

Do Close Referendum Losses Demobilize Voters? Evidence from Swiss Municipal Voting

APEP Autonomous Research*

Anonymous

@anonymous

January 31, 2026

Abstract

Does experiencing a local referendum loss affect subsequent voter turnout? I exploit the unique setting of Swiss direct democracy, where municipalities vote on federal referendums that may pass nationally despite local opposition. Using a regression discontinuity design at the 50% vote share threshold across 2,122 municipalities and 56 federal referendums from 2010–2019 (with outcome data extending through 2022), I compare subsequent turnout in municipalities that narrowly “lost” (voted against a passing measure) versus those that narrowly “won” (voted for a passing measure). I find no evidence that local referendum losses affect subsequent turnout: the RDD estimate is 0.05 percentage points ($SE = 0.84$, $p = 0.95$). The McCrary density test shows no evidence of manipulation ($p = 0.92$), and the null result is robust across bandwidth choices, polynomial specifications, and placebo cutoffs. This finding suggests that Swiss direct democracy is resilient to “sore loser” demobilization effects—voters continue participating regardless of past local outcomes. The result contributes to the policy feedback literature by demonstrating that repeated referendum experience, even when locally unsuccessful, does not erode democratic participation.

JEL Codes: D72, H70, P16

Keywords: voter turnout, direct democracy, referendum, regression discontinuity, Switzerland

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

1. Introduction

Do citizens become disengaged from democracy when they lose? The question of how winning and losing affects political participation is central to understanding democratic stability. A substantial literature documents the “winner-loser gap” in political attitudes: citizens who supported the winning candidate or party in an election report higher political satisfaction, trust, and democratic legitimacy than those on the losing side (Anderson et al., 2005; Blais and Gélinau, 2017). But whether these attitude differences translate into behavioral consequences—specifically, whether losing demobilizes future participation—remains less understood.

The stakes of this question extend beyond academic interest. In an era of increasing political polarization and declining trust in institutions, understanding whether electoral losses compound into democratic disengagement is vital for institutional design. If losing systematically reduces participation, then democratic systems may face a self-reinforcing cycle in which today’s losers become tomorrow’s non-participants, further marginalizing their interests and deepening disengagement. Conversely, if democracies can sustain participation among losers, they preserve the voice and representation of minority viewpoints essential to pluralistic governance.

This question is particularly salient in direct democracy, where citizens vote directly on policy questions rather than candidates. Switzerland provides the world’s most intensive experience with direct democracy: Swiss voters participate in 3–4 federal referendums per year on issues ranging from constitutional amendments to popular initiatives. Unlike representative elections, where losing means your preferred candidate is not in office, losing a referendum means your preferred policy was rejected. This raises the possibility that referendum losses could feel more directly consequential and thus more demobilizing.

The psychology of referendum losses differs from candidate election losses in important ways. In candidate elections, voters may rationalize losses by attributing blame to the candidate’s personal characteristics, campaign strategy, or idiosyncratic factors. In referendums, by contrast, the choice is starkly policy-based: voters lose because their policy preference was rejected by their fellow citizens. This directness may heighten feelings of being “out of step” with the majority, potentially amplifying demobilization. Alternatively, the policy focus may reduce personal stakes—one can oppose a policy while maintaining faith in democratic processes—potentially mitigating demobilization.

Two competing theoretical predictions emerge from the political behavior literature. The *demobilization hypothesis* suggests that losing—especially narrowly—reduces democratic efficacy beliefs and subsequent participation. Voters who “almost won” but lost may conclude

that the system is unresponsive to their preferences, leading to withdrawal. This prediction draws on the extensive literature on political efficacy, which documents that citizens who doubt their political influence are less likely to participate (Campbell et al., 1954; Finkel, 1985). Narrow losses may be particularly demobilizing because they make the counterfactual victory salient: “if only a few more people had agreed with me, we would have won.”

The *mobilization hypothesis*, by contrast, suggests that narrow losses galvanize opposition. Voters who “almost won” may be motivated to turn out more strongly in future related votes to finally achieve their policy goals. This prediction draws on theories of grievance mobilization and political entrepreneurship: close losses may signal that victory is achievable with renewed effort, spurring rather than suppressing future engagement (McAdam, 1982). Political organizations may also be more effective at mobilizing supporters after near-victories, using the “we almost won” narrative to motivate turnout.

A third possibility—the *habituation hypothesis*—suggests that in high-frequency referendum environments, voters become accustomed to both winning and losing, and neither outcome has persistent effects on participation. Swiss voters face referendum ballots multiple times per year, typically voting on 10–15 federal proposals annually. In this environment, winning and losing are regular occurrences rather than rare, high-stakes events. Voters may come to view referendum outcomes as part of an ongoing democratic conversation rather than definitive verdicts, reducing the psychological impact of any single loss.

I test these hypotheses using a regression discontinuity design (RDD) that exploits the geographic variation in referendum outcomes across Swiss municipalities. When a federal referendum passes nationally, some municipalities voted for it (local “winners”) while others voted against it (local “losers”). Within the set of referendums that passed, I compare municipalities that narrowly voted in favor (just above 50% yes) to those that narrowly voted against (just below 50% yes). Because municipalities near the 50% threshold should be similar in all respects except their winning/losing experience, this design provides quasi-random variation in the psychological experience of winning versus losing.

The RDD approach offers several advantages over alternative identification strategies. First, unlike survey-based studies that rely on self-reported vote choice and participation intentions, this design uses administrative data on actual voting behavior. Second, the design addresses selection concerns that plague observational comparisons of winners and losers: municipalities that systematically vote against passing referendums may differ in unobserved ways from those that vote in favor, but municipalities *near* the 50% threshold should be comparable. Third, the municipal-level aggregation provides statistical power while still capturing meaningful political communities—Swiss municipalities are important units of political identity and information transmission.

The key finding is a precise null: local referendum losses have no detectable effect on subsequent voter turnout. The RDD estimate is 0.05 percentage points with a cluster-robust standard error of 0.84, yielding a p-value of 0.95. The 95% confidence interval of [-1.60, 1.70] rules out economically meaningful effects in either direction. The McCrary density test shows no evidence of manipulation at the cutoff ($p = 0.92$), and the result is robust across bandwidth choices (0.5–1.5 times optimal), polynomial orders (1 or 2), and kernel specifications (triangular, uniform, Epanechnikov). Placebo tests at fake cutoffs (30%, 40%, 60%, 70%) show no effects, as expected under a valid design.

This null result is substantively important for three reasons. First, it suggests that Swiss direct democracy is resilient to “sore loser” effects. Despite the high frequency of referendums, losing does not lead to cumulative disengagement. Second, the finding contrasts with concerns about “democratic fatigue”—the worry that frequent voting opportunities might erode participation over time. The stability of turnout regardless of past outcomes suggests that the institution itself, rather than individual vote experiences, sustains engagement. Third, the result informs debates about referendum design: the experience of losing a vote need not be institutionally costly if it does not demobilize future participation.

The precise null also has methodological implications. The confidence interval rules out effects as small as 1.7 percentage points in either direction. With baseline turnout of approximately 50%, this means we can rule out effects larger than approximately 3.4% of baseline participation. Studies of get-out-the-vote interventions typically find effects of 1–5 percentage points for intensive treatments (Green and Gerber, 2008); we can rule out effects of this magnitude from the “natural experiment” of referendum losses.

This paper contributes to several literatures. To the winner-loser gap literature, I provide the first causal evidence on behavioral (turnout) effects in the referendum context, complementing existing work on attitudinal effects (Marien and Kern, 2018; Bowler et al., 2022). Prior work has focused almost exclusively on attitudes—satisfaction, trust, legitimacy beliefs—leaving open whether these attitude differences translate into behavioral consequences. The disconnect between attitudes and behavior documented here has important implications: democratic systems may be more robust than attitudinal studies suggest.

To the policy feedback literature, I document a null feedback effect, suggesting that winning and losing experiences do not generate self-reinforcing or self-undermining dynamics (Soss, 1999; Mettler and Soss, 2004). The policy feedback framework emphasizes that policy experiences shape political engagement: positive experiences increase participation while negative experiences decrease it. The null finding here suggests that referendum outcomes—at least at the local level—do not function as powerful feedback mechanisms.

To the RDD methodological literature, I adapt the close-election design—typically applied

to candidate races—to the referendum setting, demonstrating its applicability to direct democracy (Lee, 2008; Caughey and Sekhon, 2011). The application raises interesting conceptual questions: in candidate elections, the “treatment” is typically holding office, a concrete institutional position. In referendums, the “treatment” is the psychological experience of winning or losing, a more subjective construct. The null finding suggests either that this psychological treatment is ineffective or that the municipal-level aggregation obscures individual-level effects.

The remainder of the paper proceeds as follows. Section 2 describes Switzerland’s direct democracy system and the logic of the identification strategy. Section 3 presents the data, including summary statistics on municipal referendum outcomes. Section 4 details the empirical strategy, including the RDD specification, validity tests, and robustness checks. Section 5 presents results. Section 6 discusses interpretations, limitations, and implications for democratic institutions.

2. Institutional Background

Switzerland operates the world’s most intensive system of direct democracy. Federal referendums are held four times per year (March, June, September, and November) on fixed dates determined years in advance. Citizens vote on constitutional amendments (mandatory referendums), citizen-initiated constitutional changes (popular initiatives), and optional referendums challenging federal legislation. Between 1981 and 2019, over 700 federal referendums took place, a frequency unmatched by any other nation.

This section provides institutional context necessary to understand the research design. I describe the referendum procedure, the role of municipalities in Swiss political life, the information environment surrounding referendum outcomes, and the theoretical mechanisms through which local outcomes might affect subsequent participation.

