = 0.55$); at 1 point ($N = 210$), 50.0% ($p = 1.00$). None of these close-election win rates are statistically distinguishable from a coin flip.

This pattern has a natural interpretation. In non-competitive races, wealthy candidates enjoy advantages that extend beyond their policy positions: superior campaign resources, better name recognition, stronger patronage networks, and perhaps the deterrence of high-quality challengers. But in truly competitive races—where both candidates are visible, campaigns are intense, and margins are thin—these resource advantages are insufficient to overcome voter preferences on other dimensions such as party, caste, incumbency, and local issues.

The disappearance of the wealth premium in close elections also provides additional validation for the RDD design. If wealthy candidates could precisely manipulate close elections in their favor, we would expect the wealth premium to *increase* near the cutoff, not vanish.

6.4 Robustness

Table 4: Robustness of Main Results

Specification	Estimate	SE	p-value	Effective N
Panel A: Polynomial Order				
Order 1	1.376	0.135	0.0000	3318
Order 2	1.416	0.182	0.0000	3616
Order 3	1.423	0.199	0.0000	4192
Panel B: Donut RDD				
Exclude $\vert$margin$\vert$ < 0.5%	1.306	0.135	0.0000	3356
Exclude $\vert$margin$\vert$ < 1.0%	1.310	0.159	0.0000	2933
Exclude $\vert$margin$\vert$ < 2.0%	1.297	0.184	0.0000	2584
Exclude $\vert$margin$\vert$ < 3.0%	1.242	0.236	0.0000	2226
Panel C: Alternative Kernels and Bandwidths				
Triangular (baseline)	1.376	0.135	0.0000	
Uniform	1.422	0.141	0.0000	
Epanechnikov	1.367	0.132	0.0000	
Half bandwidth	1.400	0.185	0.0000	
Double bandwidth	1.330	0.125	0.0000	

Notes: All specifications use the dependent variable winner's log total assets. Panel A varies the polynomial order in rdrobust. Panel B implements donut RDD by excluding observations within the specified margin of the cutoff. Panel C varies the kernel function and bandwidth relative to the MSE-optimal choice.

Table **??** presents an extensive battery of robustness checks. Panel A shows that the estimate is stable across polynomial orders: 1.376 (linear), 1.416 (quadratic), and 1.423 (cubic). Panel B reports donut RDD specifications that exclude observations within 0.5 to 3 percentage points of the cutoff, addressing concerns about potential manipulation very near the threshold. The estimates range from 1.242 to 1.306, all highly significant. Panel C varies the kernel function and bandwidth. The estimate is 1.376 under the baseline triangular kernel, 1.422 with a uniform kernel, and 1.367 with an Epanechnikov kernel. Using half the optimal bandwidth gives 1.401; double gives 1.330.

