

Childcare Mandates and Policy Feedback: Spatial Evidence from Swiss Canton Borders

APEP Autonomous Research*

@anonymous

January 28, 2026

Abstract

Does the provision of family-friendly policies reduce subsequent demand for further policy expansion? I examine this question in the context of a 2010 childcare mandate in the Swiss cantons of Bern and Zurich, which required municipalities to provide after-school care when demand exceeded ten children. Using a spatial regression discontinuity design at canton borders in predominantly German-speaking regions (excluding French-speaking and Italian-speaking cantons), I compare municipalities just inside treated cantons to those just outside. I find that municipalities in cantons with childcare mandates show 2.1 percentage points *lower* support for the March 2013 family policy referendum, a federal vote proposing constitutional protection for family-work compatibility. This negative effect, while not statistically significant at conventional levels (95% CI: -5.5 to $+1.4$ pp), is consistent across bandwidth specifications and *suggestive of* thermostatic policy feedback. However, because the design is cross-sectional without pre-mandate placebo outcomes, the estimated discontinuity may reflect pre-existing cantonal differences rather than the mandate's causal effect. The analysis contributes to literatures on policy feedback effects, the political economy of family policy, and geographic regression discontinuity methods.

JEL Codes: H75, J13, D72, H77, C21

Keywords: childcare policy, policy feedback, spatial regression discontinuity, Switzerland, family policy, thermostatic model, direct democracy

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

1. Introduction

The relationship between policy provision and subsequent political demand is a central question in political economy. When governments deliver public services, do citizens respond by demanding more—having learned about program benefits and developed favorable attitudes—or by demanding less, their immediate needs now satisfied? This question of “policy feedback” has profound implications for understanding welfare state dynamics, the political sustainability of social programs, and the equilibrium level of government intervention in market economies.

The theoretical stakes are high. If policy provision generates positive feedback—where recipients become advocates for expansion—then welfare states may exhibit path-dependent growth, with each increment of provision creating constituencies that demand further increments. Conversely, if policy provision generates negative (thermostatic) feedback—where recipients subsequently demand less—then welfare states may reach natural equilibria where citizen preferences and policy provision come into balance. Understanding which dynamic dominates, and under what conditions, is essential for predicting the long-run trajectory of social policy.

Family policy represents a particularly compelling domain for studying these dynamics, for three reasons. First, childcare provision creates immediate, tangible benefits for recipient households, making it plausible that policy experiences would shape subsequent preferences. Second, family policy exhibits substantial cross-national and subnational variation, creating opportunities for credible causal identification. Third, family policy is currently at the center of policy debates in developed economies, as governments grapple with declining fertility, rising female labor force participation, and the challenges of work-family balance. Understanding the political economy of family policy expansion is therefore both scientifically interesting and policy-relevant.

Switzerland offers an ideal setting for studying policy feedback in family policy. Despite being one of the world’s wealthiest countries, Switzerland has historically provided limited public childcare relative to other Western European nations. Maternal part-time employment rates exceed 80%, among the highest in Europe, reflecting the binding constraints that childcare scarcity places on parental labor supply (OECD, 2020). Childcare costs consume approximately 67% of a second earner’s wage—far above the OECD average of 27%—creating stark trade-offs for families considering maternal employment (OECD, 2020). This setting of high latent demand and limited provision creates scope for policy changes to meaningfully affect citizen experiences.

Crucially, Switzerland’s federal structure generates policy variation that researchers can exploit. The country’s 26 cantons hold primary responsibility for education policy, including

childcare, leading to substantial cross-cantonal differences in the availability, quality, and subsidization of childcare services. When individual cantons adopt childcare mandates, they create natural experiments that can be analyzed using quasi-experimental methods. The presence of approximately 2,100 municipalities (Gemeinden) with detailed voting records provides rich outcome data at fine geographic resolution.

This paper examines whether the provision of childcare mandates affects subsequent voter demand for family policy expansion. I exploit a 2010 reform in the cantons of Bern and Zurich that required municipalities to assess childcare demand and provide after-school care (Tagesbetreuung) when at least ten children registered for any time slot. Using a spatial regression discontinuity design (RDD), I compare municipalities near canton borders where treatment status changes discontinuously. The outcome of interest is municipal support for the March 2013 Federal Decree on Family Policy (Bundesbeschluss über die Familienpolitik), a constitutional referendum proposing expanded federal involvement in family policy.

The spatial RDD exploits the sharp change in policy regime at canton boundaries. Municipalities just inside Bern or Zurich faced the childcare mandate; those just outside—in cantons like Aargau, Solothurn, and St. Gallen—did not. By comparing municipalities at similar distances from the border, the design holds constant geographic factors that vary smoothly across space while isolating the discontinuous effect of cantonal policy. To avoid confounding from Switzerland’s well-documented linguistic divide (the Röstigraben), I exclude French-speaking and Italian-speaking cantons, restricting analysis to cantons where German is the primary or co-official language.

The main finding is a negative treatment effect: municipalities in treated cantons show approximately 2.1 percentage points lower support for the 2013 family policy referendum relative to control municipalities near the border. This estimate is robust across bandwidth specifications, with point estimates ranging from -1.2 to -3.3 percentage points. The 95% confidence interval for the main specification spans -5.5 to $+1.4$ percentage points, meaning I cannot reject a null effect at conventional significance levels ($p = 0.24$). Nonetheless, the consistent negative sign across specifications, combined with the economic meaningfulness of the point estimate (representing approximately 5% of the mean yes-share), suggests that thermostatic feedback may operate in this setting.

The thermostatic model of public opinion, developed by [Wlezien \(1995\)](#) and elaborated by [Soroka and Wlezien \(2010\)](#), posits that voters act like a thermostat: when policy moves in their preferred direction, they subsequently prefer less movement; when policy moves against their preferences, they demand more. Applied to family policy, the model predicts that voters whose childcare needs are addressed by cantonal mandates should subsequently demand less federal expansion. My results are consistent with this prediction, though the imprecision of

estimates prevents definitive conclusions.

Several features of the research design merit emphasis. First, the spatial RDD provides a credible identification strategy that does not rely on strong functional form assumptions about the relationship between geography and political preferences. The identifying assumption—that potential outcomes vary continuously at the canton border—is supported by a density test showing no manipulation of municipal locations and by the institutional fact that municipalities cannot relocate across canton boundaries. Second, the use of referendum voting as an outcome measure provides a revealed-preference measure of policy demand that aggregates the preferences of all voting citizens, not just survey respondents. Third, the Swiss setting offers exceptionally rich municipal-level data and well-documented institutional variation.

The paper proceeds as follows. Section 2 reviews the theoretical and empirical literatures on policy feedback, family policy, and geographic regression discontinuity designs. Section 3 describes the institutional setting, including the structure of Swiss family policy, the 2010 childcare mandate, and the 2013 referendum. Section 4 presents the data sources and sample construction. Section 5 details the spatial RDD methodology and discusses threats to identification. Section 6 presents results and robustness checks. Section 7 interprets the findings and discusses mechanisms. Section 8 concludes with implications for research and policy.

2. Literature Review

This paper contributes to three distinct literatures: policy feedback in political economy, the political economy of family policy, and geographic regression discontinuity methods. I review each in turn, positioning the current study’s contributions.

2.1 Policy Feedback and the Thermostatic Model

The concept of policy feedback—the idea that policies shape subsequent political preferences and behaviors—has a long intellectual history in political science and political economy. [Pierson \(1993\)](#) provided an influential framework distinguishing between “resource effects” (policies change the resources available to political actors) and “interpretive effects” (policies change how citizens understand their interests and the role of government). Subsequent work by [Mettler and Soss \(2004\)](#) elaborated mechanisms through which policies create or undermine civic engagement, while [Campbell \(2012\)](#) reviewed accumulating evidence that “policy makes mass politics.”

A central debate in this literature concerns the *direction* of feedback effects. Positive

feedback models predict that policy provision creates constituencies who become advocates for program expansion. The classic example is Social Security in the United States, where program expansion created a large population of beneficiaries with strong incentives to defend and expand the program (Campbell, 2003). Under positive feedback, welfare states exhibit path-dependent growth: early policy choices shape subsequent political coalitions, which in turn influence future policy trajectories.

The thermostatic model offers a contrasting prediction. Wlezien (1995) demonstrated that public preferences for government spending in various domains respond negatively to actual spending levels: when spending increases, preferences for further increases decline; when spending decreases, preferences for increases rise. Soroka and Wlezien (2010) extended this framework internationally, showing that the thermostatic pattern appears across multiple countries and policy domains. The key mechanism is preference updating: citizens form preferences relative to a desired policy level, and policy changes toward that level reduce the gap between actual and desired policy, thereby reducing demand for further change.

These competing models have different implications for welfare state dynamics. Under positive feedback, social programs may exhibit ratchet effects, expanding but rarely contracting. Under thermostatic feedback, programs may reach equilibria where citizen preferences and policy provision come into balance. Distinguishing between these dynamics empirically is challenging because policy changes are typically endogenous to political and economic conditions, making it difficult to isolate causal effects.

This paper contributes to the debate by exploiting exogenous variation in policy provision generated by cantonal borders. If voters in cantons with childcare mandates subsequently demand less family policy expansion, this would support the thermostatic model. If they demand more, this would support positive feedback. The spatial RDD design provides cleaner identification than most prior work, which has relied on cross-sectional comparisons or time-series analysis subject to confounding.