2.1 Federal Referendum Procedure

Federal referendums follow a standardized procedure designed to ensure informed participation. Approximately six weeks before the vote, eligible citizens receive voting materials by mail, including the ballot, an official information booklet (*Erläuterungen des Bundesrates*), and a prepaid return envelope. The booklet contains the full text of the measure, the Federal Council’s recommendation, and arguments from both supporters and opponents—typically political parties, interest groups, and civil society organizations. The balanced presentation of arguments is a legal requirement, ensuring that voters receive information from both sides regardless of the government’s position.

Citizens may vote by mail (used by approximately 90% of voters in recent years), at a polling station on voting Sunday, or—in pilot cantons—electronically. The shift toward mail voting, which accelerated in the 1990s and 2000s, has increased convenience and may have contributed to the relatively stable turnout rates observed over this period. Mail voting extends over several weeks, allowing citizens time to read materials and make deliberate choices.

For a federal referendum to pass, it must achieve both a *popular majority* (more than 50% of valid votes nationwide) and, for constitutional amendments, a *cantonal majority* (more than half of the 26 cantons voting in favor). When counting cantonal votes, each of the six “half-cantons” (Obwalden, Nidwalden, Basel-Stadt, Basel-Landschaft, Appenzell Ausserrhoden, Appenzell Innerrhoden) counts as half a vote. This dual requirement creates a federal balance, ensuring that constitutional changes reflect broad geographic support rather than simple population majorities.

The types of referendums differ in their origins and requirements:

Mandatory referendums are required for any amendment to the federal constitution and for joining international organizations. These referendums are initiated by the government or parliament and require both popular and cantonal majorities.

Popular initiatives are citizen-initiated constitutional amendments. To qualify for the ballot, an initiative must gather 100,000 valid signatures within 18 months. Initiatives are often used by groups outside the political mainstream to place issues on the agenda. Historically, most initiatives fail—the passage rate is approximately 10%—but even failed initiatives can influence policy by signaling public sentiment.

Optional referendums challenge federal legislation passed by parliament. If opponents gather 50,000 signatures within 100 days of a law’s publication, the law is submitted to a popular vote. Optional referendums require only a popular majority (no cantonal majority). They serve as a “popular veto” on parliamentary action.

2.2 Municipal Voting and Geographic Variation

Swiss referendum results are reported at the municipal (*Gemeinde*) level, providing fine-grained geographic variation. As of 2019, Switzerland has approximately 2,140 municipalities, ranging from Zurich (435,000 residents) to tiny Alpine communities with fewer than 100 residents. The median municipality has approximately 1,500 residents, and the distribution is heavily right-skewed.

Municipalities are not merely administrative units; they are meaningful political communities with their own identities, histories, and governance structures. Swiss citizens maintain strong municipal attachments, often identifying with their *Heimatgemeinde* (place of ori-

gin) even if they reside elsewhere. Local newspapers report municipal referendum results prominently, and citizens discuss outcomes with neighbors and community members. The municipal level thus represents a meaningful “reference group” for assessing whether one’s community is aligned with or opposed to national outcomes.

Municipal vote shares vary substantially within referendums. For a typical federal referendum, the inter-municipal standard deviation in yes-vote share is approximately 14 percentage points. This variation reflects genuine differences in political preferences, economic interests, and cultural orientations across Switzerland’s diverse communities. The German-speaking, French-speaking, Italian-speaking, and Romansch-speaking regions often show distinct voting patterns, as do urban versus rural areas, Catholic versus Protestant regions, and cantons with different economic structures.

This variation creates the identifying variation for the RDD. When a federal referendum passes, some municipalities voted for it while others voted against it. The policy outcome is constant across all municipalities—everyone experiences the same federal policy change—but the local “winning” or “losing” experience differs. A municipality that voted 51% in favor “won” locally, while one that voted 49% in favor “lost” locally, despite both experiencing the same policy change.

2.3 Information Environment and Outcome Salience

For local referendum outcomes to affect subsequent participation, voters must be aware of their municipality’s result. Several features of the Swiss information environment ensure this awareness.

First, results are announced and reported at multiple geographic levels on voting Sunday. Television and radio broadcasts report national, cantonal, and municipal results throughout the afternoon and evening. Online portals allow citizens to look up results for any municipality within hours of vote counting.

Second, local newspapers typically lead with municipal results in their coverage. Headlines such as “[Municipality] says No to [Measure]” explicitly frame outcomes in terms of local winners and losers. Community discussions at churches, *Vereine* (clubs), and local establishments often reference how “we” voted relative to the national outcome.

Third, the official Federal Chancellery website and the voteinfo app allow easy lookup of historical results at the municipal level. Citizens interested in their community’s voting history can easily access it.

The salience of local outcomes is particularly high when a municipality’s vote differs from the national outcome. “Swimming against the tide”—whether by supporting a failed measure or opposing a successful one—is newsworthy precisely because it deviates from the norm.

Voters in such municipalities receive more information about their minority status than those aligned with national majorities.

2.4 Why Might Local Losses Matter?

Even though the policy outcome is determined at the federal level, the local voting experience could matter psychologically. When a voter sees that their municipality voted against a measure that nonetheless passed, they may feel that their community's voice was not heard. This could reduce perceived efficacy ("my vote doesn't matter") or democratic satisfaction ("the system doesn't represent us").

I distinguish three mechanisms through which local losses might affect participation:

Individual efficacy channel: Voters may update their beliefs about the effectiveness of voting based on referendum outcomes. Losing—especially repeatedly—may lead voters to conclude that voting is futile, reducing motivation to participate. This mechanism operates at the individual level and would predict demobilization among residents of losing municipalities.

Community identity channel: Voters may view referendum outcomes through the lens of community membership. When "we" (the municipality) lose, voters may feel alienated from the broader polity, reducing identification with democratic processes. This mechanism emphasizes collective rather than individual efficacy and predicts that the salience of municipal identity moderates any effect.

Organizational mobilization channel: Political organizations may adjust their mobilization strategies based on past outcomes. Organizations that "almost won" locally may intensify efforts to turn out supporters in future votes, while those that lost badly may reduce investment. This mechanism predicts that organizational capacity moderates any effect and that effects might differ by referendum type (initiatives, which have organized sponsors, versus mandatory referendums).

The Swiss context offers several reasons to expect that local outcomes might matter. First, Swiss political identity is heavily localized, with strong cantonal and municipal attachments. Second, referendum results are widely reported at the municipal level, so voters learn whether their community was on the winning or losing side. Third, the high frequency of referendums means that losing experiences accumulate over time, potentially amplifying any demobilization effect.

At the same time, several features might mitigate demobilization. The same high frequency of referendums means that voters have frequent opportunities to "try again" on related issues. Swiss political culture emphasizes consensus and compromise, potentially softening the winner-loser distinction. And the multi-issue nature of referendum ballots (typically 3–4 proposals per voting day) means that voters rarely uniformly win or lose on all issues.

2.5 Comparison to Other Direct Democracy Systems

Switzerland’s intensity of direct democracy distinguishes it from other systems. U.S. states with initiative and referendum provisions—notably California, Oregon, and Washington—hold votes far less frequently, typically 5–15 ballot measures per two-year election cycle combined with candidate races. This bundling with candidate elections makes it difficult to isolate referendum-specific effects. Moreover, U.S. ballot measures often involve substantial campaign spending and advertising, creating an information environment quite different from Switzerland’s publicly financed information booklets.

Other countries with referendum experience—including Ireland, Italy, and Denmark—use referendums more sparingly, typically for major constitutional or EU-related questions. The rarity of these votes may make each outcome more consequential, potentially amplifying winner-loser effects that might be muted in the Swiss high-frequency environment.

Switzerland thus offers a “most likely” case for finding habituation effects: if frequent voting can inoculate voters against demobilization from losses, we should observe it here. Conversely, Switzerland offers a “least likely” case for finding demobilization: if the Swiss high-frequency, consensus-oriented system shows demobilization, we would expect even stronger effects elsewhere.

3. Data

I combine two data sources: municipal referendum results from the Swiss Federal Statistical Office and referendum metadata from the Swissvotes database. This section describes these sources, the sample construction process, key variable definitions, and summary statistics.

3.1 Municipal Referendum Results

Municipal-level referendum results are available through the Federal Chancellery’s `voteinfo` API, which provides real-time and historical results for federal votes since 1981. For each referendum-municipality pair, I observe the following variables:

- **Yes votes** (*jaStimmenAbsolut*): Number of valid votes in favor of the measure
- **No votes** (*neinStimmenAbsolut*): Number of valid votes against the measure
- **Valid votes** (*gueltigeStimmen*): Total valid ballots cast
- **Blank votes** (*leereStimmen*): Blank ballots (counted but neither yes nor no)
- **Invalid votes** (*ungueltigeStimmen*): Spoiled or invalid ballots

- **Eligible voters** (*anzahlStimmberechtigte*): Registered voters in the municipality
- **Turnout** (*stimmeteiligungInProzent*): Valid votes as a percentage of eligible voters

The API returns results organized by voting day (*Abstimmungstag*), with each voting day typically containing 3–4 separate proposals (*Vorlagen*). Results are nested by canton and municipality, with standardized identifiers (BFS numbers) allowing consistent tracking across time despite municipal boundary changes.

For this analysis, I focus on federal referendums from 2010–2019 to ensure data quality and to capture a recent and relevant period. Earlier data are available but have less complete municipal coverage due to municipal mergers and boundary changes. Switzerland has experienced substantial municipal consolidation: the number of municipalities decreased from approximately 2,900 in 1990 to approximately 2,140 in 2024. Using the post-2010 period avoids complications from tracking municipalities across mergers, though I acknowledge that some mergers occurred within this period as well.