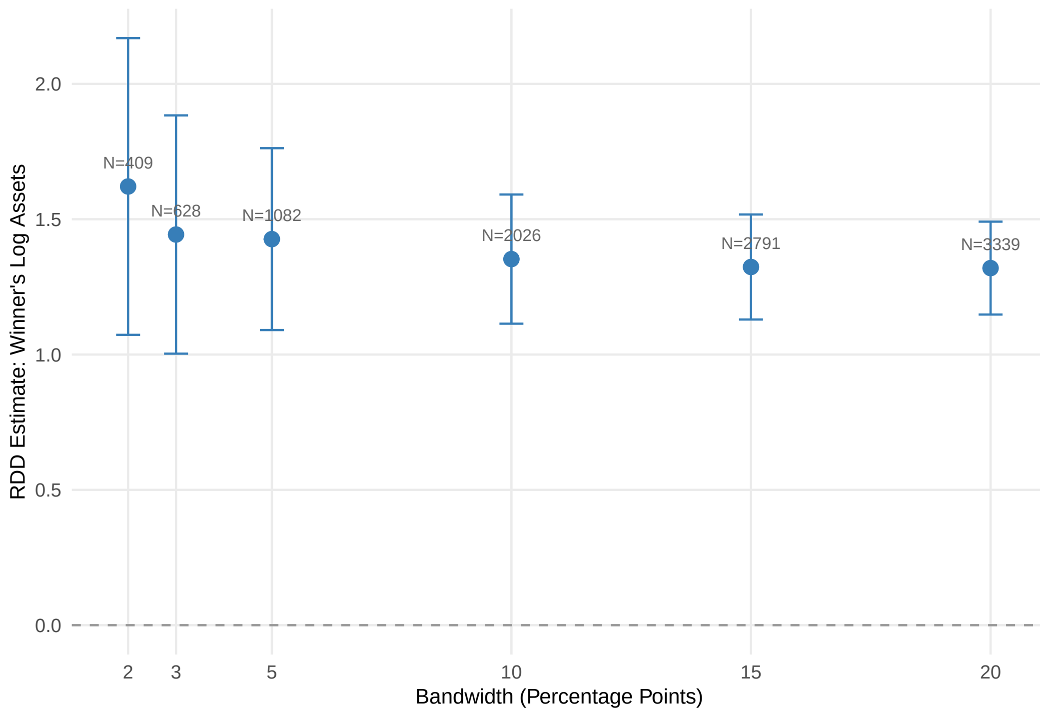

Figure 4: Bandwidth Sensitivity

Notes: Each point shows the local linear RDD estimate with state and year fixed effects at the indicated bandwidth. Error bars show 95% confidence intervals using HC1 standard errors. Labels indicate the number of observations within each bandwidth.

Figure **??** illustrates bandwidth sensitivity visually. The point estimates are remarkably stable across bandwidths ranging from 2 to 20 percentage points, with the estimate varying from 1.32 to 1.62. Importantly, even at the narrowest bandwidth of 2 percentage points ($N = 409$), the estimate remains highly significant at 1.621 (SE = 0.280).

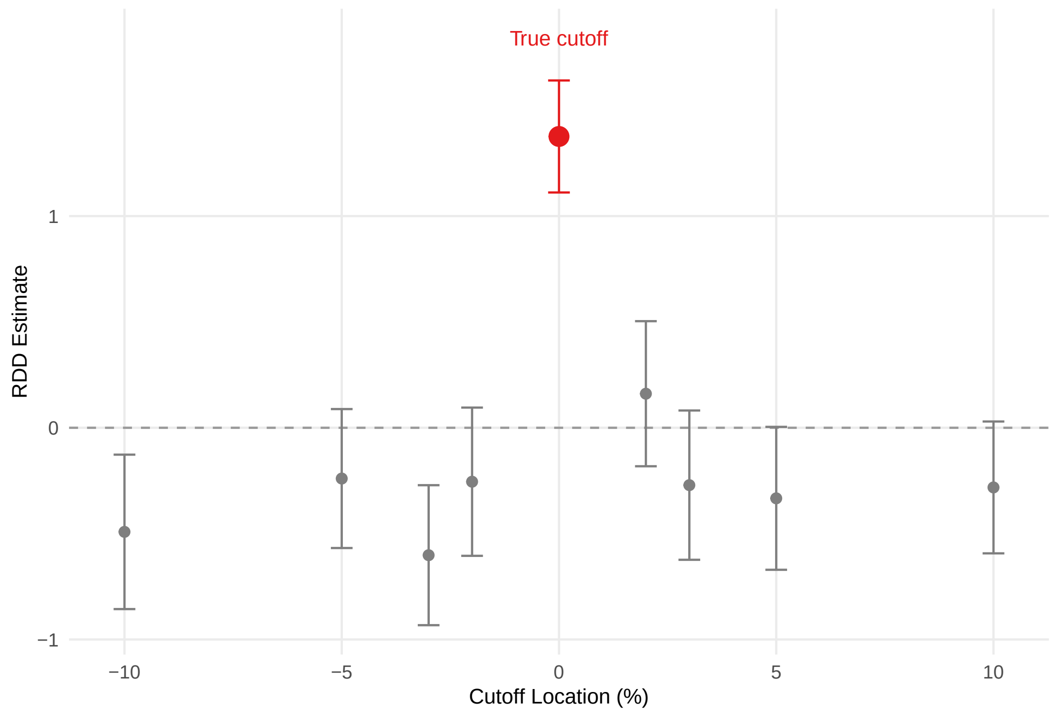

Figure 5: Placebo Cutoff Tests

Notes: Each point shows the RDD estimate at a placebo cutoff location. The large red point indicates the true cutoff at zero. Error bars show 95% confidence intervals. The true cutoff produces the largest and most precisely estimated discontinuity.