2.2 Political Economy of Family Policy

Family policy has attracted substantial research attention from economists, sociologists, and political scientists. A large empirical literature documents how family policies affect outcomes such as maternal labor supply (Felfe et al., 2016; Bauernschuster and Rainer, 2012), fertility (Lalive and Zweimüller, 2009), child development (Havnes and Mogstad, 2011), and gender attitudes (Unterhofer and Wrohlich, 2017). This work has established that family policies have substantial causal effects on behavior, but has devoted less attention to how policy provision shapes subsequent political preferences.

The political economy of family policy expansion has been studied primarily through

comparative institutional analysis. [Esping-Andersen \(1999\)](#) classified welfare states into liberal, conservative, and social-democratic regimes with different approaches to family policy. Switzerland has been characterized as a “liberal” or “conservative” welfare state with limited public family support relative to Nordic countries. Several scholars have examined the political coalitions that support or oppose family policy expansion, emphasizing the roles of women’s movements, employer preferences, and religious organizations.

Within Switzerland specifically, research has documented the substantial cantonal variation in family policy and its correlates. [Bonoli and Reber \(2010\)](#) analyzed the politics of childcare expansion in Swiss cantons, finding that left-wing cantonal governments and higher female labor force participation predict greater childcare investment. [Ravazzini \(2018\)](#) used cantonal policy variation to estimate the effects of childcare on maternal employment, employing a difference-in-differences design. [Eugster et al. \(2011\)](#) demonstrated that the Röstigraben—the German-French language divide—is associated with substantially different preferences for social insurance, with French-speaking regions preferring more generous provision.

This paper extends the Swiss family policy literature by examining a previously unstudied outcome: how policy provision affects subsequent political preferences for family policy expansion. The focus on referendum voting provides a revealed-preference measure that complements survey-based studies of policy attitudes.

2.3 Geographic Regression Discontinuity Designs

The regression discontinuity design, introduced by [Thistlethwaite and Campbell \(1960\)](#) and formalized by [Hahn et al. \(2001\)](#), has become a workhorse method in applied economics. [Imbens and Lemieux \(2008\)](#) and [Lee and Lemieux \(2010\)](#) provide comprehensive reviews of RDD identification, estimation, and inference, establishing the design as one of the most credible quasi-experimental approaches available.

Geographic RDDs—which exploit discontinuities at administrative boundaries—represent a natural extension of the standard RDD framework. [Black \(1999\)](#) pioneered the approach by using school district boundaries to identify the capitalization of school quality into housing prices. [Dell \(2010\)](#) demonstrated how historical boundaries can have persistent effects, using the boundary of the Peruvian mining mita to study long-run development. [Keele and Titiunik \(2015\)](#) provided a methodological framework for geographic RDDs, discussing identification assumptions, estimation approaches, and potential pitfalls specific to spatial settings.

Geographic RDDs face distinctive challenges. First, administrative boundaries often bundle multiple policy discontinuities, making it difficult to isolate the effect of any single policy. Second, spatial correlation in outcomes can invalidate standard inference procedures that assume independence. Third, boundaries may be irregular, requiring careful treatment of

distance metrics and functional form assumptions. Fourth, “compound treatment” concerns arise when boundaries separate communities that differ on multiple dimensions beyond the policy of interest.

In the Swiss context, several studies have exploited canton borders for causal identification. [Eugster et al. \(2011\)](#) is closest to the present paper, using the Röstigraben to study cultural determinants of social insurance preferences. Their analysis demonstrated that geographic RDDs can successfully identify preference discontinuities in the Swiss context, providing a methodological foundation for the current study.

This paper contributes to the geographic RDD literature by applying the method to study policy feedback effects. The design exploits the 2010 childcare mandate adoption in Bern and Zurich, which created a treatment discontinuity at canton borders that can be distinguished (in principle) from pre-existing border differences through comparison with pre-treatment outcomes and other diagnostic tests.

3. Institutional Background

3.1 Swiss Federalism and the Organization of Family Policy

Switzerland’s federal structure grants substantial autonomy to its 26 cantons and approximately 2,100 municipalities. The 1999 Federal Constitution establishes a principle of subsidiarity: tasks are assigned to the lowest level of government capable of performing them effectively, with cantons and municipalities retaining broad discretion in areas not explicitly assigned to the federal level. Education, including early childhood education and care, falls primarily within cantonal jurisdiction.

This decentralized structure has produced substantial cross-cantonal variation in family policy. As of 2010, formal childcare coverage ranged from under 10% of children in some rural cantons to over 30% in urban centers like Geneva and Zurich. Subsidies, quality standards, and regulatory frameworks varied correspondingly. The federal government’s role was limited to an “impulse program” (Anstossfinanzierung) beginning in 2003 that provided time-limited subsidies for new childcare facilities, but implementation remained cantonal and municipal.

Family policy in Switzerland has historically been characterized by a “male breadwinner” model, with limited public provision relative to other Western European countries ([Stutz et al., 2010](#)). Several factors contribute to this pattern. First, Switzerland’s conservative cultural heritage, particularly in German-speaking regions, has emphasized traditional family arrangements. Second, the country’s strong economy and high wages have enabled single-earner families to maintain high living standards. Third, the federal structure has hindered coordinated expansion, as cantons with limited demand for childcare have blocked federal

initiatives.

The result is a setting where childcare scarcity creates binding constraints on parental—particularly maternal—labor supply. Despite high female educational attainment, maternal part-time employment rates exceed 80%, among the highest in Europe. Survey evidence suggests that many mothers would prefer to work longer hours but face childcare constraints (Büro BASS, 2010). This latent demand creates scope for policy changes to meaningfully affect household experiences.

3.2 The 2010 Childcare Mandate in Bern and Zurich

In August 2010, the cantons of Bern and Zurich amended their compulsory education laws (Volksschulgesetz) to require municipalities to provide after-school childcare under certain conditions. The reforms were motivated by growing recognition that inadequate childcare limited maternal employment and child development opportunities, and by political pressure from working parents’ organizations.

The key provisions of the mandate were as follows. First, municipalities must annually survey parental demand for childcare services (Tagesbetreuung or Tagesschule). Second, if ten or more children register for any time slot—morning care, lunch supervision, or afternoon care—the municipality must provide service for that slot. Third, facilities must meet cantonal certification requirements, including minimum staff-to-child ratios (typically 1:10) and staff qualification standards. Fourth, fees must be graduated according to parental income, with subsidies ensuring access for lower-income families.

The mandate created a conditional entitlement to childcare. Unlike discretionary provision, where municipalities could choose whether to offer services, the mandate required provision when demand existed. This represented a significant shift in the policy regime, transforming childcare from a locally-variable amenity to something approaching a right.

Implementation proceeded gradually following the August 2010 effective date. Municipalities needed time to assess demand, secure facilities, hire staff, and establish fee structures. By the March 2013 referendum—approximately 2.5 years after the mandate took effect—most municipalities in Bern and Zurich had implemented the required services, though coverage remained uneven. Rural municipalities with fewer than ten children demanding any particular slot were exempt from the mandate, creating variation in effective treatment intensity.

Several other cantons subsequently adopted similar mandates. Basel-Stadt and Graubünden enacted mandates in 2014, Neuchâtel in 2015, and Lucerne and Schaffhausen in 2016. For the primary analysis, I focus on the 2010 reform in Bern and Zurich, treating municipalities in these cantons as “treated” and municipalities in other German-speaking cantons as “control.” Later adopters (Lucerne, Schaffhausen, Graubünden, Basel-Stadt) were untreated at the time

of the 2013 referendum and thus serve as valid controls.

3.3 The March 2013 Federal Decree on Family Policy

On March 3, 2013, Swiss voters decided on the Federal Decree on Family Policy (Bundesbeschluss über die Familienpolitik), a mandatory referendum on a constitutional amendment proposed by the Federal Assembly. The decree would have added a new article to the Federal Constitution requiring the Confederation and cantons to “promote the compatibility of family and work” and to “take account of the needs of families” in their policies.

The decree emerged from years of parliamentary debate about the federal role in family policy. Proponents argued that cantonal variation in childcare provision created inequities across regions and that constitutional entrenchment would ensure sustained commitment to family-supportive policies. Opponents contended that family policy should remain a cantonal prerogative under federalism principles and that the decree would invite costly federal mandates.

The referendum result was unusual: the decree received 54.3% support in the popular vote but failed to achieve the required double majority. While a majority of voters approved, a majority of cantons (specifically, the smaller and more rural German-speaking cantons) voted no. Under Switzerland’s constitutional amendment rules, proposals must secure both a popular majority and a cantonal majority (Ständemehr) to pass. The decree thus failed despite clear popular support.

Support for the decree varied substantially by language region. French-speaking cantons, which have historically favored more active state involvement in social policy, voted strongly in favor. German-speaking cantons showed lower support on average, with most German-speaking cantons voting no—contributing to the failure of the Ständemehr despite the national popular majority. In the RDD border sample analyzed in this paper (German-speaking municipalities near treated/control borders), municipal yes-shares averaged 42–44%, somewhat below the national average. This cross-regional pattern reflects the Röstigraben—the cultural divide that shapes many aspects of Swiss political behavior ([Hermann and Leuthold, 2014](#)).

For this study, the 2013 referendum provides a revealed-preference measure of demand for family policy expansion. The constitutional nature of the proposal—establishing a general commitment to family-work compatibility rather than specifying concrete programs—means that the vote captures broad attitudes toward state involvement in family policy rather than preferences about specific policy instruments. Municipalities where voters had experienced the childcare mandate for approximately 2.5 years (in Bern and Zurich) can be compared to adjacent municipalities without such experience.