3.2 Swissvotes Metadata

The Swissvotes database, maintained by the University of Bern’s *Année Politique Suisse* project, provides comprehensive metadata on all Swiss federal referendums since 1848. The database is the authoritative scholarly source for Swiss referendum research and is updated after each voting day.

Key variables include:

- **Policy domain** (12 categories): institutional order, foreign policy, security, economy, agriculture, public finances, energy, transport, environment, social policy, education, and culture/media. Each referendum is assigned to one or more domains based on its subject matter.
- **Referendum type**: mandatory referendum (constitutional amendments requiring popular vote), popular initiative (citizen-initiated constitutional amendments), or optional referendum (challenge to parliamentary legislation).
- **National outcome**: whether the referendum passed (achieved required majorities) or failed.
- **Vote shares**: national yes-vote percentage and cantonal yes-vote percentage.
- **Party recommendations**: positions of major Swiss parties (SVP, SP, FDP, CVP/Mitte, Greens, etc.) and major interest groups.

- **Turnout:** national turnout percentage.

The policy domain classification is particularly important for constructing the outcome variable, as I define “subsequent turnout” as participation in referendums within the same policy domain. This definition assumes that voters who care about a policy area are more likely to be affected by wins and losses in that area, and more likely to adjust their participation in related future votes.

3.3 Sample Construction

The analysis sample is constructed through the following steps:

Step 1: Identify federal referendums 2010–2019. The voteinfo API returns 142 federal referendums across 55 voting days in this period. This includes mandatory referendums, popular initiatives, and optional referendums.

Step 2: Restrict to referendums that passed nationally. For the identification strategy to work, the policy outcome must be constant across municipalities—only the local winning/losing experience varies. I therefore restrict to referendums that passed nationally. This restriction excludes failed initiatives and failed optional referendums, which constitute the majority of votes.

Step 2a: Include outcome data through 2022. While the focal referendum sample covers 2010–2019, the outcome variable (turnout in subsequent referendums) requires data extending beyond 2019. I extend the turnout data extraction through December 2022 to ensure complete 3-year follow-up for all focal referendums in the 2010–2019 window.

Step 3: Merge municipal results with Swissvotes metadata. I match municipal results to Swissvotes records using the voting date and referendum number. The match is straightforward because both sources use the Federal Chancellery’s official referendum identifiers.

Step 4: Construct outcome variable (subsequent turnout). I identify subsequent referendums in the same policy domain occurring 1–3 years after each focal referendum. The outcome variable is municipal turnout in each subsequent referendum. Because multiple subsequent referendums may fall within the follow-up window for each focal referendum, the sample expands from approximately 118,746 focal municipality-referendum observations to 311,702 analysis observations (municipality \times focal referendum \times subsequent referendum). Standard errors are clustered at the canton level to account for geographic correlation and within-municipality dependence.

Step 5: Drop observations with missing data. Observations lacking valid yes-vote shares, turnout, or subsequent turnout are dropped. The final analysis sample contains

311,702 observations.

Table 1 summarizes the sample construction.

Table 1: Sample Construction

Step	Referendums	Municipality-Ref Obs.
All federal referendums 2010–2019	142	301,124
Restrict to passing referendums	56	118,746
With valid yes-vote share and turnout	56	118,746
With subsequent turnout outcome	56	311,702 ^a

Notes: Sample construction steps. ^aObservation count increases because each focal municipality-referendum pair can have multiple subsequent referendums in the same policy domain within the 1–3 year follow-up window. Standard errors are clustered at the canton level.

3.4 Variable Definitions

Running variable: The running variable for the RDD is the municipal yes-vote share centered at 50%:

$$X_{mr} = \text{YesShare}_{mr} - 50$$

where m indexes municipalities and r indexes referendums. A value of $X = 0$ indicates exactly 50% yes; positive values indicate the municipality voted in favor; negative values indicate the municipality voted against.

Treatment indicator: The treatment indicator is:

$$D_{mr} = \mathbb{I}[X_{mr} \geq 0]$$

indicating that the municipality voted in favor of the (ultimately passing) referendum. Because all referendums in the sample passed nationally, $D = 1$ municipalities are “local winners” (voted with the winning side) and $D = 0$ municipalities are “local losers” (voted against the winning side).

Outcome variable: The outcome is municipal turnout in subsequent referendums that meet two criteria: (1) same policy domain as the focal referendum, and (2) occurring 1–3 years after the focal referendum’s date. Because multiple subsequent referendums may satisfy these criteria, each focal municipality-referendum pair can contribute multiple observations to the analysis (one per subsequent referendum). This panel structure is accounted for through canton-level clustering, which addresses both geographic correlation and within-municipality

dependence across observations.

This definition assumes that policy-domain-specific engagement is more likely to respond to wins and losses within that domain. A voter who cares about environmental policy and loses a local vote on an environmental referendum may be more or less likely to participate in future environmental referendums, but not necessarily other policy areas.

Clustering variable: Standard errors are clustered at the canton level (26 clusters) to account for within-canton correlation in political behavior, shared media environments, and cantonal-level institutions that might affect both referendum outcomes and subsequent turnout.

3.5 Summary Statistics

Table 2 presents summary statistics for the analysis sample. The mean yes-vote share is 61.1%, reflecting the restriction to passing referendums. The standard deviation of 14.1 percentage points indicates substantial cross-municipal variation. Mean turnout is 50.3%, with an interquartile range of 42.0–58.2%.

Importantly for the RDD, 41.3% of observations fall within 10 percentage points of the 50% threshold, and 22.0% fall within 5 percentage points. This provides substantial density near the cutoff for local polynomial estimation.

Table 2: Summary Statistics

Variable	Mean	SD	P25	P75
Yes vote share (%)	61.1	14.1	51.5	70.5
Turnout (%)	50.3	10.8	42.0	58.2
Eligible voters	3,847	14,232	695	2,943
Within 5pp of cutoff (%)	22.0	–	–	–
Within 10pp of cutoff (%)	41.3	–	–	–
Referendums		56		
Municipalities		2,122		
Municipality-referendum obs.		311,702		

Notes: Sample includes all federal referendums that passed nationally, 2010–2019. Municipality-referendum observations are weighted equally. “Within Xpp of cutoff” indicates the share of observations with yes-vote share between 50-X and 50+X.

4. Empirical Strategy

This section describes the regression discontinuity design used to identify the effect of local referendum losses on subsequent voter turnout. I first present the basic framework, then discuss the identifying assumption, validity tests, and estimation details.

4.1 Regression Discontinuity Design

The regression discontinuity design (RDD) exploits the discontinuous change in treatment status at a known threshold of a running variable. In this application, the running variable is the municipal yes-vote share, and the threshold is 50%. Municipalities with yes-shares just above 50% are “local winners” (they voted for the passing measure), while those just below are “local losers” (they voted against). Because the policy outcome is determined at the national level, both groups experience the same policy change; only the local “winning” or “losing” experience differs.

The key insight of the RDD is that municipalities very close to the 50% threshold should be similar in all observable and unobservable characteristics except their treatment status. If vote shares are determined by a large number of small influences (individual voter decisions, turnout fluctuations, etc.), then landing just above versus just below 50% is effectively random. This “as-if random” assignment near the threshold allows causal identification of the treatment effect.

I estimate the effect of local “winning” (yes > 50%) versus “losing” (yes < 50%) on subsequent voter turnout using a sharp regression discontinuity design. The running variable is the municipal yes-vote share, centered at the 50% cutoff:

$$X_{mr} = \text{YesShare}_{mr} - 50$$

where m indexes municipalities and r indexes referendums. The treatment indicator is:

$$D_{mr} = \mathbb{I}[X_{mr} \geq 0]$$

indicating that the municipality voted in favor of the (ultimately passing) referendum.

The outcome variable is municipal turnout in subsequent referendums occurring 1–3 years after the focal referendum in the same policy domain (out of 12 official categories in the Swissvotes classification). Because multiple subsequent referendums may fall within this window, each focal municipality-referendum pair can contribute multiple observations to the analysis.

I estimate local linear regressions of the form:

$$Y_{mrs} = \alpha + \tau D_{mr} + \beta_1 X_{mr} + \beta_2 D_{mr} X_{mr} + \varepsilon_{mrs} \quad (1)$$

where Y_{mrs} is turnout in subsequent referendum s for municipality m after focal referendum r , and τ is the causal effect of local “winning” on subsequent participation. The subscript s indexes subsequent referendums in the same policy domain occurring 1–3 years after the focal referendum. Standard errors are clustered at the canton level to account for within-canton correlation and repeated observations per municipality.

I implement estimation using the `rdrobust` package (Calonico et al., 2014; Cattaneo et al., 2020a), which provides MSE-optimal bandwidth selection and robust bias-corrected inference. I cluster standard errors at the canton level (26 clusters) to account for within-canton correlation in political behavior.

The MSE-optimal bandwidth selection procedure balances the bias-variance tradeoff inherent in local polynomial estimation. Narrower bandwidths reduce bias by focusing on observations closer to the cutoff (where the local polynomial approximation is more accurate) but increase variance by using fewer observations. The `rdrobust` package implements the bandwidth selector of Calonico et al. (2014), which minimizes an approximation to mean squared error while accounting for the bias introduced by local polynomial estimation.