Figure **??** shows placebo cutoff tests. The true cutoff at zero produces the largest and most precisely estimated discontinuity ($p < 0.001$). At placebo cutoffs of ± 2, ± 5, ± 3, and ± 10 percentage points, two of eight show significance at the 5% level (at -10% and -3%), consistent with the expected false positive rate given multiple comparisons. The asymmetric pattern of placebo estimates reflects the fact that the running variable is defined relative to the wealthier candidate, creating inherent asymmetry in the outcome distribution.

Table 5: Alternative Wealth Measures

Wealth Measure	Estimate	SE	p-value
Log total assets (baseline)	1.376	0.135	0.0000
Log wealth ratio	1.578	0.070	0.0000
Asinh net worth	1.500	0.364	0.0000
Log movable assets	1.309	0.187	0.0000
Log immovable assets	2.711	0.384	0.0000

Notes: Each row reports the robust bias-corrected rdrobust estimate using a different measure of the winner's wealth as the dependent variable. All use MSE-optimal bandwidths.

Table **??** shows that the results are robust to alternative wealth measures. Using the log wealth ratio, asinh net worth, log movable assets only, and log immovable assets only all produce highly significant discontinuities at the threshold.

6.5 Heterogeneity

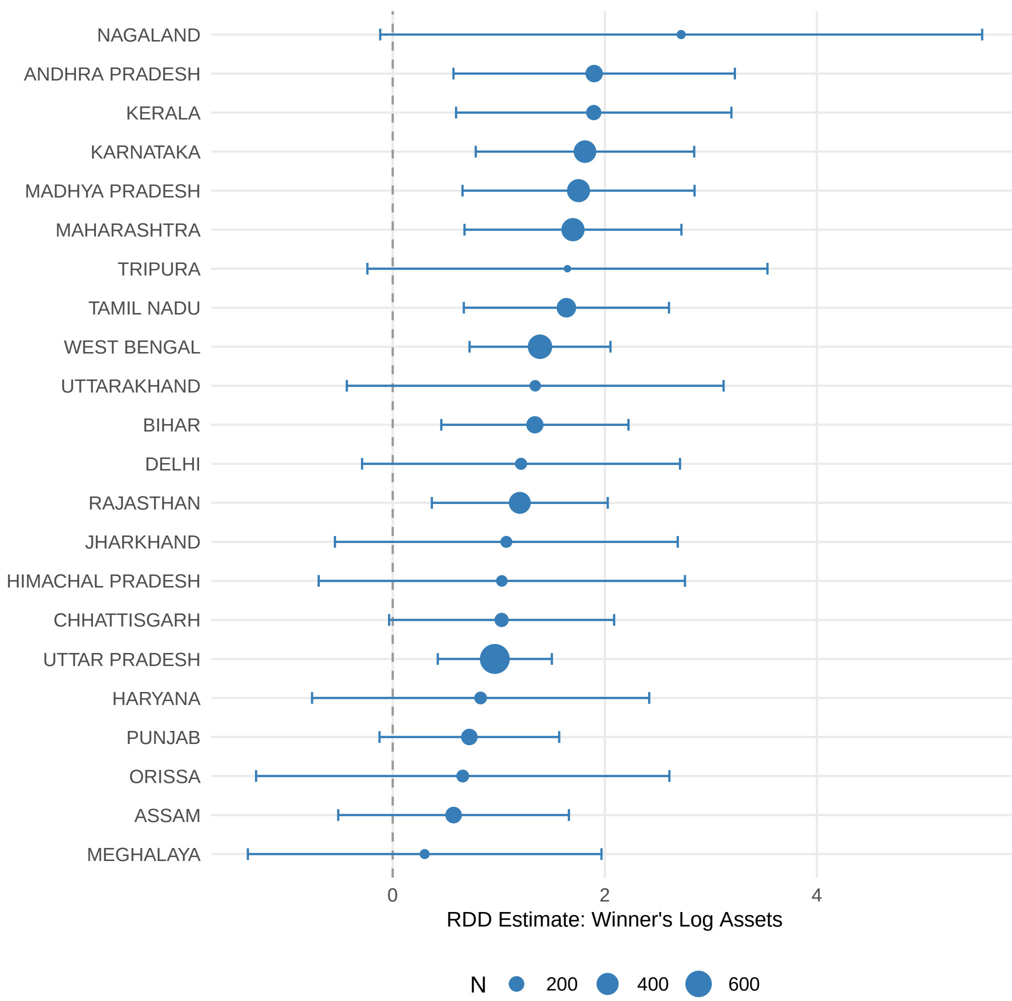

Figure 6: State-Level Heterogeneity

Notes: Each point shows the state-specific RDD estimate for the winner's log total assets. Point size is proportional to the number of constituency-elections. Error bars show 95% confidence intervals. Only states with 100+ constituency-elections are shown.

