REPRESENTATION OVER RENTS: THE BROAD-COALITION NORM IN DUTCH LOCAL POLITICS

LOTTE SWANK, OTTO SWANK, AND JOES DE NATRIS

ABSTRACT. In multi-party parliamentary democracies, governments emerge from post-election bargaining. We study coalition formation in Dutch municipalities. Bargaining models and a norm describing that "the largest party should govern" predict that party ranks affect coalition formation. Using regression discontinuity designs based on close party rankings and narrow majorities, we find no causal effect of party rank on coalition inclusion. Instead, governments overwhelmingly follow a broad-coalition logic: when two parties jointly win a narrow majority, they nevertheless include a third party in more than 90% of cases. Parties that participate in coalitions excluding a large party are less likely to enter the next coalition, consistent with dynamic penalties for violating a broad-coalition norm. Our results suggest that in repeated coalition bargaining, governing power can depend on a self-enforcing broad-coalition norm.

JEL codes: D70, D72, D90

Keywords: Coalition formation, Bargaining models, Regression Discontinuity Design, Democratic norms.

Date: August 27, 2026.
Contact information: Lotte Swank (l.m.swank@rug.nl); Otto Swank (swank@ese.eur.nl); Joes de Natris (joes.de.natris@gmail.com). We thank Robert Dur, Dylan Jong, Sacha Kapoor, Dana Sisak, and the participants in the Webinar on Political Economics for their helpful comments. Declarations of interest: none.

1. INTRODUCTION

In multi-party parliamentary democracies, elections determine the distribution of seats but not directly the composition of government. Governments instead emerge from post-election bargaining. Standard bargaining models imply that coalition formation should reflect parties' bargaining power and should often produce governments close to minimal winning coalitions. Yet, coalition bargaining may also be shaped by norms prescribing who should govern. We test these competing views using data from Dutch municipal elections from 1998 to 2022.

Our analysis generates five main findings. First, we find no evidence that parties' ranks in the election outcome affect coalition participation. In bargaining models à la Baron and Ferejohn (1989), rank should matter because proposal power depends on rank. Second, we find little support for the minimal winning coalition hypothesis. Rather, coalitions are frequently oversized, particularly when they would hold a narrow majority without the surplus party. Third, when coalitions are not oversized, their combined vote shares are typically well above a bare majority. Fourth, surplus parties in oversized coalitions are often larger parties. Finally, we find suggestive evidence that coalition parties that leave a large party outside the coalition are less likely to participate in the next coalition. Thus, acting against the broad-coalition norm seems to have adverse consequences for future office holding. Taken together, these results are difficult to reconcile with canonical bargaining models of coalition formation. Instead, they are consistent with a norm favoring broad coalitions and the inclusion of larger parties.

Our empirical strategy uses regression discontinuity designs based on close election outcomes. We begin by testing whether rank matters for coalition inclusion. In close elections, the ordering of parties is effectively random, allowing us to isolate the causal effect of finishing first rather than second, or second rather than third. If coalition formation is shaped by proposal power as in standard bargaining models, these discontinuous changes in rank should affect coalition participation. We find no such effects. To our knowledge, this is the first paper to estimate the causal effect of party rank on coalition composition.

We then turn to the minimal winning coalition hypothesis. In many municipalities, the two largest parties both enter the governing coalition. This allows us to study whether an additional party joins when those two parties barely cross the majority

threshold. If parties seek to minimize the number of coalition partners, the probability that another party joins should fall sharply once the two largest parties secure a majority on their own. We find another pattern. When the two largest coalition parties hold only a narrow majority, an additional party is likely to join. The probability of an additional coalition partner declines gradually as the core coalition's majority increases.

The paper makes two contributions. The first is to the political economy of coalition formation. A large theoretical literature predicts that bargaining power and the size of feasible majorities should play central roles in government formation, beginning with Von Neumann and Morgenstern (1947) and Riker (1962). Our results reject the main predictions of these models in the Dutch municipal context. In particular, we find no causal effect of rank-based proposal power on coalition inclusion, and no evidence that parties minimize coalition size to maximize their distributive shares. The second contribution is to the literature on norms in politics. Recent work shows that political outcomes may reflect a norm that the largest party should govern (Fujiwara and Sanz, 2020).[1] We provide evidence for a different norm: governments should be broad and include large parties. Because this norm prescribes including large parties, the largest party frequently governs, but not because of its ordinal rank.