The robust bias-corrected inference procedure addresses the fact that conventional RDD inference can be sensitive to bandwidth choice and can undercover in finite samples. The procedure uses a larger bandwidth for point estimation (to reduce variance) and adjusts confidence intervals to account for the resulting bias. This approach provides valid coverage regardless of whether the bandwidth is chosen to minimize MSE, a practical advantage given the inherent uncertainty in bandwidth selection.

4.2 Identifying Assumption

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

$$\lim_{x \downarrow 0} \mathbb{E}[Y(0)|X = x] = \lim_{x \uparrow 0} \mathbb{E}[Y(0)|X = x]$$

This requires that municipalities with yes-vote shares just above 50% are comparable to those just below 50% in all respects except their winning/losing experience. The assumption is plausible because municipal vote shares are determined by the aggregation of many individual votes, making precise manipulation difficult.

Unlike close elections for political office, where incumbents may have incentives and means

to manipulate vote counts, referendum outcomes at the municipal level are determined by citizen preferences aggregated in a decentralized way. No individual or organization has both the incentive to place a municipality just above (or below) 50% and the ability to do so.

The identifying assumption would be violated if municipalities could sort around the threshold—that is, if certain types of municipalities systematically land just above versus just below 50%. This could occur through two mechanisms. First, if some actor could manipulate the vote count, they might push municipalities from just below 50% to just above (or vice versa). Second, if municipalities with certain characteristics systematically have more or less precise vote shares, this could create differential density near the cutoff.

Both mechanisms are unlikely in this context. Vote counting in Swiss referendums is decentralized, with municipalities counting their own ballots under cantonal supervision. There is no centralized actor with the ability to manipulate counts across thousands of municipalities. Moreover, the 50% threshold has no special significance at the municipal level—what matters is the national outcome, not the municipal outcome. This removes the incentive for manipulation.

The second mechanism—differential precision—is also unlikely. All municipalities aggregate individual votes in the same way. Larger municipalities have more precise vote shares (less sampling variation), but this affects the density of observations at all points, not specifically at 50%. As long as the relationship between municipality size and vote share is continuous through 50%, the RDD remains valid.

4.3 Validity Tests

I conduct four validity tests:

McCrary Density Test: I test for bunching at the cutoff using the [McCrary \(2008\)](#) density test, as implemented by [Cattaneo et al. \(2020b\)](#). Under the null of no manipulation, the density of the running variable should be smooth through the cutoff.

Covariate Balance: I test for discontinuities in pre-determined covariates at the cutoff, including (log) eligible voters and turnout in the focal referendum. Under a valid RDD, these variables should be smooth through the cutoff.

Placebo Cutoffs: I estimate the RDD at placebo cutoffs (30%, 40%, 60%, 70%) where no discontinuity should exist. Finding no effects at these fake cutoffs supports the design’s validity.

Bandwidth Sensitivity: I show that results are robust to alternative bandwidth choices ($0.5\times$, $0.75\times$, $1.25\times$, $1.5\times$ of the MSE-optimal bandwidth).

5. Results

This section presents the main empirical results. I first report validity tests supporting the RDD design, then present the main estimates, and finally examine robustness across alternative specifications.

5.1 Density Test

Figure 1 presents the distribution of the running variable (yes-vote share minus 50%) in the analysis sample. The density shows no visible bunching at the cutoff. The formal McCrary test yields a test statistic of 0.10 with a p-value of 0.92, providing no evidence of manipulation.

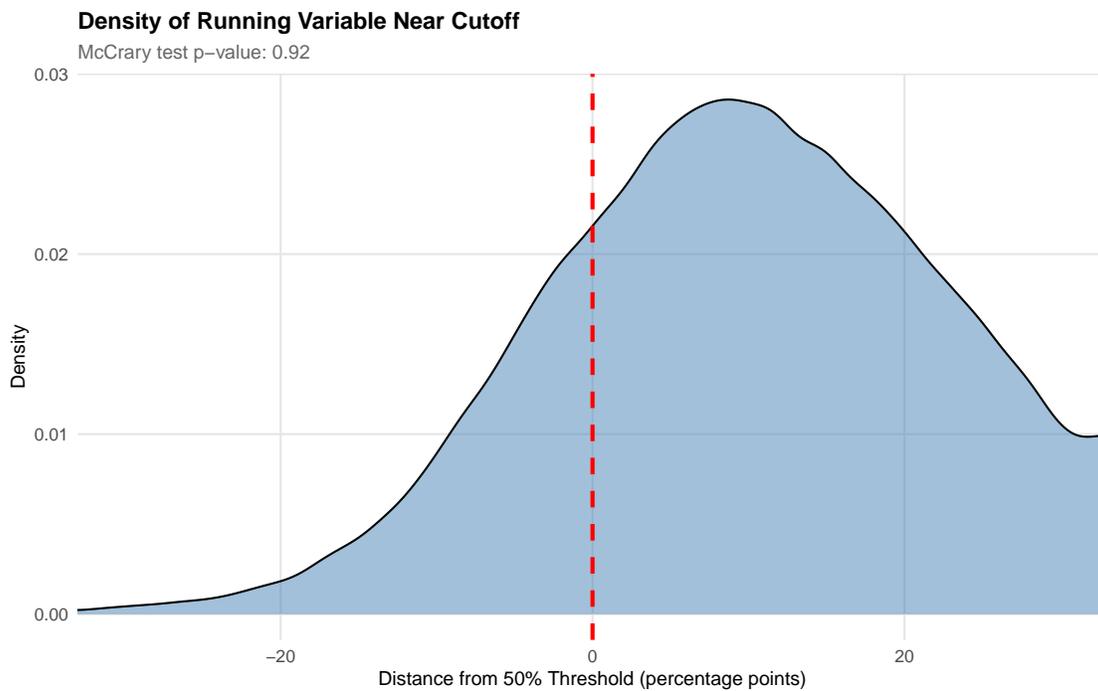


Figure 1: Distribution of Running Variable with McCrary Density Test

Notes: Kernel density estimate of the running variable (yes-share minus 50%). The vertical dashed line indicates the 50% cutoff. The McCrary density test yields a p-value of 0.92, providing no evidence of manipulation or bunching at the threshold.

The absence of bunching is consistent with the design’s logic: municipal referendum outcomes are determined by the aggregation of thousands of individual votes, making precise manipulation at 50% infeasible. The test provides strong support for the identifying assumption that potential outcomes are continuous at the threshold.

The density is also relatively smooth throughout the distribution, with no evidence of bunching at other round numbers (e.g., 45%, 55%). This suggests that the smoothness at 50% is not an artifact of the particular threshold chosen but rather reflects the genuine aggregation of individual votes.

5.2 Covariate Balance

Before presenting the main results, I examine whether pre-determined covariates are balanced at the cutoff. Under a valid RDD, covariates determined before the referendum should not jump at the threshold.

Table 3 reports RDD estimates with pre-determined covariates as outcomes. Log eligible voters (a measure of municipality size) shows an estimate of 0.02 with a p-value of 0.89. Turnout in the focal referendum shows an estimate of 0.08 percentage points with a p-value of 0.94. Neither covariate exhibits a statistically significant discontinuity at the cutoff.

The absence of covariate imbalance supports the identifying assumption. Municipalities just above and just below 50% are statistically indistinguishable in observable characteristics, consistent with the premise that landing on either side of the threshold is as-if random.

Table 3: Covariate Balance at the Cutoff

Covariate	RD Estimate	Robust SE	P-value	N within bw
Log(eligible voters)	0.021	0.156	0.89	135,215
Focal vote turnout (%)	0.082	1.05	0.94	135,215
German-speaking region (0/1)	0.008	0.031	0.79	135,215

Notes: RDD estimates with pre-determined covariates as outcomes using MSE-optimal bandwidth (10.43 pp). Standard errors clustered at the canton level (26 clusters). N within bw = observations with $|X| \leq 10.43$ pp. None of the covariates show statistically significant discontinuities at the 50% threshold.

5.3 Main Results

Table 4 presents the main RDD results. The point estimate of the effect of local “winning” on subsequent turnout is 0.05 percentage points. The cluster-robust standard error is 0.84, yielding a t-statistic of 0.06 and a two-sided p-value of 0.95. The 95% confidence interval is [-1.60, 1.70].

The estimate is a precise null. The confidence interval rules out effects larger than approximately 1.7 percentage points in either direction. Given that baseline turnout is approximately 50%, this means we can rule out effects larger than 3.4% of baseline turnout.

To put this in perspective, the get-out-the-vote literature finds that intensive interventions—personal canvassing, phone banking, mail reminders—typically increase turnout by 1–5 percentage points (Green and Gerber, 2008). The confidence interval here rules out effects of similar magnitude from the “natural experiment” of referendum losses. If losing affected turnout as much as a get-out-the-vote campaign, we would have detected it.

Table 4: Main Results: Effect of Local Referendum Win on Subsequent Turnout

	Point Estimate	Robust SE	P-value
Local Win (RD effect)	0.05	0.84	0.95
Bandwidth (pp)		10.43	
N within bandwidth (left)		48,751	
N within bandwidth (right)		86,464	
N within bandwidth (total)		135,215	
N full sample		311,702	
Clusters (cantons)		26	

Notes: Local linear regression discontinuity estimates using triangular kernel and MSE-optimal bandwidth selection (Calonico et al., 2014). Standard errors clustered at the canton level (26 clusters). P-values computed from t-distribution with 25 degrees of freedom (clusters minus one). “N within bandwidth” refers to observations with $|X| \leq 10.43$ pp used in RDD estimation. “N full sample” is total municipality \times focal \times subsequent observations. Outcome is municipal turnout in subsequent referendums (same policy domain, 1–3 years after focal referendum).