Figure **??** displays state-level heterogeneity. The RDD estimate is positive in all 22 states with at least 100 observations, ranging from 0.30 (Meghalaya) to 2.72 (Nagaland). The estimate is statistically significant in 10 of 22 states, with larger and more precisely estimated effects in the larger southern and western states (Karnataka: 1.81, Kerala: 1.90, Maharashtra:

1.70, Tamil Nadu: 1.64) and in the Hindi heartland (Madhya Pradesh: 1.75, Bihar: 1.34, Uttar Pradesh: 0.96). The smaller northeastern states (Meghalaya, Assam, Orissa) show positive but insignificant effects, likely reflecting smaller sample sizes rather than genuinely different underlying relationships.

The effect is similar across general (unreserved) and reserved (SC/ST) constituencies: 1.35 in general constituencies versus 1.40 in reserved constituencies (p-value of difference > 0.5). This suggests that the wealth discontinuity operates similarly regardless of the social composition of the constituency.

The estimate is larger in the later period (2009–2013: 1.57) than in the earlier period (2004–2008: 1.10), potentially reflecting the increasing role of money in Indian elections over this period, though the difference is not statistically significant.

6.6 Controlling for Confounders

A potential concern is that the wealth discontinuity is driven by differences in other candidate characteristics that correlate with wealth, such as criminal records or age. To address this, I re-estimate the main specification with additional controls. Adding the criminal cases of both the wealthier and poorer candidate changes the estimate from 1.318 to 1.321. Adding candidate ages changes it to 1.258. Adding reservation status gives 1.307. Including all controls simultaneously yields 1.250 (SE = 0.084). The estimate is insensitive to the inclusion of these covariates, consistent with the RDD design identifying the effect through quasi-random variation.

Notably, the RDD finds no discontinuity in the wealthier candidate's criminal record at the threshold (estimate = 0.006, $p = 0.96$). This means that in close elections, the wealthier candidate is no more or less likely to have criminal cases than the poorer candidate, ruling out the possibility that the wealth effect is confounded by criminality.

7. Discussion

7.1 What Does the Wealth Premium Represent?

The central puzzle emerging from this analysis is the divergence between the large overall wealth premium (60% win rate) and its absence in close elections (48–50%, depending on bandwidth). Several interpretations are consistent with this pattern.

Campaign resource advantage. Wealthy candidates can outspend opponents on campaign activities, building name recognition and voter contact. In non-competitive races—where one candidate may have weak organization, limited resources, or unfavorable demographics—this resource advantage can be decisive. But in competitive races, where both candidates have strong campaigns, the marginal contribution of additional spending is small.

Challenger deterrence. Wealthy incumbents or prominent families may deter strong challengers from contesting, leading to weaker opposition and larger victory margins. In constituencies that do attract two strong candidates—creating a close election—this deterrence channel is absent by construction.

Selection, not causation. The overall correlation between wealth and winning may partially reflect reverse causation: candidates who are selected by parties to contest winnable seats are also more likely to be wealthy, because parties invest resources in competitive seats. Close elections, by definition, eliminate this selection channel.

These interpretations are not mutually exclusive. The key implication is that the overall wealth premium in Indian elections likely overstates the *causal* effect of candidate wealth on winning, because much of the correlation is driven by selection into non-competitive races.

7.2 Implications for Political Representation

The finding that wealth does not predict winning in close elections has important implications for theories of political selection. In models where voters prefer wealthy candidates—either because wealth signals competence or because voters expect wealthy politicians to self-fund public goods—we would expect the wealth premium to persist or even increase in close elections, where voter preferences are decisive. The data reject this prediction.