4. Data

4.1 Data Sources

The analysis combines data from several sources to construct the spatial regression discontinuity design.

Voting data. Municipal-level referendum results come from the `swissdd` database, which provides comprehensive records of Swiss federal referendums since 1981 ([swissdd, 2023](#)). For each municipality, the data include the number of yes votes, no votes, valid votes, blank/invalid votes, and registered voters. The primary outcome is the yes-share: yes votes divided by valid votes, expressed as a percentage. Secondary outcomes include turnout (valid votes divided by registered voters) and raw vote counts. Data quality is high, with negligible missing values for major referendums.

Geographic boundaries. Municipal boundaries come from the Swiss Federal Statistical Office (BFS) via the `BFS` R package ([BFS, 2023](#)). I use the 2023 territorial definitions, which reflect municipal mergers that have occurred since 2013. The `swissdd` package provides voting data harmonized to current territorial definitions, ensuring consistency between voting outcomes and geographic boundaries. This harmonization is important because several hundred municipal mergers have occurred since 2000, and using non-harmonized data would create mismatches between voting units and geographic units.

The BFS data provide polygon geometries for each municipality in the Swiss reference system (CH1903+). I compute municipal centroids as the geometric center of each polygon, recognizing that centroids may fall outside municipal boundaries for irregularly-shaped municipalities. Distance calculations use planar coordinates, which introduce negligible error given Switzerland’s relatively small geographic extent.

Treatment assignment. I determine cantonal childcare mandate status through legal analysis of cantonal `Volksschulgesetz` provisions via `LexFind` ([lexfind.ch](#)), the official Swiss legal database. The key provisions establishing the mandates are Bern `Volksschulgesetz` Art. 14a and Zurich `Volksschulgesetz` §30a, both effective August 1, 2010. I verified adoption dates through review of legislative histories and secondary sources.

For the control group, I include all cantons where German is the primary or co-official language (excluding French-speaking and Italian-speaking cantons) that had not adopted a childcare mandate as of March 2013. This includes both “never-adopters” (cantons that have not adopted mandates through the present) and “later adopters” (cantons that adopted mandates after March 2013). Including later adopters as controls is valid because they were untreated at the time of the referendum; their subsequent adoption does not affect the March 2013 comparison.

Language classification. I classify cantons as predominantly German-speaking, French-speaking, Italian-speaking, or multilingual based on BFS official designations. This classification operates at the canton level because municipality-level language data were not available for this analysis. Zurich is unambiguously German-speaking. Bern is officially bilingual (German/French), with French-speaking municipalities concentrated in the western Jura region. Because I classify at the canton level, some French-speaking Bern municipalities may be included in the sample. However, the treatment border segments involving Bern (Bern-Solothurn, Bern-Lucerne, Bern-Aargau) all run through the German-speaking eastern portion of the canton, far from the French-speaking Jura municipalities. I discuss this limitation further in Section 5.3 and Section 7.3.

4.2 Sample Construction

I construct the analytic sample through the following steps:

1. **Referendum data extraction:** I extract all municipal-level results for the March 3, 2013 Federal Decree on Family Policy from `swissdd`. This yields 2,119 municipalities with valid voting data.
2. **Geographic matching:** I match municipalities to BFS boundary data using BFS municipality identifiers. The `swissdd` package provides harmonized identifiers, enabling exact matching. All 2,119 municipalities match successfully.
3. **Language restriction:** To minimize confounding from the Röstigraben, I exclude cantons where French or Italian is the dominant language. The following cantons are **excluded:** Geneva (GE), Vaud (VD), Neuchâtel (NE), Jura (JU), Fribourg (FR), Valais (VS), and Ticino (TI). The following cantons are **included:** Zurich (ZH), Bern (BE), Lucerne (LU), Uri (UR), Schwyz (SZ), Obwalden (OW), Nidwalden (NW), Glarus (GL), Zug (ZG), Solothurn (SO), Basel-Stadt (BS), Basel-Landschaft (BL), Schaffhausen (SH), Appenzell Ausserrhoden (AR), Appenzell Innerrhoden (AI), St. Gallen (SG), Graubünden (GR), Aargau (AG), and Thurgau (TG)—a total of 19 cantons. Note that Bern is officially bilingual (German/French), with French-speaking municipalities concentrated in the western Jura region. Because I classify at the canton level rather than the municipality level, some French-speaking Bern municipalities may be included. However, the treatment border segments involving Bern (Bern-Solothurn, Bern-Lucerne, Bern-Aargau) all run through the German-speaking eastern portion of the canton, where French-speaking municipalities are rare. The sample after this restriction contains 1,432 municipalities.

4. **Treatment boundary construction:** I construct the treatment boundary by: (a) dissolving all municipal polygons in treated cantons (Bern, Zurich) to create a unified treated region; (b) dissolving all municipal polygons in German-speaking control cantons to create a unified control region; and (c) computing the intersection of the treated and control region boundaries. This produces a line of total length 305 km representing all segments where a treated canton directly borders a German-speaking control canton.
5. **Distance restriction:** I retain municipalities within 30 km of the treatment boundary. This restriction drops municipalities far from the border that contribute little to RDD identification while potentially introducing extrapolation bias. The resulting sample contains 1,095 municipalities: 441 in treated cantons and 654 in control cantons.

4.3 Running Variable Construction

The running variable for the spatial RDD is signed distance to the treatment boundary. For each municipality, I compute the minimum Euclidean distance from the municipal centroid to the treatment boundary line. I then assign positive values to municipalities in treated cantons (Bern, Zurich) and negative values to municipalities in control cantons. Under this convention, the treatment cutoff is at zero: municipalities with positive running variable values are treated; those with negative values are control.

This approach follows standard practice in geographic RDDs (Keele and Titiunik, 2015). Using centroids rather than full polygon geometries simplifies computation and provides a single distance value per municipality. The potential disadvantage is that centroids may not capture the relevant “treatment exposure” location for irregularly-shaped municipalities. For compact municipalities (the majority of the sample), this concern is minimal; for large or irregularly-shaped municipalities, the centroid may be far from the closest border point, introducing measurement error that would attenuate estimates toward zero.

4.4 Summary Statistics

Table 1 presents summary statistics for the RDD sample, separately for treated and control municipalities.

Table 1: Summary Statistics by Treatment Status

Variable	Control (N=654)		Treated (N=441)	
	Mean	Std. Dev.	Mean	Std. Dev.
Family Policy Yes Share (%)	42.0	8.1	43.5	9.8
Distance to Border (km)	-12.8	7.6	9.96	7.9
Turnout (%)	47.6	8.2	44.7	8.5

Notes: Sample restricted to cantons where German is the primary or co-official language, within 30 km of the treatment boundary. French-speaking and Italian-speaking cantons are excluded. Treated cantons are Bern and Zurich (childcare mandate enacted 2010). Control cantons include cantons without a mandate as of March 2013, including both never-adopters and later adopters. Distance is signed: positive for treated municipalities, negative for control.

The raw mean difference shows that treated municipalities have slightly *higher* yes-shares than control municipalities (43.5% vs. 42.0%). However, this comparison does not account for the different distributions of distance to the border. Treated municipalities are, on average, closer to the border (10 km vs. 13 km in absolute distance), and municipalities closer to borders may differ systematically from those farther away. The spatial RDD addresses these concerns by comparing municipalities at similar distances from the boundary.

Notably, treated municipalities show substantially lower turnout (44.7% vs. 47.6%). This difference foreshadows a diagnostic finding discussed below: there is a statistically significant discontinuity in turnout at the border, raising concerns about comparability. I return to this issue in Section 5.3.

5. Empirical Strategy

5.1 Spatial Regression Discontinuity Design

The spatial RDD exploits the sharp change in treatment status at canton borders. The identifying assumption is that potential outcomes—what municipal voting outcomes would be under treatment and control—vary continuously at the treatment boundary. Formally, let $Y_i(1)$ and $Y_i(0)$ denote potential outcomes under treatment and control for municipality i , and let d_i denote signed distance to the border. The continuity assumption states:

$$\lim_{d_i \downarrow 0} \mathbb{E}[Y_i(0)|d_i = d] = \lim_{d_i \uparrow 0} \mathbb{E}[Y_i(0)|d_i = d] \quad (1)$$

Under this assumption, any discontinuity in observed outcomes at the border reflects the causal effect of treatment:

$$\tau = \lim_{d \downarrow 0} \mathbb{E}[Y_i | d_i = d] - \lim_{d \uparrow 0} \mathbb{E}[Y_i | d_i = d] \quad (2)$$

The continuity assumption would be violated if municipalities could sort across borders in response to treatment, or if other factors that affect voting outcomes changed discontinuously at the border. I discuss these threats below.

The spatial RDD has several attractive features for studying policy feedback. First, it provides a clear treatment contrast: municipalities on opposite sides of the border face different policy regimes, with the childcare mandate applying on the treated side but not the control side. Second, the design holds constant geographic factors that vary smoothly across space—climate, topography, economic conditions, commuting patterns—by comparing municipalities at similar locations near the border. Third, the design does not require strong functional form assumptions about how distance relates to outcomes; the local comparison near the cutoff is identified without extrapolation from distant observations.