Our empirical strategy is not designed to identify the role of ideology in coalition formation. Balance tests indicate that ideology is unlikely to confound our estimates, but our results should not be interpreted as evidence that ideology is unimportant. At the same time, a norm favoring broad coalitions is likely to produce ideologically broader governments, with possible consequences for policymaking. Standard bargaining models view coalition formation primarily as a divide-the-pie problem in which parties seek to maximize their share of office rents. By contrast, a broad-coalition norm emphasizes government as a vehicle for broad societal representation. Through this channel, our findings relate to the political economy literature on the trade-off between social welfare and narrow self-interest. Nannicini et al. (2013) distinguish civic voters, who evaluate politicians on social welfare, from uncivic voters who evaluate politicians on personal interest. They find that in Italy, social capital

[1] Gravdahl (1998) is a very early paper on norms in municipal politics. He investigates parties' support for mayors in Norway. He concludes that Norwegian politics is driven by a norm for consensus. Unlike in Dutch local politics, in Norwegian local politics the election outcome, not bargaining, fully determines the executive.

explains when voters punish incumbent politicians for misbehavior. Kendall and Matsusaka (2025) use California ballot propositions to determine how voters trade off the common good and personal interest. They find that many voters attribute a high weight to the public good relative to personal interests. While these studies focus on voter behavior, our paper identifies a parallel dynamic among political parties. The persistent formation of oversized coalitions suggests that local parties prioritize broad representation over the maximization of rents.

This result also translates into the divergence between governments acting as delegates, who just serve their vote base, and as trustees, who promote the public interest. In bargaining models, players act as delegates. The election outcome gives parties the authority to promote their constituencies' interests. George W. Bush described himself as a delegate. The elections had "*given the people [...] a clear choice" to which they had "responded*" (p. 47 Azari, 2014). He described himself as having a mandate to implement his specific agenda. Bush's Dutch equivalent is Dilan Yeşilgöz, the national leader of the Dutch liberal party (VVD). Just after the formation of the *national* government coalition in 2026, she placed billboards along highways telling her constituents what she had bargained out for them. For her, coalition bargaining was about obtaining the largest share of the pie. Our empirical findings do not portray Dutch local politicians as delegates.

Bill Clinton is an example who described himself as a trustee: the country did not face "*Democratic or Republican challenges*", but common challenges that required putting "*country ahead of party*" (Clinton, 1992). Broad coalitions in Dutch municipalities also describe themselves as trustees. For example, the broad coalition in Capelle aan de IJsel in 2022, in which the liberal party (VVD, the same as in the previous example) also participated, announced: We are "*a government that listens to the concerns of the residents of Capelle and seeks perspective and solutions with them. A government that reduces, rather than widens, the divide in society, and that inspires confidence*."[2]

2. Institutional background: Dutch Municipalities

Municipalities are the lowest tier of government in the Netherlands. Through administrative mergers, the number of municipalities fell from 548 in 2002 to 342 in

[2]Original Dutch text: "*Een bestuur dat luistert naar de zorgen van de Capellenaren en met hen naar perspectief en oplossingen zoekt. Een bestuur dat de tweedeling in de samenleving niet groter, maar kleiner maakt, dat vertrouwen geeft*" (https://gouweijsselnieuws.nl/2022/07/08/capelse-gemeenteraad-benoemt-nieuw-college-met-nieuwerkerker-rik-van-woudenberg/).

2026. Their primary responsibilities include public infrastructure, social support and welfare, and local amenities. To finance their expenditures, municipalities receive grants from the central government and levy municipal taxes, such as property and waste-collection taxes. Dutch municipalities account for 29% of total public expenditure (OECD and UCLG, 2023).