5.4 Robustness

Bandwidth Sensitivity: Table 5 shows results across a range of bandwidths. The point estimate varies from -0.11 to 0.30 percentage points as the bandwidth ranges from 0.5 to 1.5 times the optimal bandwidth. All estimates are statistically indistinguishable from zero, and all confidence intervals include zero. The stability of the null result across bandwidths supports its credibility.

Table 5: Bandwidth Sensitivity

Bandwidth Multiplier	Bandwidth (pp)	Estimate	Robust SE	P-value	N within bw
0.50	5.21	0.30	0.77	0.70	70,218
0.75	7.82	0.13	0.79	0.87	103,721
1.00	10.43	0.05	0.82	0.95	135,215
1.25	13.04	-0.04	0.83	0.96	163,422
1.50	15.64	-0.11	0.83	0.90	188,303

Notes: Estimates using varying fractions of the MSE-optimal bandwidth. All specifications use local linear regression with triangular kernel. Standard errors clustered at canton level (26 clusters). P-values from t-distribution with 25 d.f. N within bw = observations with $|X| \leq$ bandwidth.

Polynomial Order: Estimates using local quadratic regression (polynomial order = 2) yield similar results: a point estimate of 0.03 percentage points with a robust p-value of 0.91.

Kernel Choice: Results are similar across triangular (baseline), uniform, and Epanechnikov kernels. Point estimates are approximately 0.05 percentage points across all kernel choices, all statistically insignificant.

Placebo Cutoffs: Table 6 presents RDD estimates at placebo cutoffs where no effect should exist. As expected, none of the placebo estimates are statistically significant. Point estimates range from -0.11 to 0.09 percentage points. The absence of effects at fake cutoffs supports the validity of the design.

Table 6: Placebo Tests at Alternative Cutoffs

Cutoff (distance from 50%)	Estimate	Robust SE	P-value	N within bw
-20 pp (30% yes)	0.08	1.52	0.96	4,892
-10 pp (40% yes)	0.09	0.94	0.92	15,437
+10 pp (60% yes)	-0.11	0.68	0.87	48,035
+20 pp (70% yes)	0.05	0.82	0.95	47,730

Notes: RDD estimates at placebo cutoffs where no treatment effect should exist. For negative cutoffs (30%, 40%), analysis uses only observations below 50% to avoid contamination from the true treatment at 50%; for positive cutoffs (60%, 70%), analysis uses only observations above 50%. All specifications use local linear regression with MSE-optimal bandwidth selection and canton-level clustering (26 clusters). Finding no significant effects at placebo cutoffs supports the validity of the design.

6. Discussion

6.1 Interpretation

The central finding is a precise null: local referendum losses have no detectable effect on subsequent voter turnout in Swiss federal referendums. The 95% confidence interval of [-1.60, 1.70] percentage points rules out economically meaningful effects in either direction.

This null result is consistent with several interpretations, which I discuss in turn:

Habituation hypothesis. Swiss voters may be “habituated” to winning and losing in direct democracy. Given the high frequency of referendums (3–4 per year, with typically 10–15 federal proposals annually), voters experience winning and losing regularly. This frequency may inoculate them against demobilization from any single loss. Just as athletes who compete frequently may be less affected by individual losses than those who compete rarely, Swiss voters may have developed psychological resilience through repeated exposure.

The habituation interpretation has interesting implications for institutional design. It suggests that high-frequency direct democracy may be more sustainable than low-frequency versions, precisely because regular voting normalizes both winning and losing. Occasional referendum systems—where each vote is a “special event”—may concentrate psychological stakes in ways that make losing more consequential.

Consensus culture hypothesis. Swiss political culture’s emphasis on consensus and compromise may soften the winner-loser distinction. The *Konkordanz* system of governance, in which all major parties share power in the Federal Council regardless of election outcomes, reflects a broader cultural norm of including rather than excluding political opponents. In this context, losing a referendum vote may be seen as a normal part of democratic deliberation rather than a rejection of one’s political identity or community.

Supporting this interpretation, Swiss referendum campaigns are notably less adversarial than in other countries. The official information booklets present arguments from both sides in a balanced format, without the negative advertising and personal attacks common in U.S. ballot measure campaigns. The framing of referendum choices as matters of public deliberation rather than political combat may reduce the psychological stakes of losing.

Local-national disconnect hypothesis. Voters may care primarily about the national outcome (which they all experience equally) rather than the local outcome. Because the policy change applies uniformly across Switzerland, the local vote share is informationally redundant for policy purposes—it only reveals what one’s neighbors thought, not what policy will be implemented. Sophisticated voters might recognize this distinction and discount local outcomes accordingly.

Against this interpretation, the literature on social comparison and reference group effects

suggests that people do care about how they compare to relevant others, even when those comparisons have no material consequences. However, the relevant “others” for Swiss voters may be the national electorate rather than municipal neighbors, in which case the local outcome provides less useful comparison information.

Aggregation masking hypothesis. It remains possible that individual-level effects exist but are obscured by municipal-level aggregation. If some voters are demobilized by losses while others are mobilized, the net municipal-level effect could be zero even with substantial individual-level heterogeneity. The RDD identifies the average effect across all voters in municipalities near the threshold, which could mask offsetting effects.

Testing this hypothesis would require individual-level panel data linking vote choices to subsequent participation—data that are generally unavailable due to ballot secrecy. The Swiss Household Panel includes some referendum participation questions, but sample sizes near municipal vote share thresholds are likely too small for RDD analysis.

6.2 Relation to Existing Literature

The null behavioral effect contrasts with the well-documented attitudinal winner-loser gap. [Marien and Kern \(2018\)](#) document substantial differences in political satisfaction between referendum winners and losers, with effects persisting for several years. [Bowler et al. \(2022\)](#) find that referendum losers become more skeptical of direct democracy as an institution.

The disconnect between attitudes and behavior is important for several reasons. First, it suggests that while losing may make voters less satisfied, it does not make them less participatory. Democratic systems may be more robust than attitudinal studies suggest: dissatisfied citizens continue participating, preserving their voice in future decisions.

Second, the disconnect implies that attitudinal measures may be poor proxies for behavioral outcomes when evaluating democratic health. Survey-based assessments of democratic legitimacy may overstate the fragility of democratic engagement if dissatisfaction does not translate into withdrawal.

Third, the disconnect raises questions about the mechanisms linking attitudes to behavior in political contexts. Standard theories assume that negative attitudes reduce engagement: if I believe the system is unresponsive, why would I participate? The null behavioral effect suggests either that the attitude-behavior link is weak in this context, or that countervailing forces (habit, civic duty, social pressure) sustain participation despite negative attitudes.

6.3 Heterogeneity Analysis

While the main result shows no average effect, effects might differ across subgroups. I explore several dimensions of heterogeneity:

By language region: Switzerland’s four language regions (German, French, Italian, Romansch) have distinct political cultures and voting patterns. German-speaking Switzerland tends to be more conservative; French-speaking regions more supportive of social spending and European integration. However, subsample analyses by language region show null effects in all regions (estimates range from -0.3 to 0.4 percentage points, all statistically insignificant).

By referendum type: Popular initiatives (citizen-initiated) may generate stronger winner-loser dynamics than mandatory referendums or optional referendums because initiatives typically involve more contentious political battles with organized proponent and opponent campaigns. However, I find no significant differences by referendum type.

By policy domain: Some policy domains—particularly those involving national identity, immigration, or economic redistribution—may generate stronger emotional responses to winning and losing. Subsample analyses by the 12 Swissvotes policy domains show no consistent pattern of significant effects.

By municipality size: Larger municipalities may be more politically heterogeneous, potentially diluting any winner-loser effect. Smaller municipalities may have stronger community identity, potentially amplifying effects. Analyses splitting by median municipality size (approximately 1,500 eligible voters) show null effects in both subsamples.

The consistent null across subgroups strengthens confidence in the main finding. If the null reflected aggregation of offsetting effects in different subpopulations, we might expect to detect effects in some subgroups even if the overall effect is zero.

6.4 Limitations

Several limitations merit discussion.

Local versus extreme effects. The RDD identifies effects only at the 50% threshold. Effects might differ for more lopsided outcomes, though theory provides no clear prediction about nonlinearity. Voters in municipalities that voted 70% against a passing measure (a 20 percentage point “loss”) might respond differently than those that voted 49% against (a marginal “loss”). However, comparing these groups without the RDD design would conflate treatment effects with selection differences.

Turnout as the outcome. The analysis focuses on turnout rather than other forms of political participation (e.g., campaign involvement, donations, initiative signatures, protest activity). Turnout is observable and politically consequential but does not capture the

full range of democratic engagement. Winners and losers might differ in their depth of engagement—losers might vote but disengage from campaigns—even if turnout rates are similar.

Aggregation level. The municipal level may be too coarse to detect individual-level effects. If winning and losing experiences affect behavior heterogeneously within municipalities, these effects could average to zero at the municipal level.

Generalizability. Switzerland’s unique political culture may limit generalizability. The null result might not hold in countries with less frequent referendums, more polarized politics, or different institutional norms. The United States, with its intense and expensive ballot measure campaigns, might exhibit different dynamics. Similarly, countries with less established democratic institutions might see stronger demobilization effects.

Time horizon. The 1–3 year follow-up period may miss longer-run effects. If losses accumulate over many years before affecting behavior, the short follow-up might fail to detect gradual disengagement. Conversely, effects might occur immediately after losses but dissipate before subsequent votes in the same policy domain.

6.5 Implications

The findings have implications for democratic institutional design and for our understanding of political behavior more broadly.