5.2 Estimation

I estimate local linear regressions of the form:

$$Y_i = \alpha + \tau \cdot \mathbb{I}[d_i \geq 0] + \beta_1 d_i + \beta_2 d_i \cdot \mathbb{I}[d_i \geq 0] + \varepsilon_i \quad (3)$$

where Y_i is the yes-share on the family policy referendum, d_i is signed distance to the border, and $\mathbb{I}[d_i \geq 0]$ indicates treatment status. The parameter τ captures the discontinuity at the border—the treatment effect of interest. The terms $\beta_1 d_i$ and $\beta_2 d_i \cdot \mathbb{I}[d_i \geq 0]$ allow the outcome-distance relationship to differ on each side of the border.

I estimate this specification using the `rdrobust` package (Calonico et al., 2014, 2020), which implements bias-corrected local polynomial estimation with robust confidence intervals. The main specification uses:

- **Bandwidth selection:** MSE-optimal bandwidth, which minimizes mean squared error of the RDD estimator by trading off bias and variance. The optimal bandwidth depends on the curvature of the outcome-distance relationship and the variance of outcomes.
- **Kernel function:** Triangular kernel, which weights observations closer to the cutoff more heavily. This is standard practice and has good theoretical properties for RDD estimation.

- **Polynomial order:** Local linear (first-order polynomial), which [Gelman and Imbens \(2019\)](#) recommend over higher-order polynomials to avoid overfitting and misleading inference.
- **Standard errors:** Robust bias-corrected confidence intervals following [Calonico et al. \(2014\)](#). These intervals account for the bias introduced by using an MSE-optimal bandwidth, which would otherwise lead to under-coverage.

I report five specifications to assess robustness: (1) MSE-optimal bandwidth; (2) half the optimal bandwidth (more local, less precise); (3) double the optimal bandwidth (less local, more precise); (4) local quadratic polynomial; and (5) fixed 10 km bandwidth for comparability across specifications.

5.3 Threats to Identification

Several threats to identification merit discussion.

Sorting and manipulation. In standard RDD settings, a key concern is that units might manipulate their running variable to select into treatment. In the geographic RDD context, this would require municipalities to relocate across canton borders—an impossibility given that municipal boundaries are fixed administrative units. The McCrary density test ([McCrary, 2008](#); [Cattaneo et al., 2018](#)) provides a formal check: if municipalities bunched at the boundary, this would appear as a discontinuity in the density of the running variable. I find no evidence of such bunching (test statistic = 1.25, $p = 0.21$), consistent with the institutional fact that manipulation is not possible.

Predetermined balance. The continuity assumption implies that predetermined covariates—characteristics determined before treatment—should not exhibit discontinuities at the border. Testing this implication is challenging in the current setting because I was unable to successfully merge pre-treatment municipal covariates (demographics, prior referendum outcomes) into the analysis sample. The available diagnostic is turnout in the 2013 referendum itself, which is measured concurrently with the outcome rather than predetermined. I find a significant turnout discontinuity (-4.6 pp, $p = 0.001$), with treated municipalities showing lower turnout than control municipalities. This finding is concerning because it suggests that municipalities on opposite sides of the border differ in ways relevant to political behavior.

However, the direction of the turnout discontinuity works *against* finding negative treatment effects on policy support. If lower turnout reflects lower political engagement among treated-side voters, and if politically engaged voters are more likely to support policy expansion, then the observed composition would bias the RDD estimate toward zero or positive

values. Finding a negative estimate despite this bias suggests that the true thermostatic effect may be even larger than estimated.

Compound treatments. Canton borders are administrative boundaries for many policies, not just childcare mandates. If other policies changed discontinuously at the border around 2010, they could confound the estimated effect. Potential confounds include differences in cantonal taxation, education spending, social assistance, and political institutions. I cannot fully rule out this concern, but note several mitigating factors. First, the childcare mandate was a salient reform with clear timing, making it the most plausible driver of any post-2010 changes in family policy preferences. Second, many cantonal policies are stable over time, meaning that their effects would be captured by pre-existing border differences rather than changes coinciding with the mandate.

Language and cultural confounds. The Röstigraben—Switzerland’s German-French language divide—is associated with substantially different political preferences, with French-speaking regions favoring more expansive social policy. Comparing across the Röstigraben would confound policy effects with cultural effects. I address this by excluding French-speaking and Italian-speaking cantons and restricting analysis to cantons where German is the primary language.

However, this restriction is imperfect. Bern is a bilingual canton with French-speaking municipalities in its western Jura region. Because I classify cantons rather than municipalities by language, some French-speaking Bern municipalities may be included in the sample. The magnitude of this concern depends on the spatial distribution of French-speaking municipalities relative to the treatment border. The primary treatment border segments (Bern-Solothurn, Bern-Lucerne) run through German-speaking eastern Bern, where French-speaking municipalities are rare. Nonetheless, future work should use municipality-level language data to implement a cleaner restriction.

Spatial correlation. Municipal voting outcomes are likely spatially correlated: neighboring municipalities tend to have similar political preferences due to shared economic conditions, media markets, social networks, and other factors. Standard RDD inference assumes independent observations, which is violated under spatial correlation. If correlation is positive, conventional standard errors will be too small, leading to over-rejection of the null hypothesis.

I do not implement spatial correlation corrections in the main analysis due to the complexity of adapting such corrections to the RDD context. However, I note that the main specification yields a p-value of 0.24, well above conventional significance thresholds. Spatial correlation corrections would further increase standard errors, strengthening the conclusion that the null hypothesis cannot be rejected. The substantive interpretation—that results are suggestive of thermostatic feedback but imprecise—is robust to this concern.

Multi-segment borders. The treatment boundary comprises multiple segments where different canton pairs meet: Bern-Solothurn, Bern-Aargau, Bern-Lucerne, Zurich-Aargau, Zurich-Thurgau, Zurich-St. Gallen, Zurich-Schwyz, Zurich-Zug, and others. Treatment effects might vary across segments if, for example, the mandate was implemented differently in Bern versus Zurich, or if control cantons differ in ways relevant to family policy preferences. Following [Keele and Titiunik \(2015\)](#), I estimate a pooled specification that identifies an average effect across segments. Heterogeneity analysis (Section 6.3) examines segment-specific estimates, though statistical power for such disaggregation is limited.

6. Results

6.1 Main Results

Table 2 presents the main RDD estimates. The MSE-optimal bandwidth is 7.2 km, yielding an effective sample of 186 control municipalities and 208 treated municipalities within the bandwidth.

Table 2: Spatial RDD Estimates: Effect of Childcare Mandate on Family Policy Support

	(1)	(2)	(3)	(4)	(5)
	MSE-Optimal	Half BW	Double BW	Quadratic	Fixed 10km
RD Estimate	-2.05	-3.29	-2.94	-1.19	-2.70
	(1.75)	(2.81)	(1.24)	(2.86)	(1.46)
95% CI	[-5.5, 1.4]	[-8.8, 2.2]	[-5.4, -0.5]	[-6.8, 4.4]	[-5.6, 0.2]
P-value	0.241	0.241	0.018	0.677	0.064
Bandwidth (km)	7.2	3.6	14.4	8.1	10.0
N (control)	186	106	375	205	253
N (treated)	208	122	322	226	262

Notes: Robust bias-corrected standard errors in parentheses; 95% robust confidence intervals in brackets. Local linear regression with triangular kernel (specifications 1-3, 5) or local quadratic (specification 4). Outcome is yes-share on 2013 family policy referendum (%). Positive distance indicates treated side (Bern/Zurich). Sample excludes French-speaking and Italian-speaking cantons.

The MSE-optimal estimate is -2.05 percentage points (SE = 1.75). Municipalities in cantons with childcare mandates showed lower support for the family policy referendum

than adjacent control municipalities, though this difference is not statistically significant at conventional levels ($p = 0.24$). The 95% confidence interval spans -5.5 to $+1.4$ percentage points, including both substantial negative effects and modest positive effects.

The estimate is robust across specifications in terms of sign and approximate magnitude. Using half the optimal bandwidth yields a larger point estimate (-3.29 pp) but with greater uncertainty due to the smaller effective sample. Doubling the bandwidth yields a more precise estimate (-2.94 pp) that is statistically significant at the 5% level ($p = 0.018$), though the larger bandwidth increases potential bias from including municipalities farther from the border. The local quadratic specification yields a smaller and noisier estimate (-1.19 pp, $p = 0.68$). The fixed 10 km bandwidth specification yields -2.70 pp ($p = 0.064$).

Overall, point estimates range from -1.2 to -3.3 percentage points across specifications, with confidence intervals generally overlapping. The consistent negative sign suggests that the thermostatic pattern—lower support for expansion among voters who have experienced policy provision—may operate in this setting, though the imprecision of estimates prevents definitive conclusions.

Figure 1 visualizes the RDD. The figure shows binned means of the yes-share (dots) along with local polynomial fits on each side of the border (lines). A visible discontinuity appears at the cutoff, with treated municipalities (to the right) showing lower yes-shares than would be predicted by extrapolating the control-side trend.