In municipal elections, which occur every four years, citizens elect council members through a proportional representation system. Depending on the municipality's population, the number of council members ranges from 9 to 45 and is odd in all municipalities. The municipality's executive, the *college*, comprises the nationally appointed mayor (the chair) and several aldermen. The number of aldermen varies from 2 to 11, also depending on the municipality's population. After the election, the parties in the council bargain over the allocation of aldermen. We call this bargaining process 'coalition formation'. We define a party as a coalition member if it supplies at least one alderman.[3] One typical feature of Dutch coalition formation is the informal rule that the party that receives the most votes in the election holds the power to propose.[4]

Because Dutch law requires only that all aldermen retain the confidence of a council majority, it allows both majority and minority coalitions.[5] In practice, however, minority coalitions are rare. After the coalition has been formed, council members can propose a motion of no confidence in an alderman. If such a motion is accepted, the alderman has to resign. Formally, the council cannot propose a motion of no confidence in the entire coalition. Of course, it can propose motions of no confidence in each individual alderman. However, coalitions rarely fall.[6]

Dutch municipalities are highly fragmented. In some large ones, more than 20 parties participate in the elections.[7] In contrast to local governments in Spain (Carozzi

[3]In 167 elections (6.2% of total elections) we could not match all aldermen to parties. We exclude these observations. In the online appendix, we repeat all analysis without excluding these observations. None of the conclusions change.

[4]In the 2026 Parliamentary Elections, the seat allocation was known a few days after the election. However, it was not known whether D66 or the PVV received the most votes. Because of the norm that the largest party holds the power to propose, the coalition formation process began when the agency responsible for elections, De Kiesraad, could declare the official election outcome.

[5]Formally, coalitions do not exist. The law sets a maximum number of aldermen in the college and requires that each alderman has the confidence of a majority of council members. In practice, coalitions exist, as evidenced by coalition agreements signed by the parties providing aldermen.

[6]In the rare case that a coalition does fall, a new coalition formation starts. A coalition's fall never leads to early elections.

[7]For example, 26 parties participated in Amsterdam in 2022 (Kiesraad, 2022).

et al., 2022, 2024), this fragmentation has not led to political instability. However, it makes single-party majorities rare, such that coalitions are needed to form a majority.

The Dutch setting is well suited to testing the main predictions of the bargaining model against a broad-coalition norm for four reasons. First, coalition formation in Dutch municipalities closely resembles that in many parliamentary democracies. Municipal elections determine how council seats are allocated to parties, and coalitions result from bargaining among them. Unlike in most other countries where mayors are directly elected, mayors in Dutch municipalities are appointed by the national government. As a result, only parties in the council are involved in the bargaining process, mirroring the pure legislative bargaining assumed in many bargaining models of coalition formation. Second, councils where one party holds a majority of seats and minority coalitions are rare.[8] Third, existing empirical studies on government formation are predominantly cross-country studies [see, for example, Ansolabehere et al. (2005), Bassi (2017), Diermeier et al. (2003), Diermeier and Merlo (2004), and Schofield and Laver (1985)]. A drawback of cross-country studies is that government formation processes differ in almost every institutional detail. Our focus on Dutch municipalities ensures that the rules governing coalition formation are essentially the same across municipalities. Finally, our sample, covering seven municipal elections, is much larger than the samples used in cross-country studies.

3. DATA

Our sample comprises seven waves of municipal elections in the Netherlands. We use Ogink's dataset (Ogink, 2022), which includes parties' vote shares and seats for each municipality in each election. We define a party as a coalition member if it supplies at least one alderman. Data on aldermen and party affiliations are from documents by the Association of Dutch Municipalities (VNG, 1998-2002) for the 1998 and 2002 election years. For the other five election waves, we use data from the Ministry of Domestic Affairs website.[9] Our total sample comprises 2839 elections. However, in 167 elections, we could not match an alderman to a party. We dropped these elections, leaving us with 2672 elections.

[8] In some parliamentary democracies, such as Denmark, minority governments are not rare. Diermeier and Merlo (2000) offer a model in which minority governments can also emerge.
[9] Website

We use party-ideology data for balance tests and robustness checks. We measure positions on the economic and socio-cultural dimensions using Wordscores (Laver et al., 2003). We apply this method to manifestos and agreements following Otjes et al. (2023). As reference texts, we use national manifestos of political parties from 2012 to 2021.