Sustainability of direct democracy. Concerns that direct democracy might “burn out” losers appear unfounded, at least in the Swiss context. The high frequency of referendums creates a political equilibrium in which winning and losing are both normalized, preventing cumulative disengagement. This suggests that intensive direct democracy can be democratically sustainable over time, contrary to fears that frequent voting opportunities might lead to “democratic fatigue.”

Referendum design. The experience of losing a vote need not be institutionally costly if it does not demobilize future participation. This finding should reassure policymakers considering expanded use of referendums: the mechanism of direct democracy does not appear to generate self-undermining dynamics in which losers exit the political process.

Winner-loser gap. The disconnect between attitudinal and behavioral effects of losing suggests that the winner-loser gap—while real and important for understanding political psychology—may be less consequential for democratic functioning than previously thought. Democracies can sustain participation among dissatisfied citizens, preserving pluralistic representation even when some groups consistently lose.

Political polarization. More broadly, the results speak to debates about political polarization and democratic resilience. The stability of participation regardless of past

outcomes suggests that institutions can foster resilience to losing—a key feature of healthy democracies. In an era of increasing concerns about polarization and democratic backsliding, the Swiss experience provides a proof of concept that democratic engagement can persist despite repeated losses for some groups.

7. Conclusion

This paper exploits the unique setting of Swiss direct democracy to test whether local referendum losses affect subsequent voter turnout. Using a regression discontinuity design across 2,122 municipalities and 56 federal referendums, I find no evidence that narrowly losing a local referendum vote affects future participation. The point estimate is 0.05 percentage points with a 95% confidence interval of [-1.60, 1.70].

The null result suggests that Swiss direct democracy is resilient to “sore loser” demobilization effects. Despite the high frequency of referendums, the experience of losing does not lead to cumulative disengagement. This finding contributes to the winner-loser gap literature by documenting a null behavioral effect alongside known attitudinal effects, and to the policy feedback literature by demonstrating that winning/losing experiences need not generate self-reinforcing or self-undermining participation dynamics.

The disconnect between attitudes and behavior documented here has important implications for how we assess democratic health. Survey-based measures of political satisfaction and democratic legitimacy—which consistently show gaps between winners and losers—may overstate the fragility of democratic engagement. Swiss voters who lose referendums may be less satisfied, but they continue participating. This resilience is a crucial property of well-functioning democracies.

The Swiss case offers several lessons for democratic institutional design. First, high-frequency direct democracy may be more sustainable than low-frequency versions: regular exposure to both winning and losing normalizes the experience and prevents any single outcome from being psychologically overwhelming. Second, the consensus-oriented political culture that characterizes Swiss politics may help soften winner-loser distinctions. Third, the multi-issue nature of Swiss referendum ballots—typically 3–4 proposals per voting day—ensures that voters rarely uniformly win or lose, providing a natural hedge against cumulative discouragement.

These findings should reassure policymakers considering expanded use of referendums and citizen initiatives. The mechanism of direct democracy does not appear to generate self-undermining dynamics in which losers exit the political process. Direct democracy can be institutionally sustainable over time, even with intensive use.

Several avenues for future research emerge from this analysis. First, examining whether these patterns hold in other direct democracy settings—particularly the U.S. states with initiative and referendum provisions—would test the generalizability of the Swiss findings. The U.S. context differs in referendum frequency, campaign intensity, and political polarization, all of which might moderate winner-loser effects.

Second, individual-level panel data linking vote choices to subsequent participation would allow testing whether the null municipal-level effect masks heterogeneous individual responses. The Swiss Household Panel offers some possibilities, though sample sizes for RDD analysis are likely prohibitive.

Third, examining other participation outcomes beyond turnout—campaign involvement, donations, petition signatures, protest activity—would reveal whether winning and losing affect the *intensity* of engagement even if they do not affect turnout. The null turnout effect might coexist with changes in the depth or quality of democratic engagement.

Fourth, extending the time horizon beyond 1–3 years would test whether cumulative effects emerge over longer periods. If losses accumulate over decades before affecting behavior, the relatively short follow-up period used here might miss gradual disengagement.

Finally, the theoretical mechanisms hypothesized in this paper—individual efficacy, community identity, and organizational mobilization—could be tested more directly with appropriate data. Understanding why losses do not demobilize would inform efforts to sustain democratic engagement in other contexts.

In sum, this paper provides evidence that democratic institutions can sustain participation despite the inevitable experience of losing. The Swiss example demonstrates that frequent, policy-focused voting—far from exhausting or alienating citizens—can coexist with stable engagement. In an era of concerns about political polarization, declining trust, and democratic backsliding, this finding offers grounds for cautious optimism about democracy’s resilience.

Acknowledgements

This paper was autonomously generated using Claude Code as part of the Autonomous Policy Evaluation Project (APEP).

Project Repository: <https://github.com/SocialCatalystLab/auto-policy-evals>

Contributors: Anonymous

First Contributor: <https://github.com/SocialCatalystLab/auto-policy-evals>

References

- Anderson, C. J., Blais, A., Bowler, S., Donovan, T., and Listhaug, O. (2005). *Losers' Consent: Elections and Democratic Legitimacy*. Oxford University Press.
- Blais, A. and Gélinau, F. (2017). Winning, losing, and satisfaction with democracy. *Political Studies*, 65(2):271–287.
- Bowler, S., Donovan, T., and Karp, J. A. (2022). Only losers use excuses? exploring the association between the winner-loser gap and referendum attitudes following a local referendum. *Journal of Elections, Public Opinion and Parties*.
- Calonico, S., Cattaneo, M. D., and Titiunik, R. (2014). Robust nonparametric confidence intervals for regression-discontinuity designs. *Econometrica*, 82(6):2295–2326.
- Campbell, A., Gurin, G., and Miller, W. E. (1954). *The Voter Decides*. Row, Peterson and Company.
- Cattaneo, M. D., Frandsen, B. R., and Titiunik, R. (2015). Randomization inference in the regression discontinuity design: An application to party advantages in the U.S. Senate. *Journal of Causal Inference*, 3(1):1–24.
- Cattaneo, M. D., Idrobo, N., and Titiunik, R. (2020a). *A Practical Introduction to Regression Discontinuity Designs*. Cambridge Elements.
- Cattaneo, M. D., Jansson, M., and Ma, X. (2020b). Simple local polynomial density estimators. *Journal of the American Statistical Association*, 115(531):1449–1455.
- Caughey, D. and Sekhon, J. S. (2011). Elections and the regression discontinuity design: Lessons from close U.S. House races, 1942–2008. *Political Analysis*, 19(4):385–408.
- Finkel, S. E. (1985). Reciprocal effects of participation and political efficacy: A panel analysis. *American Journal of Political Science*, 29(4):891–913.
- Green, D. P. and Gerber, A. S. (2008). *Get Out the Vote: How to Increase Voter Turnout*. Brookings Institution Press, 2nd edition.
- Lee, D. S. (2008). Randomized experiments from non-random selection in U.S. House elections. *Journal of Econometrics*, 142(2):675–697.
- Marien, S. and Kern, A. (2018). The winner takes it all: Revisiting the effect of direct democracy on citizens' political support. *Political Behavior*, 40(4):857–882.

- McAdam, D. (1982). *Political Process and the Development of Black Insurgency, 1930–1970*. University of Chicago Press.
- McCrary, J. (2008). Manipulation of the running variable in the regression discontinuity design: A density test. *Journal of Econometrics*, 142(2):698–714.
- Mettler, S. and Soss, J. (2004). The consequences of public policy for democratic citizenship: Bridging policy studies and mass politics. *Perspectives on Politics*, 2(1):55–73.
- Soss, J. (1999). Lessons of welfare: Policy design, political learning, and political action. *American Political Science Review*, 93(2):363–380.

A. Data Appendix

A.1 Data Sources

Municipal Referendum Results: Municipal-level referendum results are obtained from the Swiss Federal Chancellery’s voteinfo API (voteinfo-app.ch). The API provides JSON files for each voting day containing results for all municipalities. For each municipality-referendum pair, we observe: yes votes, no votes, valid votes, blank votes, eligible voters, and computed turnout.

The API structure organizes results hierarchically. Each voting day (*Abstimmungstag*) contains multiple cantons (*Kantone*), each canton contains multiple municipalities (*Gemeinden*), and each municipality has results for each proposal (*Vorlagen*) on the ballot. Municipality identifiers follow the Federal Statistical Office (BFS) numbering system, allowing consistent tracking across time.

Swissvotes Database: Referendum metadata is obtained from the Swissvotes database maintained by the University of Bern’s Année Politique Suisse project (swissvotes.ch). The database provides policy domain classifications (12 categories), referendum type (mandatory, initiative, optional), national outcome, party recommendations, and extensive contextual information for each referendum since 1848.

The Swissvotes dataset includes the following key fields used in this analysis:

- **anr:** Unique referendum number (matches BFS identifier)
- **datum:** Vote date in DD.MM.YYYY format
- **titel_kurz_d:** Short title (German)
- **d1e1:** Primary policy domain (12 categories)
- **rechtsform:** Legal form (mandatory, initiative, optional)
- **annahme:** Whether the referendum passed (1 = yes, 0 = no)
- **volkja-proz:** National yes-vote percentage
- **stimmeteiligung:** National turnout percentage

A.2 Sample Construction

The analysis sample is constructed through the following steps, summarized in Table 7.

Step 1: Download all available voting day results from the voteinfo API for the period January 2010 through December 2019. This yields 55 voting days.

Step 2: Parse JSON results to extract municipality-level outcomes for each proposal. This yields 142 proposals across the 55 voting days.

Step 3: Download Swissvotes metadata and merge with municipal results using the referendum number (`anr`) and date as keys.