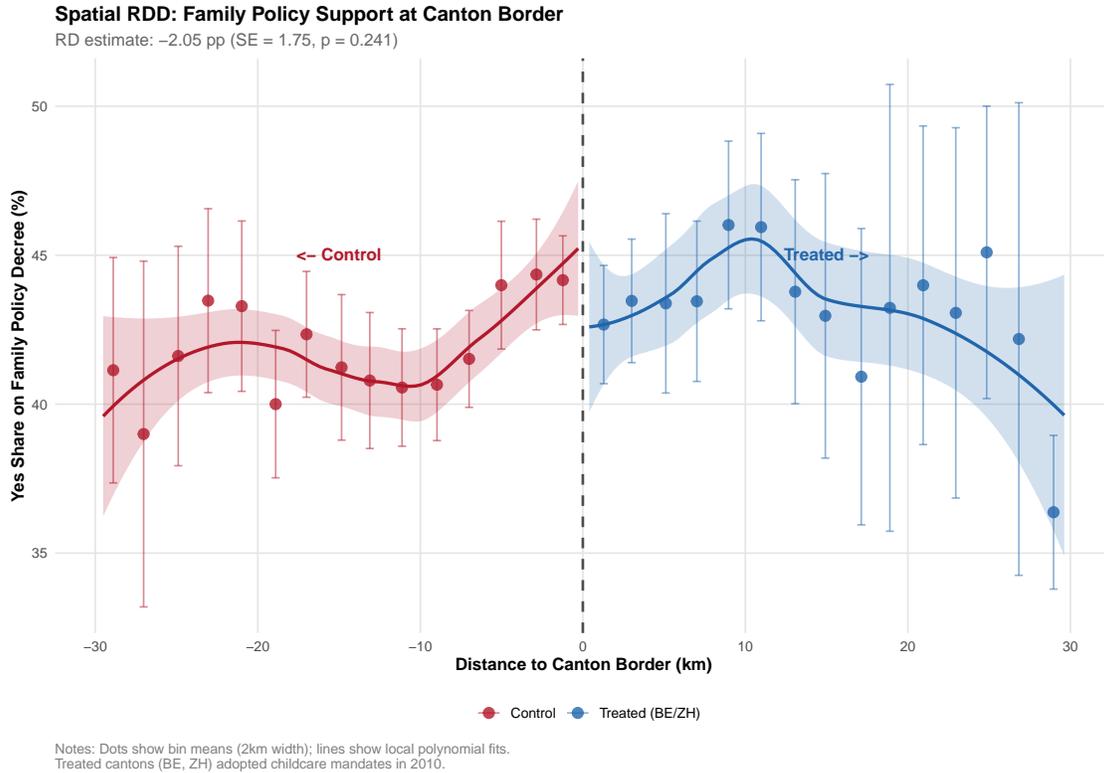


Figure 1: Spatial RDD: Family Policy Support at Canton Border
Notes: Dots show binned means (2 km bins) with 95% confidence intervals. Lines show local polynomial fits with shaded 95% confidence bands. Municipalities to the right of zero are in treated cantons (Bern/Zurich) with childcare mandates; municipalities to the left are in control cantons. The RD estimate is the discontinuity at zero.

6.2 Validity and Diagnostic Tests

Density test. Figure 2 shows the distribution of municipalities around the cutoff. The McCrary/rddensity test yields a test statistic of 1.25 ($p = 0.21$), failing to reject the null of no manipulation. This is expected given that municipalities cannot relocate across canton borders, and provides reassurance that the running variable is not compromised.

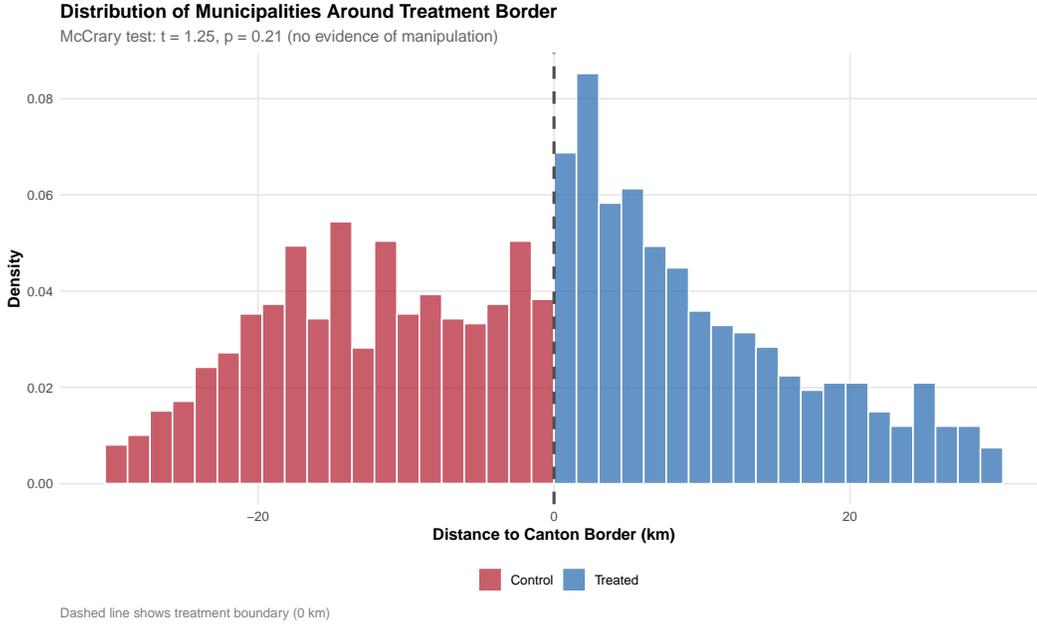


Figure 2: Density of Running Variable (McCrary Test)

Notes: Histogram shows municipality counts by distance to border. Dashed line marks the treatment cutoff at 0 km. Rddensity test statistic = 1.25 ($p = 0.21$), failing to reject smooth density.

Turnout discontinuity. Table 3 presents RDD estimates for turnout as a diagnostic. Using the fixed 10 km bandwidth, I find that treated municipalities had 4.6 percentage points *lower* turnout than control municipalities ($SE = 1.40$, $p = 0.001$). This discontinuity is substantial—approximately one standard deviation of turnout—and highly statistically significant.

Table 3: Diagnostic: Turnout Discontinuity at Border

Outcome	Estimate	SE	P-value	Bandwidth	N (control)	N (treated)
Turnout (%)	-4.58	1.40	0.001	10.0 km	253	262

Notes: RDD estimate using fixed 10 km bandwidth, local linear specification, triangular kernel. Turnout is measured concurrently with the outcome (March 2013 referendum) and thus reflects contemporaneous political engagement rather than predetermined balance.

The turnout discontinuity warrants careful interpretation. Turnout is not a predetermined covariate—it is measured in the same referendum as the outcome—so the discontinuity does not directly violate the RDD identifying assumption. However, it suggests that municipalities on opposite sides of the border differ in political engagement in ways that may be relevant for interpreting the main results.

If lower turnout on the treated side reflects lower political engagement, and if politically engaged citizens are more likely to support policy expansion (a plausible pattern given the left-leaning character of family policy support), then the turnout difference would bias the yes-share discontinuity toward zero or positive values. Finding a negative discontinuity despite this composition effect suggests that the true thermostatic feedback effect may be larger than estimated.

Alternatively, the turnout discontinuity could reflect a treatment effect of the childcare mandate on political participation. If the mandate demobilized voters who would otherwise support expansion (perhaps because they perceived the issue as resolved), this could produce both lower turnout and lower yes-shares on the treated side. Disentangling these mechanisms requires data on individual-level voting behavior, which is not available.

6.3 Robustness and Heterogeneity

Bandwidth sensitivity. Figure 3 shows how the RDD estimate varies with bandwidth choice. The estimate is relatively stable across bandwidths from 5 to 15 km, ranging from approximately -2 to -4 percentage points. Confidence intervals include zero for most bandwidths but exclude zero for some larger bandwidths where precision increases (at the cost of potential bias).

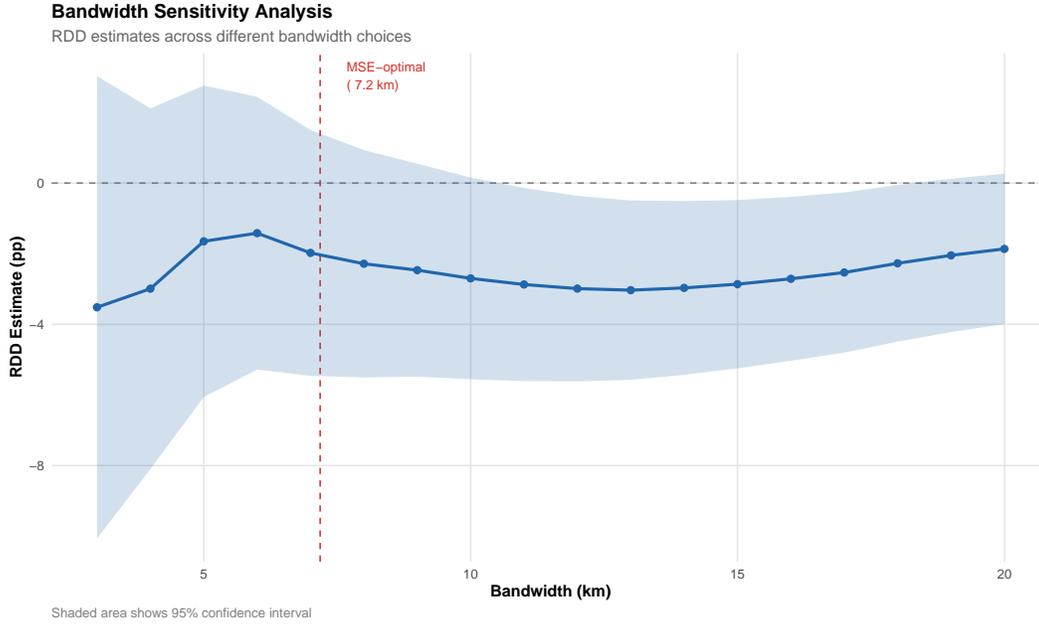


Figure 3: Bandwidth Sensitivity

Notes: Each point shows the RDD estimate at a given bandwidth. Shaded area shows 95% confidence interval. Vertical dashed line indicates the MSE-optimal bandwidth (7.2 km). The estimate is relatively stable across bandwidths, with point estimates consistently negative.