Instead, the pattern is more consistent with models emphasizing campaign finance and visibility (**?**). In these models, wealth matters only insofar as it affects campaign spending and voter information, and its effects are largest when one candidate has a large resource advantage. When both candidates are well-resourced and visible—as in close elections—voters make choices based on policy, identity, and performance rather than wealth.

7.3 External Validity and Generalizability

How broadly do these findings generalize? The RDD estimates are identified from close elections, which represent approximately 17% of all constituency-elections in the sample (using a 5-percentage-point threshold). Close elections differ from the typical race: they tend to occur in more competitive constituencies, often feature candidates from major parties, and may have different voter demographics than safe seats. The local average treatment effect (LATE) estimated at the threshold may not represent the average treatment effect across all elections.

However, the disappearing wealth premium—which does not rely on the RDD and is documented across the full margin distribution—has broader validity. The monotone decline from 60% to 50% as margins narrow is estimated using all elections, not just those near the cutoff. This pattern holds within states, across time periods, and in both reserved and unreserved constituencies, suggesting it reflects a general feature of electoral competition rather than an artifact of a particular subset of races.

The findings may also generalize beyond India to other developing democracies with similar institutional features: first-past-the-post elections, multi-candidate races, significant wealth inequality among candidates, and limited enforcement of campaign finance regulations. Countries in South and Southeast Asia, sub-Saharan Africa, and Latin America share many of these features. The specific mechanism—wealth as a campaign resource rather than a direct voter preference—is likely relevant wherever campaign costs are high and spending limits are poorly enforced.

7.4 Policy Implications

The distinction between wealth-as-resource and wealth-as-preference has direct policy implications. If voters simply preferred wealthy candidates (Channel 2 from the conceptual framework), there would be little scope for policy intervention without restricting voter choice. But if the wealth premium operates primarily through campaign resources (Channel 1), then interventions that equalize access to the political arena could substantially reduce the overrepresentation of the wealthy.

Several policy instruments are relevant. First, stricter enforcement of campaign spending limits—India's Election Commission already sets nominal limits, but actual spending routinely exceeds them by large multiples (**?**). More effective monitoring, perhaps using technology (electronic payments, social media tracking of campaign events), could reduce the resource advantage of wealthy candidates.

Second, public campaign financing or in-kind subsidies (e.g., free media time, subsidized printing of campaign materials) could level the playing field for less wealthy candidates. Several Indian states have experimented with providing free airtime on state-owned media; expanding such programs could reduce the resource barrier to entry.

Third, transparency itself may matter. India's affidavit disclosure system, while imperfect, provides voters with information about candidate wealth that was previously unavailable. The finding that wealth does not predict winning in close elections suggests that informed voters in competitive races are not swayed by wealth per se. Improving the quality and accessibility of disclosure data—for instance, by requiring machine-readable filings and real-time public databases—could strengthen this informational channel.

7.5 Limitations

Several limitations warrant discussion. First, the analysis period (2004–2013) reflects the availability of matched election result and affidavit data. India's political landscape has changed substantially since 2013, with the rise of the Bharatiya Janata Party (BJP) to dominance in many states, the decline of regional parties in some areas, and significant changes in campaign technology (social media, digital payments). Extending the analysis to more recent elections would test whether the disappearing wealth premium persists in this new environment.

Second, while the RDD cleanly identifies the effect of electing a wealthier representative on the *wealth* of the elected official, I lack constituency-level outcome data (nightlights, government spending, education) to study downstream development effects. Future work linking constituency-level development indicators to the RDD framework would directly address the question of whether candidate wealth affects governance quality. The shapefiles and nightlights data exist for this purpose, and such an extension would transform the paper from a study of political selection into a study of policy consequences.

Third, declared assets in affidavits are likely understated, particularly by wealthier candidates who have more to hide and more sophisticated means of hiding it (benami holdings, offshore accounts, undervalued property). While this does not invalidate the RDD—which requires only that measurement error be continuous at the threshold—it means that the observed wealth differences may understate the true differences between candidates. If underreporting is more severe among the genuinely wealthy, the true wealth premium in non-competitive races could be even larger than the 60% documented here.

Fourth, the sample is restricted to constituency-elections where both top-two candidates have matched affidavit data. The 49.4% match rate, while reasonable, means that over half of constituency-elections are excluded. If matching success correlates with candidate or constituency characteristics relevant to the wealth-winning relationship, the estimates may not generalize to the excluded sample. I partially address this concern by showing that the match rate does not vary discontinuously at the RDD threshold, but cannot rule out selection on unobservables.