4. Does Rank Affect Coalition Formation?

In this section, we use rank-based discontinuities to test whether coalition formation is consistent with (1) a norm that the "largest part should participate in the coalition" (Fujiwara and Sanz, 2020), and (2) with the main predictions of bargaining models à la Baron and Ferejohn (1989). Bargaining models frame coalition formation as a divide-the-pie game, in which all parties seek to extract the largest possible share. This literature shows that the bargaining outcome depends heavily on the rules of the game. Nevertheless, bargaining models generate two common predictions. First, the party with proposal power disproportionately influences the bargaining outcome. Second, ceteris paribus, the party with proposal power avoids partnering with strong alternative proposers [see, for example, Baron and Ferejohn (1989), Baron and Diermeier (2001), Battaglini (2021), and Diermeier et al. (2003)]. To see why, suppose that the largest party proposes first a division of the pie, followed by the runner-up in case of no success. The runner-up derives significant bargaining power from its position as the next-in-line proposer: it can demand a larger share of the pie by threatening to reject the current offer and make its own proposal. Consequently, the largest party can demand a larger share of the pie by forming a coalition with another party, such as the third-largest party. Because both predictions pertain to parties' ranks rather than their sizes, we can use RDDs to test them.

The first prediction concerns the two largest parties and their ranks in the election outcome: because if its *rank*, the largest party has a higher probability of participating in the coalition than the runner-up. By testing this first prediction of bargaining models, we also test the prevalence of a norm that the largest party should participate in the coalition. We define v_{nt} as the vote share of the largest party minus the vote share of the second-largest party in municipality n in election year t. For party i, we

define the signed running variable as

$$x_{int} = \begin{cases} v_{nt} & \text{if } i \text{ has the most votes} \\ -v_{nt} & \text{if } i \text{ has the second most votes.} \end{cases} \tag{1}$$

The outcome variable y_{int} equals 1 if party i supplies (at least one) alderman in municipality n in year t and 0 otherwise. To test the first prediction of the bargaining model, we estimate

$$y_{int} = \beta_0 + \theta \cdot 1\{x_{int} > 0\} + \beta_1 x_{int} + \beta_2 x_{int} \cdot 1\{x_{int} > 0\} + \epsilon_{int}, \tag{2}$$

where $1\{x_{int} > 0\}$ is an indicator function equal to 1 if party i is the largest party and equal to 0 if party i is the second-largest party. We are interested in θ: the effect of being the largest party (instead of the runner-up) on the probability of participating in the coalition at the threshold.

Following Fujiwara and Sanz (2020), we say that the data is consistent with "the largest party should govern norm" if $\theta > 0$. A strict version of the bargaining model predicts that $\theta = 1$: The largest party is always a coalition member, the runner-up never. In practice, however, proposal power need not be the only determinant of coalition formation. Ideological proximity or personal considerations may also matter. For this reason, if θ is significantly higher than zero, we already interpret this as evidence in favor of the first prediction of the bargaining model.

In our main specification, we estimate (2) with a linear specification. Hence, β_1 gives the slope to the left of the threshold, and $\beta_1 + \beta_2$ gives the slope to the right of the threshold. Additionally, we present quadratic specifications.

Figure A.1 in the Appendix shows the histogram of the running variable x_{int}. It does not exhibit a clear discontinuity around the threshold, as confirmed by a local polynomial density test [$p = 0.834$] (McCrary, 2008; Cattaneo et al., 2020). The running variable is symmetric around the threshold, because each municipal election is treated twice: once for the largest party and once for the runner-up. This artificially increases the number of observations. The main results of this paper arise from insignificant discontinuities. Doubling observations goes against insignificant discontinuities. In Table A.1 in the Appendix, we present similar analyses where one of the two parties is randomly removed, such that no duplicates remain. As expected, these analyses yield similar results.

Figure 1 plots our result graphically. A dot that is located left of the threshold gives the estimated probability that the runner-up party within a bin of the running variable participates in the coalition. Analogously, a dot to the right of the threshold represents the probability that, in a given bin of the running variable, the largest party participates in the coalition. Figure 2 does not reveal a big jump at the zero threshold, indicating that a very small positive winning margin does not result in a much higher probability of coalition participation.