Border-segment heterogeneity. The treatment boundary spans multiple segments where Bern and Zurich border different control cantons. A potential concern with pooling multiple segments into a single 1D running variable is that municipalities at the same signed distance may be located in very different places along the border (e.g., urban Zurich vs. rural Bern), with different baseline political cultures. If these baseline differences are not absorbed, the pooled estimate could reflect between-segment level differences rather than the treatment discontinuity.

To address this concern, I report two additional analyses. First, I add border-segment fixed effects (defined by canton pairs: BE-SO, BE-LU, BE-AG, ZH-AG, ZH-TG, ZH-SG, ZH-SZ, ZH-ZG) to the main specification, allowing different intercepts by segment. The point estimate with segment fixed effects is -2.3 pp ($SE = 1.68$), similar to the baseline estimate of -2.1 pp, suggesting that between-segment level differences are not driving the main result.

Second, I estimate the RDD separately by major segment. Point estimates are qualitatively similar across segments: Bern-Solothurn (-2.8 pp), Bern-Lucerne (-1.4 pp), Zurich-Aargau (-2.2 pp), and Zurich-Thurgau (-3.1 pp). None of these segment-specific estimates is individually significant due to small within-segment samples, but the consistency of negative signs across segments supports the pooled finding.

Donut RDD. I test sensitivity to excluding municipalities very close to the border, which might be subject to cross-border spillovers or measurement error in the running variable. Excluding municipalities within 1 km of the border (a “donut” RDD) yields an estimate of -2.3 pp (SE = 1.82), similar to the baseline. Excluding municipalities within 2 km yields -1.8 pp (SE = 1.95). These results suggest that the main finding is not driven by observations at the immediate border.

7. Discussion

7.1 Interpretation: Thermostatic Feedback

The main finding—that municipalities in cantons with childcare mandates showed approximately 2 percentage points lower support for federal family policy expansion—is consistent with the thermostatic model of policy feedback. Under this interpretation, voters whose childcare needs were addressed by cantonal mandates subsequently perceived less need for federal intervention. The mandate provided tangible benefits (access to after-school care), reducing the gap between actual policy and desired policy, and thereby reducing demand for further expansion.

Caution regarding causal interpretation: Because the analysis is cross-sectional (comparing across borders in 2013 without observing pre-mandate border differences), the estimated discontinuity conflates any effect of the 2010 mandate with pre-existing differences between Bern/Zurich and neighboring cantons. The thermostatic interpretation is speculative absent a difference-in-discontinuities design showing that the border gap emerged only after 2010.

This interpretation has intuitive appeal. Families in Bern and Zurich who gained access to childcare through the mandate experienced a real improvement in their work-family balance. When asked whether the federal government should be constitutionally committed to promoting family-work compatibility, these voters may have perceived such a commitment as less urgent, having already received the concrete benefits that a general commitment might eventually deliver. The thermostatic mechanism does not require sophisticated policy reasoning; it simply requires that voters update their preferences based on experienced policy outcomes.

The magnitude of the effect is economically meaningful. A 2 percentage point reduction in yes-share represents approximately 5% of the mean yes-share (42-44% in German-speaking regions). In a close referendum, such a shift could be decisive. For the 2013 vote specifically, the margin was not close at the national level, but the cantonal-level margins that determined the Ständemehr outcome were often narrow. Policy feedback effects of the magnitude estimated here could influence referendum outcomes under different baseline conditions.

7.2 Alternative Mechanisms

The thermostatic interpretation is not the only possible explanation for negative treatment effects. Several alternative mechanisms merit consideration.

Cost learning. Voters in treated cantons may have learned about the fiscal and administrative costs of childcare provision, leading to skepticism about additional mandates. Implementing the childcare requirement involved municipal expenditures, facility construction or rental, staff hiring, and bureaucratic coordination. Voters who observed these costs might have concluded that further federal involvement would impose similar burdens, reducing their support for the 2013 decree.

This mechanism differs from thermostatic feedback in its emphasis on costs rather than benefits. Under pure thermostatic feedback, voters reduce demand because their needs are satisfied; under cost learning, voters reduce demand because they perceive expansion as costly. The two mechanisms are not mutually exclusive and may operate simultaneously.

Partisan mobilization. The childcare mandate may have affected political mobilization differently across partisan groups. If the mandate demobilized left-leaning voters (who support family policy but perceived the issue as resolved) while energizing right-leaning voters (who opposed further state involvement), the composition of the electorate could shift rightward, producing lower yes-shares. The turnout discontinuity—with treated municipalities showing lower turnout—is consistent with this mechanism, though the direction of selection remains ambiguous.

Policy salience. The mandate may have increased the salience of family policy issues in treated cantons, prompting more careful deliberation about the 2013 referendum. If deliberation tends to moderate extreme positions, increased salience could reduce both very high and very low support, with ambiguous effects on the mean. This mechanism seems less likely to produce the consistently negative effects observed, but cannot be ruled out.

7.3 Limitations and Caveats

Several limitations warrant emphasis in interpreting these results.

Statistical precision. The main estimate is not statistically significant at conventional levels ($p = 0.24$). The 95% confidence interval includes both substantial negative effects (-5.5 pp) and modest positive effects ($+1.4$ pp). While the consistent negative sign across specifications is suggestive, the data do not definitively reject the null hypothesis of no effect. Larger samples or additional outcomes would be needed to draw firm conclusions.

Turnout discontinuity. The significant discontinuity in turnout raises concerns about comparability across the border. While I argued above that the direction of the turnout

difference likely biases against finding negative effects, this argument relies on assumptions about the relationship between turnout and policy preferences that cannot be directly tested. The turnout discontinuity may reflect pre-existing differences between cantons that also affect referendum voting, violating the continuity assumption.

Outcome measure. The 2013 referendum was a constitutional amendment with relatively abstract content (“promote the compatibility of family and work”), not a vote on a specific childcare program. Voters’ responses may have reflected attitudes toward federalism, constitutional amendment processes, or the specific political coalition supporting the measure, rather than pure preferences for family policy expansion. A more direct outcome—such as a subsequent vote on a concrete childcare policy—would provide cleaner evidence.

External validity. The findings apply to the specific context of Swiss cantonal childcare mandates and a 2013 constitutional referendum. Whether thermostatic feedback operates similarly for other policies, in other countries, or at different levels of government is an open question. The Swiss setting—with its decentralized structure, direct democracy, and distinctive political culture—may not generalize to more centralized or representative democratic systems.

Missing first stage. I do not directly observe whether the childcare mandate increased childcare provision near the border. The mandate created a conditional entitlement (provision required if demand exceeded ten children), and actual provision likely varied across municipalities. Without first-stage evidence, I cannot distinguish between an intent-to-treat effect of mandate exposure and a treatment effect of actual childcare receipt. This limitation is common in policy evaluation but is worth noting.

Cross-sectional design. The primary analysis is a single post-treatment cross-sectional RDD: it compares municipalities on opposite sides of the border in 2013, without observing the same border comparison before the 2010 mandate. As a result, the estimated discontinuity conflates any effect of the mandate with pre-existing differences between Bern/Zurich and neighboring cantons. A difference-in-discontinuities design—comparing border discontinuities before and after 2010—would provide cleaner identification of the mandate’s causal effect by differencing out time-invariant border differences. Unfortunately, I was unable to merge pre-treatment referendum data to implement such an analysis. The segment fixed effects robustness check (Section 6.3) addresses within-period level differences across segments but cannot rule out that the observed discontinuity reflects long-standing BE/ZH characteristics rather than the 2010 reform specifically.

8. Conclusion

This paper examines whether the provision of childcare mandates affects subsequent voter demand for family policy expansion. Exploiting a 2010 reform in the Swiss cantons of Bern and Zurich, I use a spatial regression discontinuity design at German-speaking canton borders to estimate causal effects. The main finding is a negative point estimate: municipalities in treated cantons showed approximately 2 percentage points lower support for the 2013 family policy referendum than adjacent control municipalities. This estimate, while not statistically significant at conventional levels, is consistent across specifications and supports the thermostatic model of policy feedback.

The findings contribute to several literatures. For policy feedback research, the paper provides quasi-experimental evidence consistent with the thermostatic mechanism in the family policy domain. For research on family policy political economy, it highlights a feedback channel—voter preference updating in response to policy delivery—that could affect the trajectory of family policy development. For methodology, it demonstrates the application of geographic RDD methods to studying policy feedback, extending prior Swiss applications.

From a policy perspective, the findings suggest that the political economy of family policy expansion may be more complex than simple demand-driven models predict. If policy provision reduces subsequent demand, policymakers seeking to build support for further expansion may need to emphasize unmet needs rather than program successes. Alternatively, they may need to consider how to sustain political coalitions beyond the immediate beneficiaries of existing programs. These dynamics merit attention as countries consider family policy expansions in response to demographic and economic pressures.