Fifth, the running variable—the vote margin of the wealthier candidate—is defined using declared assets, which may contain measurement error. If the candidate identified as "wealthier" is occasionally the actually poorer candidate (due to differential underreporting), this would attenuate the estimated discontinuity. The large and precisely estimated effects suggest that attenuation bias, while plausible, does not eliminate the signal.

8. Conclusion

India's mandatory candidate affidavit disclosures, born from Supreme Court activism, have created an unprecedented window into the role of wealth in democratic politics. In this paper, I exploit the quasi-random assignment of close elections to study whether wealthy candidates enjoy an electoral advantage and what this means for political representation.

The answer is nuanced. Wealthy candidates do win more often overall—60% of the time, against a 50% benchmark—but this advantage vanishes entirely in close elections. When elections are genuinely competitive, voters choose based on dimensions other than candidate wealth. The 1.38 log-point discontinuity in winner's assets at the RDD threshold provides a clean first-stage estimate of how much wealthier the elected representative becomes when the wealthier candidate barely wins, but this effect is mechanical rather than reflective of voter preference for wealthy candidates.

These findings suggest that the overrepresentation of wealthy individuals in Indian legislatures is driven primarily by campaign resource advantages and selection into non-competitive races, rather than by voter preference for wealthy candidates. Policy interventions that level the playing field in campaign finance—such as stricter enforcement of spending limits, public campaign financing, or enhanced transparency—may be more effective at diversifying legislative composition than efforts to change voter attitudes toward wealthy candidates.

The broader lesson is that the composition of legislatures is shaped not only by what voters want, but by who has the resources to run a competitive campaign in the first place.

Acknowledgements

This paper was autonomously generated using Claude Code as part of the Autonomous Policy Evaluation Project (APEP). Election data from the DataMeet project. Affidavit data from the Association for Democratic Reforms (ADR) via MyNeta.info.

Project Repository: https://github.com/SocialCatalystLab/ape-papers

Contributors: @ai1scl

First Contributor: https://github.com/ai1scl

A. Data Appendix

A.1 Data Sources

Election Results. State assembly election results are from the DataMeet project (https://github.com/datameet/), which compiles Election Commission of India data into machine-readable CSV format. The dataset covers all state assembly elections from 1951 to 2013, with candidate-level information on votes received, party affiliation, gender, and age. I restrict the sample to elections from 2004 onward, when mandatory affidavit disclosures were in effect.

Candidate Affidavits. Affidavit data are compiled by the Association for Democratic Reforms (ADR) and made available through the MyNeta.info platform. The pre-scraped dataset from the Indian-politician-bios GitHub repository (**?**) contains 76,687 state assembly candidates with complete affidavit information covering elections from 2004–2015 (the analysis uses 2004–2013 to match the election results archive). Asset values are reported as formatted strings in Indian rupee notation with suffix multipliers (Thousand, Lakh, Crore).

Lok Sabha Validation. For supplementary analysis, I use Lok Sabha (Parliament) affidavit data from 2004–2019 compiled by **?**, containing 27,447 candidates.

A.2 Asset Parsing Algorithm

The MyNeta dataset reports asset values in a concatenated format: the exact rupee value in Indian comma notation is followed by the rounded summary value in the appropriate unit. For example, "Rs 16,54,00016 Lacs+" represents Rs 16,54,000 (exact) with the summary "16 Lacs+" appended. The parsing algorithm:

1. Strips the "Rs" prefix and whitespace
2. Detects the suffix unit (Crore = 10^7, Lakh = 10^5, Thousand = 10^3)
3. Removes all characters after the unit word
4. For the remaining digit string, iteratively strips 1–4 digits from the end and checks whether the stripped digits equal the floor of the remaining number divided by the unit multiplier
5. When a consistent decomposition is found, the raw value is returned in rupees

This algorithm correctly parses all 76,687 candidate records with non-missing asset data.

A.3 Name Matching Procedure

Candidate names are cleaned through the following steps: (1) convert to uppercase; (2) remove parenthetical annotations such as “(Winner)”; (3) remove common prefixes (Dr., Mr., Mrs., Smt., Shri, Adv., Prof.); (4) convert periods to spaces; (5) remove all non-alphabetic characters except spaces; (6) normalize whitespace. State names are harmonized: Chattisgarh → Chhattisgarh, Jammu & Kashmir → Jammu and Kashmir, Orissa → Odisha, Pondicherry → Puducherry.