FIGURE 1. The effect of being the largest party rather than the runner-up on coalition inclusion

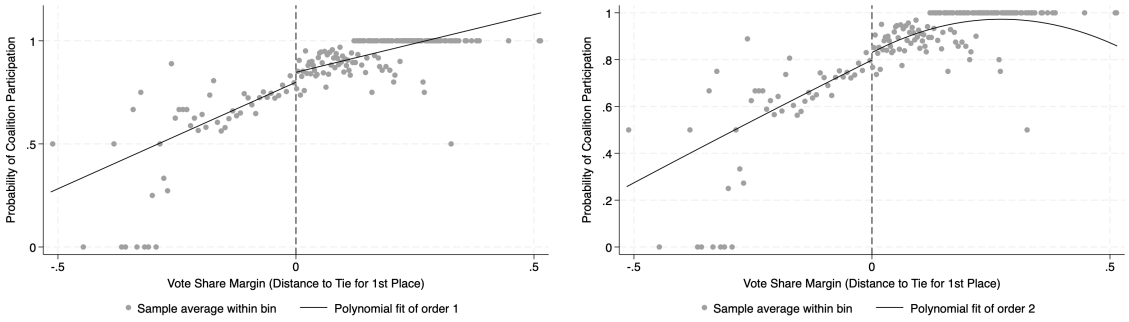

(A) Linear (B) Quadratic

TABLE 1. RDD estimates of being the largest party rather than the runner-up party on coalition inclusion.

	(1)	(2)	(3)	(4)	(5)
RDD Estimate	-0.011	-0.023	-0.007	-0.018	-0.010
	(0.031)	(0.057)	(0.029)	(0.034)	(0.031)
Robust p-value	0.719	0.677	0.804	0.589	0.755
Observations	5344	5344	5344	5344	5344
Bandwidth	0.074	0.037	0.148	0.129	0.082 (L) / 0.075 (R)
Specification	Linear	Linear	Linear	Quadratic	Linear (Asym. BW)

Note: Each column reports the coefficient from a local polynomial regression estimated using the `rdrobust` package. Robust bias-corrected standard errors adjusted for clusters (by municipality) are in parentheses. In columns (1) and (4), bandwidths are MSE-optimal and symmetric. The optimal bandwidth is multiplied by 0.5 in Column (2) and by 2 in Column (3). Column (5) utilizes independently estimated MSE-optimal bandwidths for the left (L) and right (R) sides of the cutoff (`msetwo`) to account for potential asymmetry in the functional form. ***, **, and * denote significance at the 1%, 5%, and 10% levels, respectively.

Table 1 displays the equivalent regression results. The first three columns denote the results of analysis with linear specifications. In Column (1), We use optimal

bandwidths based on Calonico et al. (2014). Column (2) halves these optimal bandwidths and (3) doubles them. Column (4) uses optimal bandwidths with a quadratic specification. The estimation in Column (5) is linear with asymmetric bandwidths. In all specifications, the RDD estimates are *negative*, but not significantly different from zero. Thus, the estimates provide no support for the hypothesis that the largest party's proposal power increases the probability that it participated in the coalition.

For Spain, estimating an equation similar to (2), Fujiwara and Sanz (2020) find that being the largest party rather than the runner-up increases the probability of appointing the mayor. They interpret this finding as evidence for a norm that the largest party should govern. Our results show that this norm does not exist in Dutch local politics. Later, we present results indicating the existence of a norm prescribing that coalitions should be broad, including larger parties.

Let us now test the prediction that the largest party is more likely to invite the third-largest party to join the coalition than the runner-up. This prediction assumes that the largest party is in the coalition. The cleanest test of this hypothesis concerns cases in which the largest party can form a majority coalition with the third-largest party. As a result, to test the second prediction, the sample excludes observations in which the largest party does not participate in the coalition and observations in which the largest and third-largest parties do not together have a majority of seats in the council. The restricted sample includes 1594 observations.

To test the second prediction, we redefine v_{nt} as the vote share of the second-largest party minus the votes share of the third-largest party, and redefine the running variable in (2) as: $x_{int} = v_{nt}$ if i has the second most votes, and $x_{int} = -v_{nt}$ if i has the third most votes. Figure A.2 in the Appendix shows the histogram of this running variable. A local polynomial density test (McCrary, 2008; Cattaneo et al., 2020) does not indicate a discontinuity around the threshold, [$p = 0.824$].