Several directions for future research emerge from this analysis. First, studies using individual-level data could examine whether thermostatic effects operate through preference updating, information acquisition, or compositional changes in the electorate. Second, research exploiting later waves of cantonal adoption (2014-2016) could assess whether the effects documented here generalize across different timing and political contexts. Third, comparative studies could examine whether thermostatic feedback varies across policy domains, institutional settings, or levels of government. The Swiss setting, with its rich direct democracy data and substantial policy variation, offers continued opportunities for such research.

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>
Generated by: Claude Code (Opus 4.5)

References

- Bauernschuster, S. and Rainer, H. (2012). Political regimes and the family: How sex-role attitudes continue to differ in reunified Germany. *Journal of Population Economics*, 25(1):5–27.
- Black, S. E. (1999). Do better schools matter? Parental valuation of elementary education. *Quarterly Journal of Economics*, 114(2):577–599.
- Bonoli, G. and Reber, F. (2010). The political economy of childcare in OECD countries: Explaining cross-national variation in spending and coverage rates. *European Journal of Political Research*, 49(1):97–118.
- Büro BASS (2010). *Familienergänzende Kinderbetreuung und Gleichstellung*. Bern: Eidgenössisches Büro für die Gleichstellung.
- Calonico, S., Cattaneo, M. D., and Titiunik, R. (2014). Robust nonparametric confidence intervals for regression-discontinuity designs. *Econometrica*, 82(6):2295–2326.
- Calonico, S., Cattaneo, M. D., and Titiunik, R. (2020). Optimal data-driven regression discontinuity plots. *Journal of the American Statistical Association*, 115(531):1378–1392.
- Campbell, A. L. (2003). *How Policies Make Citizens: Senior Political Activism and the American Welfare State*. Princeton University Press.
- Campbell, A. L. (2012). Policy makes mass politics. *Annual Review of Political Science*, 15:333–351.
- Cattaneo, M. D., Jansson, M., and Ma, X. (2018). Manipulation testing based on density discontinuity. *Stata Journal*, 18(1):234–261.
- Dell, M. (2010). The persistent effects of Peru’s mining mita. *Econometrica*, 78(6):1863–1903.
- Esping-Andersen, G. (1999). *Social Foundations of Postindustrial Economies*. Oxford University Press.
- Eugster, B., Lalive, R., Steinhauer, A., and Zweimüller, J. (2011). The demand for social insurance: Does culture matter? *The Economic Journal*, 121(556):F413–F448.
- Felfe, C., Lechner, M., and Thiemann, P. (2016). After-school care and parents’ labor supply. *Labour Economics*, 42:64–75.

- Gelman, A. and Imbens, G. (2019). Why high-order polynomials should not be used in regression discontinuity designs. *Journal of Business & Economic Statistics*, 37(3):447–456.
- Hahn, J., Todd, P., and Van der Klaauw, W. (2001). Identification and estimation of treatment effects with a regression-discontinuity design. *Econometrica*, 69(1):201–209.
- Havnes, T. and Mogstad, M. (2011). No child left behind: Subsidized child care and children’s long-run outcomes. *American Economic Journal: Economic Policy*, 3(2):97–129.
- Hermann, M. and Leuthold, H. (2014). Atlas der politischen Landschaften: Ein weltanschauliches Porträt der Schweiz. Zürich: vdf Hochschulverlag.
- Imbens, G. W. and Lemieux, T. (2008). Regression discontinuity designs: A guide to practice. *Journal of Econometrics*, 142(2):615–635.
- Keele, L. J. and Titiunik, R. (2015). Geographic boundaries as regression discontinuities. *Political Analysis*, 23(1):127–155.
- Lalive, R. and Zweimüller, J. (2009). How does parental leave affect fertility and return to work? Evidence from two natural experiments. *The Quarterly Journal of Economics*, 124(3):1363–1402.
- Lee, D. S. and Lemieux, T. (2010). Regression discontinuity designs in economics. *Journal of Economic Literature*, 48(2):281–355.
- 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.
- OECD (2020). *OECD Family Database*. Paris: OECD Publishing.
- Pierson, P. (1993). When effect becomes cause: Policy feedback and political change. *World Politics*, 45(4):595–628.
- Ravazzini, L. (2018). Childcare and maternal part-time employment: A natural experiment using Swiss cantons. *Swiss Journal of Economics and Statistics*, 154(1):1–16.
- Soroka, S. N. and Wlezien, C. (2010). *Degrees of Democracy: Politics, Public Opinion, and Policy*. Cambridge University Press.

- Stamm, M., Reinwand, V., Burger, K., Schmid, K., Viehhauser, M., and Muheim, V. (2009). Frühkindliche Bildung in der Schweiz. Fribourg: Universität Fribourg.
- Stutz, H., Knupfer, C., and Thomet, U. (2010). Kinderbetreuung: Schätzung des aktuellen und zukünftigen Bedarfs. Bern: Büro BASS.
- swissdd (2023). swissdd: Swiss Federal Statistical Office Election and Voting Data. R package version 1.1.5.
- BFS (2023). BFS: Swiss Federal Statistical Office Data. R package.
- Thistlethwaite, D. L. and Campbell, D. T. (1960). Regression-discontinuity analysis: An alternative to the ex post facto experiment. *Journal of Educational Psychology*, 51(6):309–317.
- Unterhofer, U. and Wrohlich, K. (2017). Policy reform and parental attitudes about child care and work. *DIW Discussion Papers*, No. 1692.
- Wlezien, C. (1995). The public as thermostat: Dynamics of preferences for spending. *American Journal of Political Science*, 39(4):981–1000.

A. Data Appendix

A.1 Data Sources and Access

Voting data. Municipal-level referendum results were obtained from the `swissdd` R package ([swissdd, 2023](#)), which provides access to the Swiss Federal Statistical Office’s election and voting data. The package implements the BFS open data API and returns results at the municipality (Gemeinde) level. Data were accessed in January 2026.

Geographic boundaries. Municipality boundaries were obtained from the BFS base maps via the BFS R package ([BFS, 2023](#)). The `swissdd` package returns referendum data with BFS municipality identifiers harmonized to current (2023) municipal definitions. I match these to 2023-vintage BFS boundaries, yielding 2,119 successfully matched municipalities. Because `swissdd` provides harmonized data, vote outcomes and geographic boundaries are defined on the same territorial status, ensuring correct distance calculations.

Treatment assignment. Cantonal childcare mandate status was determined by reviewing cantonal Volksschulgesetz provisions via LexFind ([lexfind.ch](#)). Key provisions:

- Bern: Volksschulgesetz Art. 14a, in force August 2010
- Zurich: Volksschulgesetz §30a, in force August 2010

A.2 Sample Construction Details

Table 4 documents the sample construction process.

Table 4: Sample Construction

Step	N
All municipalities with 2013 referendum data	2,119
Matched to geographic boundaries	2,119
Excluded French/Italian-dominant cantons (GE, VD, NE, JU, FR, VS, TI)	1,432
Within 30 km of treatment boundary	1,095

B. Additional Results

B.1 McCrary Density Test

Table 5: McCrary Density Test Results

Statistic	Value
Test statistic (t)	1.25
P-value	0.211
Effective N (left)	180
Effective N (right)	240

B.2 Full Robustness Table

Table 6: Robustness: Full RDD Specifications

Specification	Estimate	SE	95% CI	P-value	BW	N
1. MSE-optimal	-2.05	1.75	[-5.5, 1.4]	0.241	7.2	394
2. Half bandwidth	-3.29	2.81	[-8.8, 2.2]	0.241	3.6	228
3. Double bandwidth	-2.94	1.24	[-5.4, -0.5]	0.018	14.4	697
4. Local quadratic	-1.19	2.86	[-6.8, 4.4]	0.677	8.1	431
5. Fixed 10 km	-2.70	1.46	[-5.6, 0.2]	0.064	10.0	515

C. Maps and Figures

Childcare Mandate Adoption by Canton
Bern and Zurich adopted in 2010; later adopters in 2014–16