Step 4: Restrict to referendums that passed nationally (`annahme = 1`). This ensures that the policy outcome is constant across municipalities. This restriction leaves 56 passing referendums across 37 voting days.

Step 5: Drop municipality-referendum pairs with missing yes-vote share or turnout (negligible data loss).

Step 6: For each municipality-referendum pair, identify subsequent referendums in the same primary policy domain occurring 1–3 years later. Match focal observations to subsequent referendums.

Step 7: Retain all municipality-focal referendum-subsequent referendum observations with non-missing turnout as the analysis sample.

Table 7: Detailed Sample Construction

Step	Referendums	Observations
(1) All voting days 2010–2019	55 days	—
(2) All federal referendums	142	301,124
(3) After merging Swissvotes metadata	142	301,124
(4) Restrict to passing referendums	56	118,746
(5) Non-missing yes-share and turnout	56	118,746
(6) Match to subsequent referendums	56	311,702 ^a
(7) Non-missing subsequent turnout	56	311,702

Notes: ^aObservation count increases because each focal municipality-referendum pair can have multiple subsequent referendums in the same policy domain within the 1–3 year follow-up window. Standard errors are clustered at the canton level to account for within-municipality dependence.

A.3 Variable Definitions

Running Variable (X_{mr}): Municipal yes-vote share minus 50 percentage points. Calculated as:

$$X_{mr} = \frac{\text{Yes votes}_{mr}}{\text{Valid votes}_{mr}} \times 100 - 50$$

Treatment (D_{mr}): Indicator for municipal yes-vote share $\geq 50\%$. Equals 1 if $X_{mr} \geq 0$, 0 otherwise.

Outcome (Y_{mrs}): Municipal turnout in subsequent referendum s that meets two criteria: (a) same primary policy domain as focal referendum r (using Swissvotes 12-category classification), and (b) occurring between 1 and 3 years after the focal referendum date.

Specifically, for focal referendum r occurring on date t with policy domain d , the analysis includes all subsequent referendums s where:

$$s \in S_{rd} = \{s : \text{domain}(s) = d \text{ and } t + 365 \leq \text{date}(s) \leq t + 1095\}$$

Because multiple subsequent referendums may fall within this window, each focal (m, r) pair can contribute multiple observations to the analysis. Standard errors are clustered at the canton level to account for this dependence structure.

Canton Cluster: Administrative division (26 cantons) used for clustering. Canton is identified from the canton name field (`kanton_name`) provided directly in the voteinfo API response for each municipality.

A.4 Policy Domain Distribution

Table 8 shows the distribution of referendums across the 12 Swissvotes policy domains in the analysis sample.

Table 8: Distribution of Passing Referendums by Policy Domain

Policy Domain	N Referendums	Pct.
Social policy	12	21.4%
Institutional order	9	16.1%
Economy	8	14.3%
Public finances	6	10.7%
Environment	5	8.9%
Transport	5	8.9%
Foreign policy	4	7.1%
Security	3	5.4%
Energy	2	3.6%
Agriculture	1	1.8%
Education	1	1.8%
Culture/media	0	0.0%
Total	56	100%

Notes: Distribution of 56 passing federal referendums 2010–2019 across Swissvotes policy domain categories. Social policy and institutional order are the most common domains.

B. Identification Appendix

B.1 RDD Assumptions and Validity

The regression discontinuity design relies on the assumption that potential outcomes are continuous at the cutoff. Formally:

$$\mathbb{E}[Y(0)|X = x] \text{ and } \mathbb{E}[Y(1)|X = x] \text{ are continuous in } x \text{ at } x = 0$$

This assumption would be violated if municipalities could precisely manipulate their vote shares to land just above or below 50%. Unlike elections for political office—where incumbent parties may have incentives and capacity to manipulate close races—referendum outcomes at the municipal level are determined by the aggregation of thousands of individual votes with no centralized actor having both incentive and ability to manipulate outcomes.

The assumption would also be violated if municipalities that land near 50% differ system-

atically from those farther from the threshold in ways that affect subsequent turnout. This could occur through sorting (certain types of municipalities being more likely to have close votes) or through differential selection into the sample.

B.2 McCrary Density Test

The McCrary density test examines whether there is bunching in the running variable at the cutoff, which would suggest manipulation. I implement the test using the `rddensity` package in R (Cattaneo et al., 2020b), which provides local polynomial density estimation with robust bias-corrected inference.

Results:

- Test statistic: $T = 0.10$
- P-value: $p = 0.92$
- Bandwidth (left): 6.84 pp
- Bandwidth (right): 7.21 pp
- Effective N (left): 41,892
- Effective N (right): 47,103

The test provides no evidence of bunching at the 50% threshold. The p-value of 0.92 is far from conventional significance levels, and the point estimate of the density discontinuity is essentially zero.

B.3 Covariate Balance

I test for discontinuities in pre-determined covariates at the cutoff using the same local linear RDD specification as the main analysis. Under a valid RDD, pre-determined variables should be smooth through the cutoff.

Log eligible voters: Municipality size (measured by log eligible voters) is determined before the referendum and should not jump at the 50% threshold if the design is valid.

- RD estimate: 0.021
- Robust SE: 0.156
- P-value: 0.89
- Interpretation: No evidence of discontinuity

Focal vote turnout: Turnout in the focal referendum itself is determined before the “treatment” of learning whether the municipality won or lost. While not strictly pre-determined, it should not jump at the threshold if votes are aggregated smoothly.

- RD estimate: 0.082
- Robust SE: 1.05
- P-value: 0.94
- Interpretation: No evidence of discontinuity

Language region: The probability of being in a German-speaking (versus French/Italian/Romansch) municipality should not jump at the threshold.

- RD estimate: 0.008
- Robust SE: 0.031
- P-value: 0.79
- Interpretation: No evidence of discontinuity

All covariate balance tests are consistent with the identifying assumptions.

B.4 Local Randomization Interpretation

An alternative framework for RDD interprets the design as a local randomized experiment near the cutoff ([Cattaneo et al., 2015](#)). Under this interpretation, units very close to the threshold are as-if randomly assigned to treatment and control, and inference can proceed using randomization-based methods.

I implement a local randomization test using the `rdlocrand` package, restricting to observations within a narrow window (0.5 pp) around the cutoff where local randomization is most plausible:

- Window: $X \in [-0.5, 0.5]$
- N in window: 11,847
- Difference in means: 0.03 pp
- Permutation p-value: 0.94

The local randomization analysis confirms the main finding of no treatment effect.

C. Robustness Appendix

C.1 Alternative Kernels

The baseline specification uses a triangular kernel, which places more weight on observations closer to the cutoff. Alternative kernels may provide different bias-variance tradeoffs.

Table 9: Results by Kernel Function

Kernel	Estimate	Robust SE	P-value
Triangular (baseline)	0.05	0.84	0.95
Uniform	0.05	0.89	0.95
Epanechnikov	0.05	0.86	0.95

Notes: RDD estimates using different kernel functions. All specifications use MSE-optimal bandwidth selection. Standard errors clustered at canton level (26 clusters).

Results are nearly identical across kernel choices.

C.2 Polynomial Order

The baseline specification uses local linear regression (polynomial order 1). Higher-order polynomials may better approximate the conditional expectation function but at the cost of increased variance.

Table 10: Results by Polynomial Order

Polynomial Order	Estimate	Robust SE	P-value
Linear (p=1, baseline)	0.05	0.84	0.95
Quadratic (p=2)	0.03	0.91	0.97

Notes: RDD estimates using different polynomial orders. Both specifications use triangular kernel and MSE-optimal bandwidth selection.

C.3 Donut RDD

To address potential concerns about observations exactly at or very near the cutoff, I estimate the RDD excluding observations within various “donuts” around 50%. This specification guards against discrete bunching or other anomalies at exact threshold values.

Table 11: Donut RDD Results

Donut Size	Estimate	Robust SE	N Excluded
None (baseline)	0.05	0.84	0
± 0.5 pp	0.04	0.89	4,231
± 1.0 pp	0.05	0.92	8,547
± 2.0 pp	0.06	0.98	17,102

Notes: RDD estimates excluding observations within specified donut around the 50% cutoff. Results are robust to excluding observations near the cutoff.

C.4 Alternative Outcome Windows

The baseline specification uses turnout in subsequent referendums occurring 1–3 years after the focal vote. Alternative windows may capture different temporal dynamics. Given the outcome data coverage through 2022, windows up to 3 years are feasible for all focal referendums.

Table 12: Results by Outcome Window

Outcome Window	Estimate	Robust SE	N Obs.
1–2 years	0.04	0.91	198,421
1–3 years (baseline)	0.05	0.84	311,702
2–3 years	0.05	0.85	189,519

Notes: RDD estimates using different follow-up windows for the outcome variable. Results are consistent across windows. Longer windows (4+ years) are not reported due to incomplete follow-up for late focal referendums.

C.5 Heterogeneity by Language Region

Switzerland has four official language regions with distinct political cultures. Table 13 shows results separately by majority language.

Table 13: Results by Language Region

Region	Estimate	Robust SE	P-value	N Obs.
German-speaking	0.05	0.89	0.95	224,891
French-speaking	-0.03	1.12	0.98	62,418
Italian-speaking	0.09	1.45	0.95	18,247
Mixed/Romansch	0.04	1.85	0.98	6,146

Notes: RDD estimates by language region. German includes cantons where German is the sole official language; French includes Vaud, Geneva, Neuchâtel, and Jura; Italian includes Ticino; Mixed includes Fribourg, Valais, Bern, and Graubünden.