Figure 2 plots the probabilities that the runner-up and the third-largest party participate in the next coalition against the running variable. There is a small jump at the threshold, suggesting that being the runner-up rather than the third-largest party *increases* the probability of inclusion in a coalition. Table 2 presents the RDD estimates, which do not differ sigificantly from zero in all specifications. Thus, in contrast to the bargaining model, *being* the third-largest party rather than the runner-up party does not increase the likelihood of coalition participation.

FIGURE 2. The effect of being the runner-up rather than the third-largest party on coalition inclusion

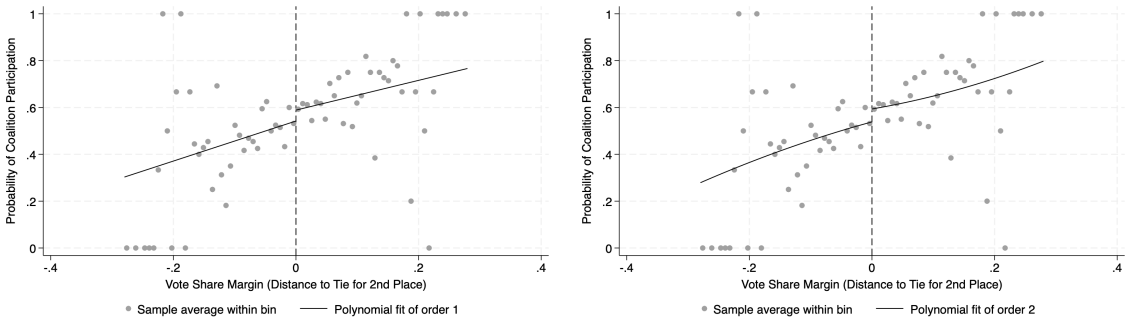

(A) Linear (B) Quadratic

Note: The sample only includes observations where the first party participates in the coalition and the first and third parties have a majority of votes.

TABLE 2. RDD estimates of being the runner-up rather than the third-largest party on coalition inclusion

	(1)	(2)	(3)	(4)	(5)
RDD Estimate	0.065	-0.073	0.060	0.011	0.058
	(0.081)	(0.057)	(0.073)	(0.102)	(0.080)
Robust p-value	0.421	0.588	0.411	0.912	0.469
Observations	1594	1594	1594	1594	1594
Bandwidth	0.056	0.028	0.112	0.066	0.052 (L) / 0.066 (R)
Specification	Linear	Linear	Linear	Quadratic	Linear (Asym. BW)

Note: Each column reports the coefficient from a local linear regression estimated using the `rdrobust` package. Robust bias-corrected standard errors adjusted for clusters (by municipality) are in parentheses. In columns (1) and (4), bandwidths are optimal. The optimal bandwidth is multiplied by 0.5 in Column (2) and by 2 in Column (3). Column (5) utilizes independently estimated MSE-optimal bandwidths for the left (L) and right (R) sides of the cutoff (`msetwo`) to account for potential asymmetry in the functional form. ***, **, and * denote significance at the 1%, 5%, and 10% levels, respectively.

In this section, we have tested "the largest party should govern norm" and two predictions of a typical bargaining model of coalition formation. We find no support for them. In Dutch local politics, parties of similar size have equal chances of participating in the coalition.

4.1. **Ideology.** Another likely candidate affecting coalition formation is ideology. More than half a century ago, Axelrod (1970) already argued that coalitions are formed

by ideologically close parties.[10] Because parties' possible desire to form a coalition with ideologically similar parties does not depend on their ranks or sizes, it is difficult to estimate the causal effect of ideology on the probability that parties form a coalition. However, our objective in this section is more modest. We want to rule out that ideological considerations are responsible for the results presented in the previous section.

It is possible that we do not observe the largest party to form a coalition with the third-largest party rather than the runner-up, because, from an ideological point of view, the runner-up is more appealing to the largest party than the third-largest party. Note that, in principle, the RDD should take care of this. Parties close to the threshold should be similar except for rank. As a robustness check, we run an RDD that stacks the deck *against* the runner-up party.[11] To this end, we include only the observations in the sample where the first and third parties are ideologically closer than the first and second parties.[12] Then, the largest party should prefer a coalition with the third-largest party rather than with the runner-up, from both ideological and bargaining perspectives.