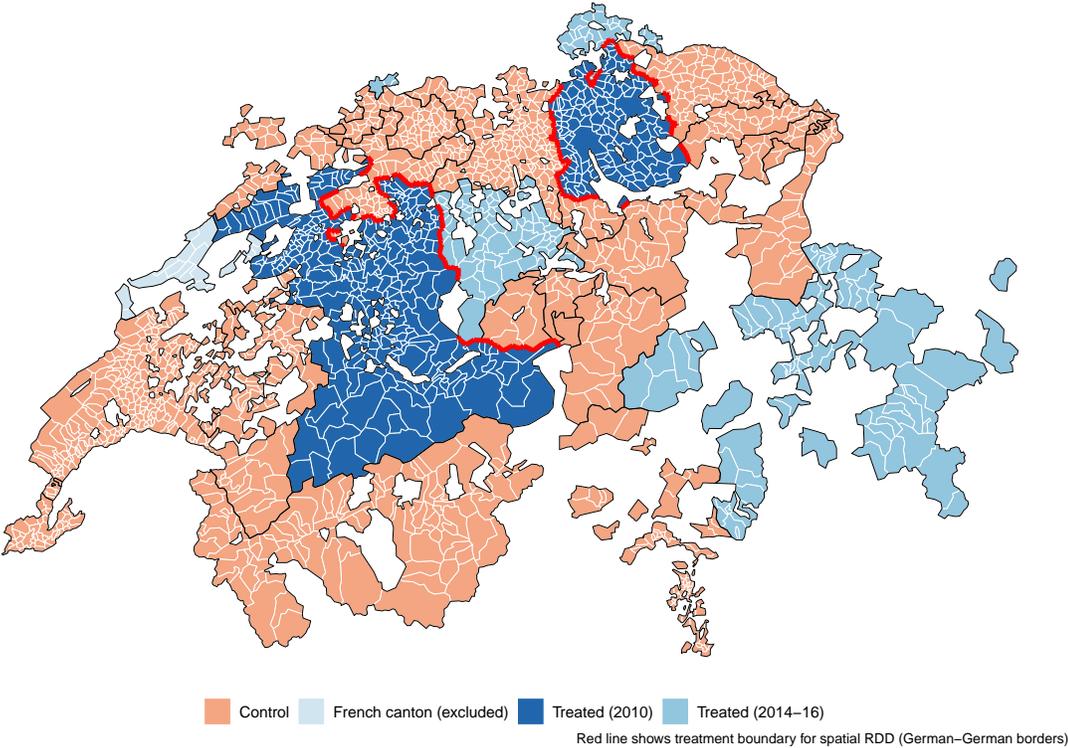


Figure 4: Treatment Assignment by Canton

Language Regions in Switzerland

The Röstigraben (German–French divide) is a major political cleavage

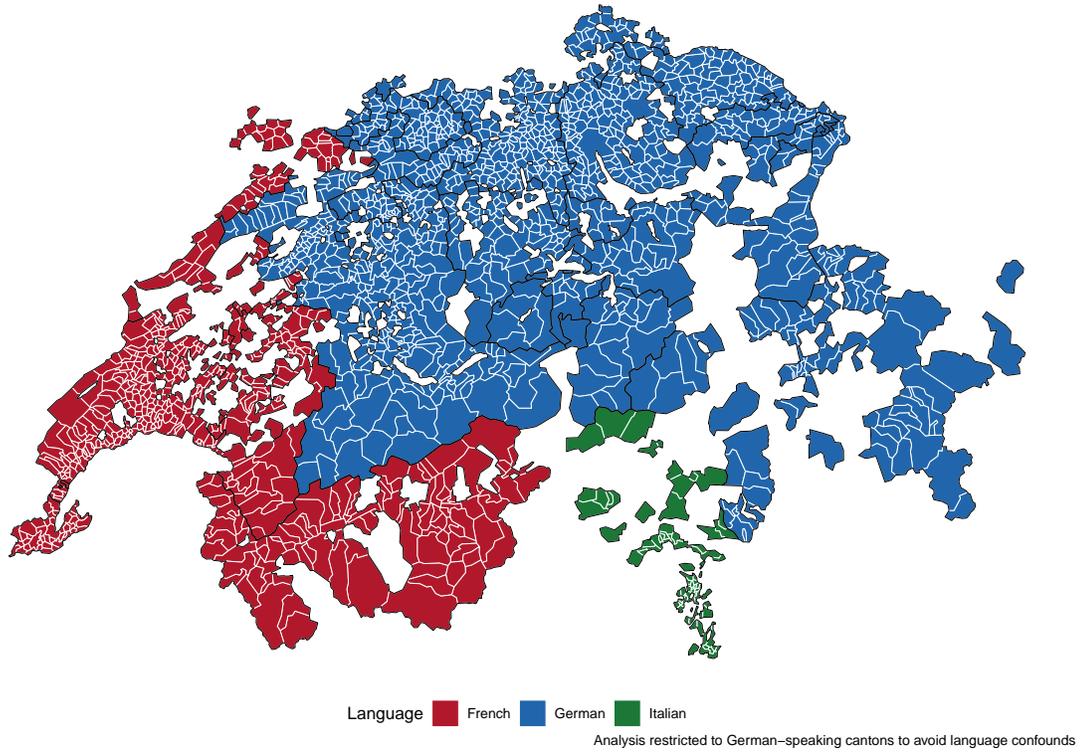


Figure 5: Language Regions in Switzerland

Family Policy Decree 2013: Yes Share by Municipality

Post-treatment outcome measuring support for family policy

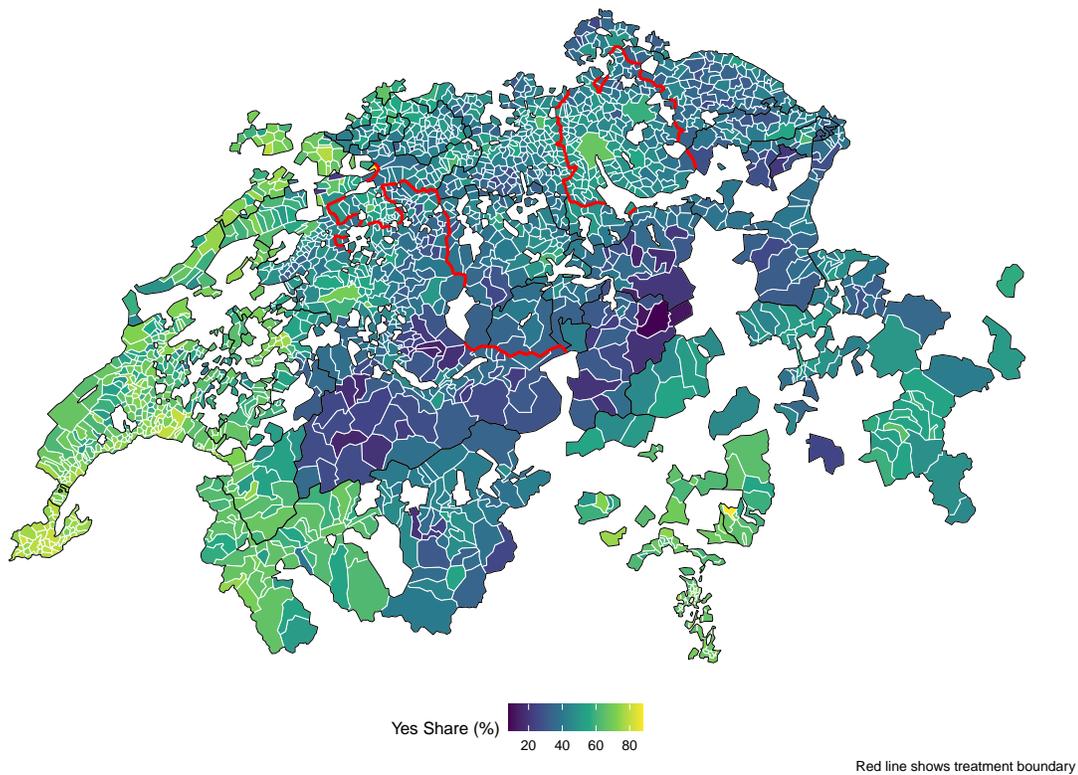


Figure 6: Family Policy Referendum Yes Share by Municipality

RDD Sample: German-Speaking Municipalities Near Border

N = 1095 municipalities within 30km of treatment boundary

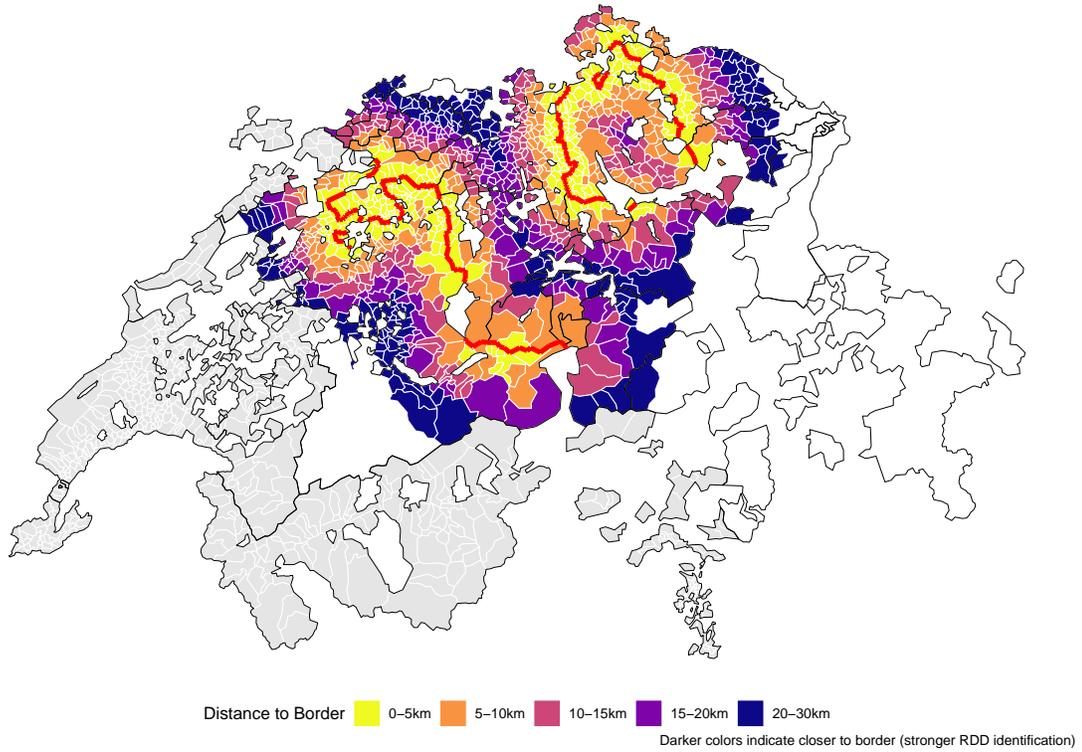


Figure 7: RDD Sample: Distance to Treatment Border