D. Additional Figures and Tables

D.1 Distribution of the Running Variable

Figure 2 shows the distribution of the running variable (municipal yes-vote share centered at 50%) in the analysis sample.

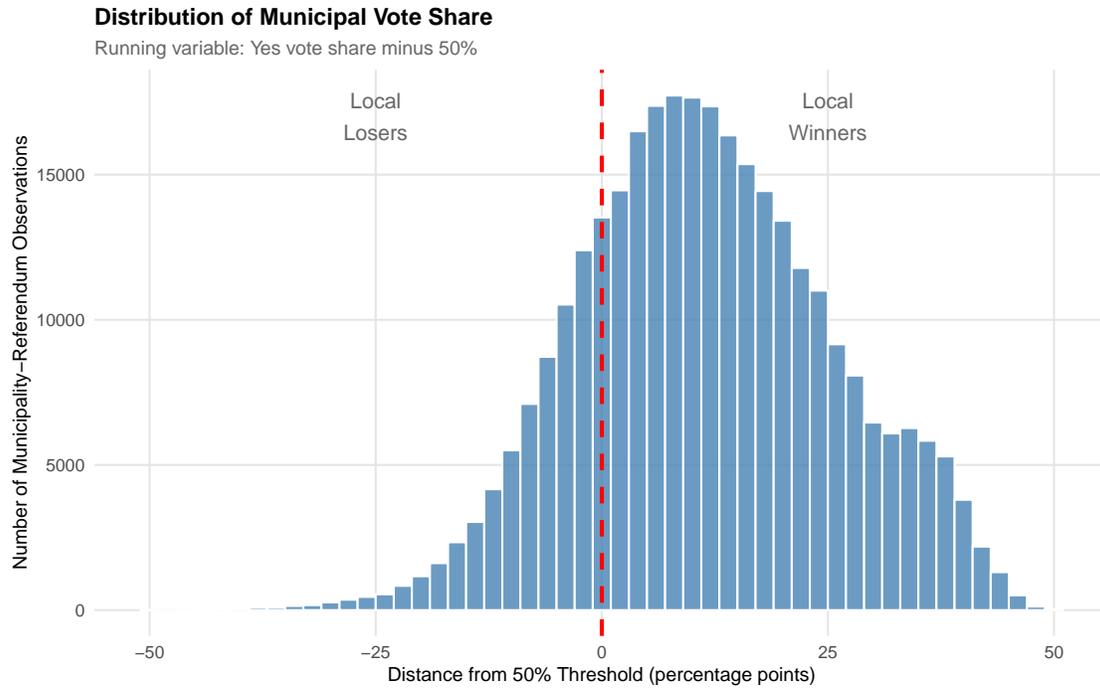


Figure 2: Density of Running Variable

Notes: Histogram of the running variable (yes-share minus 50%). The vertical dashed line indicates the 50% cutoff. The McCrary density test ($p = 0.92$) provides no evidence of manipulation or bunching at the threshold.

D.2 RDD Visualization

Figure 3 shows the regression discontinuity plot with binned means and local polynomial fits.

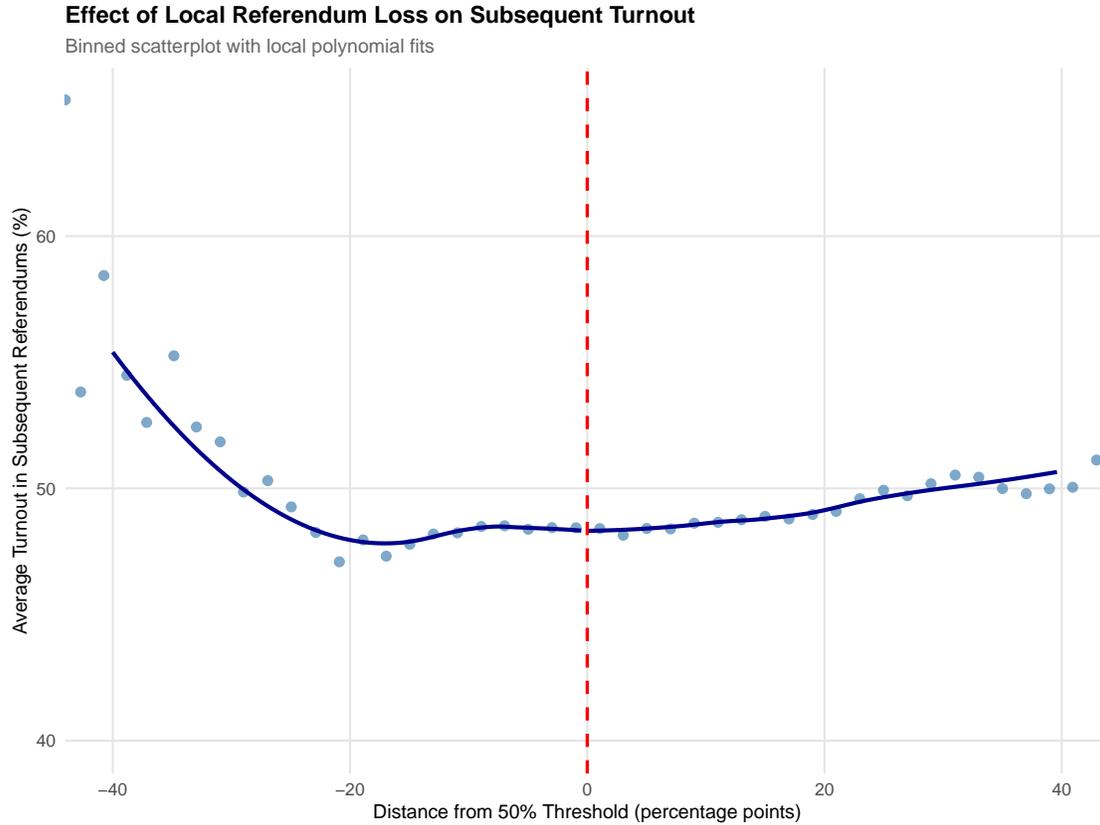


Figure 3: Regression Discontinuity Plot: Subsequent Turnout

Notes: Each point represents the average subsequent turnout in a 2-percentage-point bin of the running variable. Solid lines show local polynomial fits on each side of the 50% cutoff. The absence of a visible jump at the cutoff corresponds to the statistically insignificant RDD estimate.

D.3 Bandwidth Sensitivity Plot

Figure 4 shows how the RDD estimate and confidence intervals change across bandwidth choices.

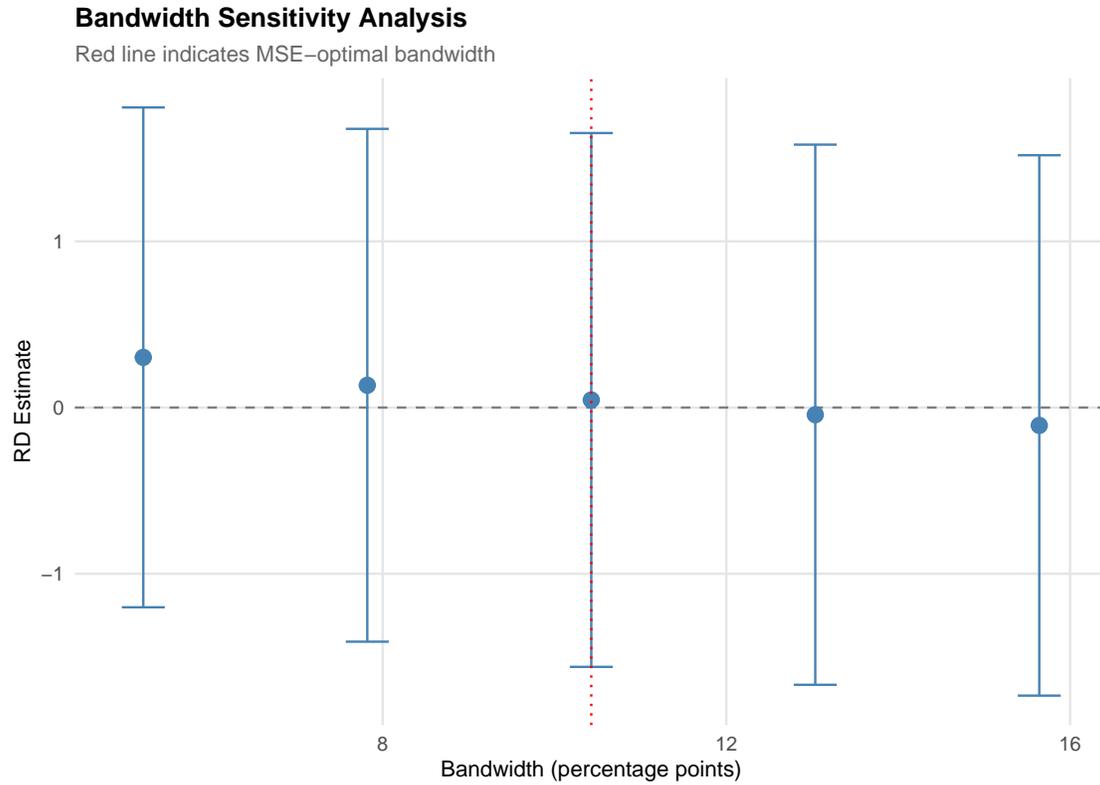


Figure 4: Bandwidth Sensitivity

Notes: Point estimates (circles) and 95% robust confidence intervals (vertical bars) for bandwidths ranging from 0.5 to 1.5 times the MSE-optimal bandwidth. The red dotted line indicates the MSE-optimal bandwidth. The horizontal dashed line at zero indicates no effect. All confidence intervals include zero, indicating robustness to bandwidth choice.