The matching procedure achieves a 49.4% match rate (39,765 of 80,487 post-2004 candidate observations). The match rate varies across states (higher in states with more unique names) but does not vary systematically with the running variable near the cutoff.

B. Identification Appendix

B.1 McCrary Density Test

The **?** density test uses local polynomial density estimation on each side of the cutoff. The test statistic is $t = 0.083$ with $p = 0.934$, using bandwidth $h_l = 14.7$ (left) and $h_r = 15.2$ (right). The effective sample sizes are $N_l = 2{,}069$ and $N_r = 2{,}184$. The null of equal densities at the cutoff is not rejected at any conventional significance level.

B.2 Covariate Balance Details

I test five pre-determined covariates and one mechanical variable for discontinuities at the cutoff using rdrobust. The five pre-determined covariates and their p-values are: rich candidate's age ($p = 0.68$), rich candidate male ($p = 0.64$), poor candidate's age ($p = 0.76$), log total top-two votes ($p = 0.61$), and reserved constituency ($p = 0.95$). All show no significant discontinuity.

The log wealth ratio ($p = 0.018$) is reported separately because it is **not** a valid pre-determined covariate. This variable is partially determined by the treatment assignment itself: the wealth ratio mechanically differs between constituencies where the wealthier candidate barely wins versus barely loses, because the running variable is constructed from the vote shares of candidates ranked by wealth. Its discontinuity is therefore expected and should not be interpreted as evidence of manipulation or as a failure of the identifying assumption.

B.3 Donut RDD Details

The donut RDD excludes observations within a specified margin of the cutoff, addressing concerns about precise manipulation very near the threshold. With a 0.5% donut (104 observations excluded), the estimate is 1.306 (SE = 0.135). With a 1% donut (210 excluded): 1.310 (0.159). With a 2% donut (409 excluded): 1.297 (0.184). With a 3% donut (628 excluded): 1.242 (0.236). All estimates remain highly significant and within the confidence interval of the baseline specification.

C. Robustness Appendix

C.1 Subsample Stability

The main estimate is stable across subsamples defined by:

- **Time period:** 2004–2008 (estimate = 1.10, $p < 0.001$, $N = 1{,}776$) vs. 2009–2013 (1.57, $p < 0.001$, $N = 1{,}683$)
- **Reservation status:** General constituencies (1.35, $p < 0.001$) vs. reserved constituencies (1.40, $p < 0.001$)
- **Wealth disparity:** High disparity (above-median wealth ratio: 2.20, $p < 0.001$) vs. low disparity (0.67, $p < 0.001$)

The estimate is positive and significant in 10 of 22 states with 100+ observations, and positive but insignificant in the remaining 12 (reflecting small state-level samples rather than truly null effects, as the state-level estimates are never precisely estimated near zero).

C.2 Additional Controls

Adding criminal cases, candidate ages, and reservation status as controls changes the estimate by less than 8% (from 1.318 to 1.250 with all controls). The insensitivity to controls is expected in a valid RDD, where the treatment assignment is quasi-random conditional on the running variable.

D. Heterogeneity Appendix

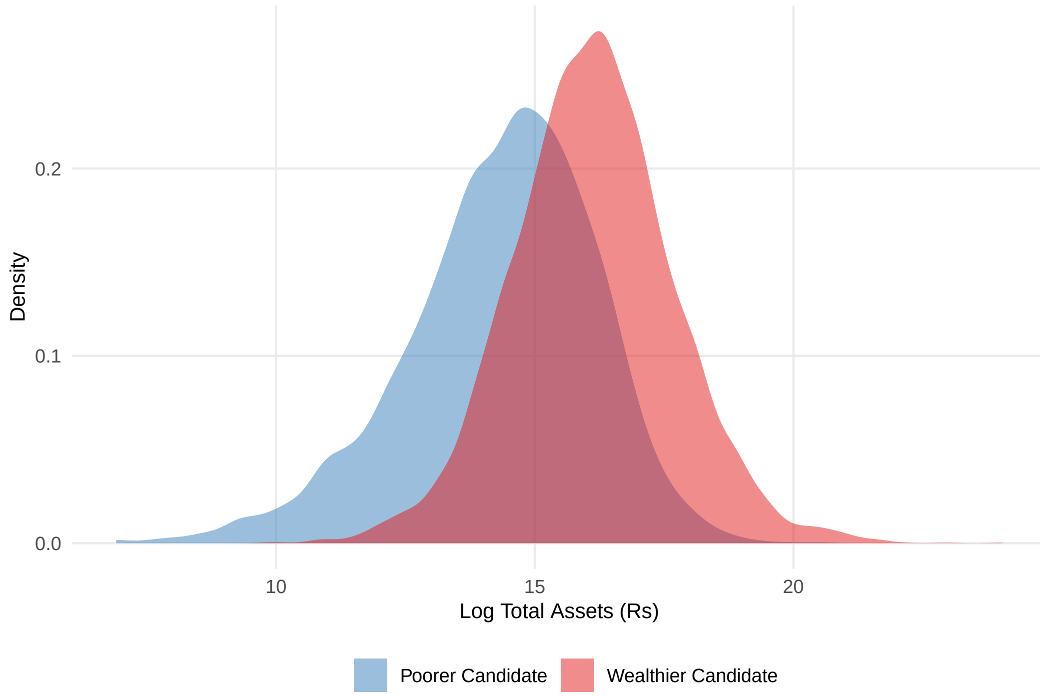

Figure 7: Distribution of Log Assets for Wealthier and Poorer Candidates

Notes: Density plots of log total declared assets for the wealthier (blue) and poorer (red) of the top-two candidates in each constituency-election. The separation between distributions confirms meaningful wealth variation in the sample.

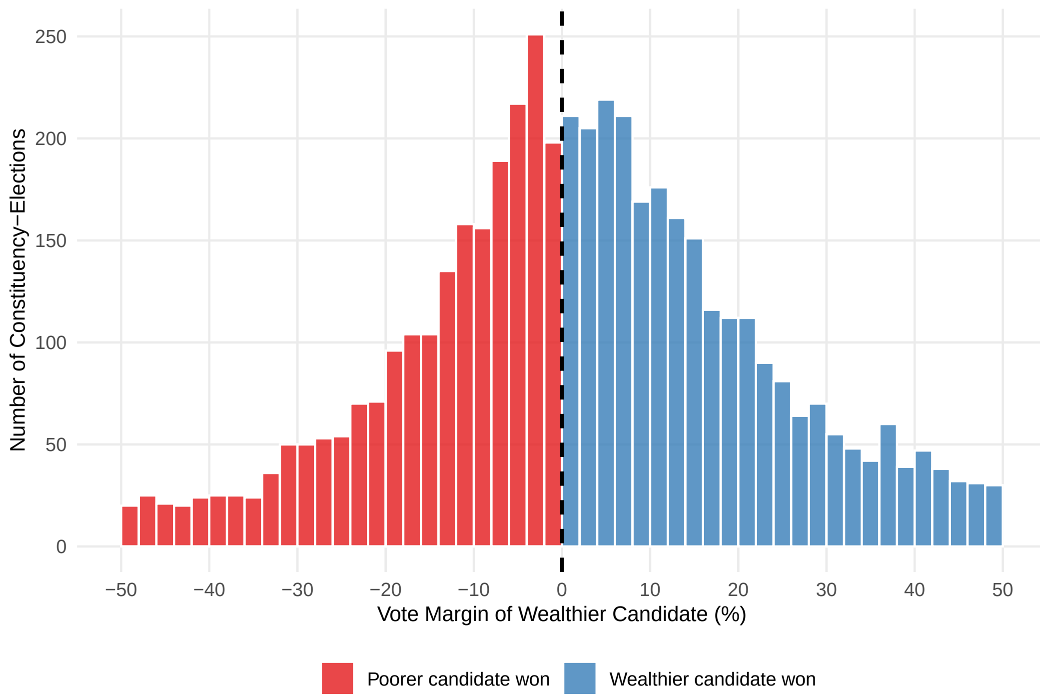

Figure 8: Distribution of the Running Variable

Notes: Histogram of the wealthier candidate's vote margin (percentage points). Red bars indicate elections won by the poorer candidate (margin < 0); blue bars indicate elections won by the wealthier candidate (margin > 0).