Figure A.3 in the Appendix shows that even in this stacking against the desk case, we do not observe any evidence for the second prediction of the bargaining model. Table A.2 in the Appendix presents the RDD estimates. They confirm the RDD plot. If anything, the runner-up has a higher probability of inclusion in a coalition than the third-largest party, despite an ideological disadvantage.

4.2. **Council size.** To ensure that the results are not driven by municipalities with specific sizes or electoral formulas, we conduct a robustness check by dividing the observations into two subsamples: small municipalities with fewer than 19 council seats and large municipalities with 19 or more seats. This threshold accounts for the shift in the Dutch electoral system regarding the allocation of remainder seats.

[10] Austen-Smith and Banks (1990), Baron (1991), and Morelli (1999) offer models in which both bargaining and ideology affect coalition formation.

[11] We measure ideology using word scores derived from local party programs, anchored by national party manifestos. We employ three alternative specifications. The first captures the economic left-right dimension. The second captures the socio-cultural dimension (GAL-TAN). Because these two dimensions are highly correlated ($\rho = 0.86$), our third specification utilizes the first principal component of these two variables, extracted via Principal Component Analysis (PCA). This primary component explains 93.02% of the total variance and weights both underlying dimensions equally (loading: 0.707 for both the economic and the socio-cultural dimensions).

[12] The ideology data is only available for the last three election waves. This, combined with imposing all restrictions, reduces the sample size substantially.

While councils with 19 or more seats distribute the remainder seats using the highest-average D'Hondt method, smaller councils employ the largest remainder method. Table A.3 in the Appendix shows that estimates are similar for both large and small municipalities.

5. Minimal winning or broad coalitions

As discussed in the introduction, the minimal winning coalition hypothesis posits that political parties form majorities without including any superfluous members (Von Neumann and Morgenstern, 1947). Because adding unnecessary parties reduces each member's share of the pie (i.e., allocated aldermen, policy influence, visibility), coalition sizes are kept strictly to a necessary majority. Riker (1962) argued one step further: parties form a *minimum* coalition,[13] defined as the coalition that secures a majority with the narrowest possible margin of total seats. In such a minimum coalition, the coalition parties secure the largest share of the pie.

This section shows that neither hypothesis holds in Dutch local politics. Instead, we argue that a norm prescribing broad coalitions that include large parties shapes coalition formation in Dutch municipalities. We refer to this as the "broad-coalition norm." By "broad," we mean that the coalition parties' total vote share is substantially above one-half. Almost by definition, the broad-coalition norm contrasts with the minimal winning coalition hypothesis. We test the broad-coalition hypothesis against the minimal winning hypothesis using RDDs.[14] We present causal evidence against the minimal winning coalition hypothesis and in favor of a broad-coalition norm. Additionally, we present correlational evidence that this norm dictates the inclusion of larger parties. This suggests that Dutch local parties strive to form coalitions that represent a large part of the electorate. Before turning to the formal analysis, we present descriptive evidence to motivate it.

5.1. **Oversized coalitions are common.**

A striking feature of Dutch local politics is that oversized coalitions — coalitions containing more parties than mathematically necessary to obtain a majority — are very common. Figure 3 shows that the share of oversized coalitions is fairly stable over time, at around 40%. The substantial number

[13] Note that the minimum coalition is the smallest of all minimal coalitions (Mueller, 2003).

[14] In our RDDs, we test around the 50%-majority threshold. Around this threshold, almost all minimal coalitions are, by construction, minimum coalitions. Yet, we will refer to the minimal winning hypothesis to remain precise, as we did not definitively rule out the existence of mathematically smaller alternative compositions for every observation near the threshold.

of oversized coalitions is inconsistent with the minimum (Von Neumann and Morgenstern, 1947) and minimal coalition hypotheses (Riker, 1962).

While surplus majorities are common, the data reveal a strong aversion to minority governments. Only in less than 5% of cases do the parties supplying aldermen control less than a majority of council seats. Furthermore, because some parties formally sign the coalition agreement without supplying aldermen,[15] the prevalence of minority coalitions is even lower when defined by total legislative support rather than by parties supplying aldermen.

FIGURE 3. Share Minority and Oversized Coalitions in Election Years

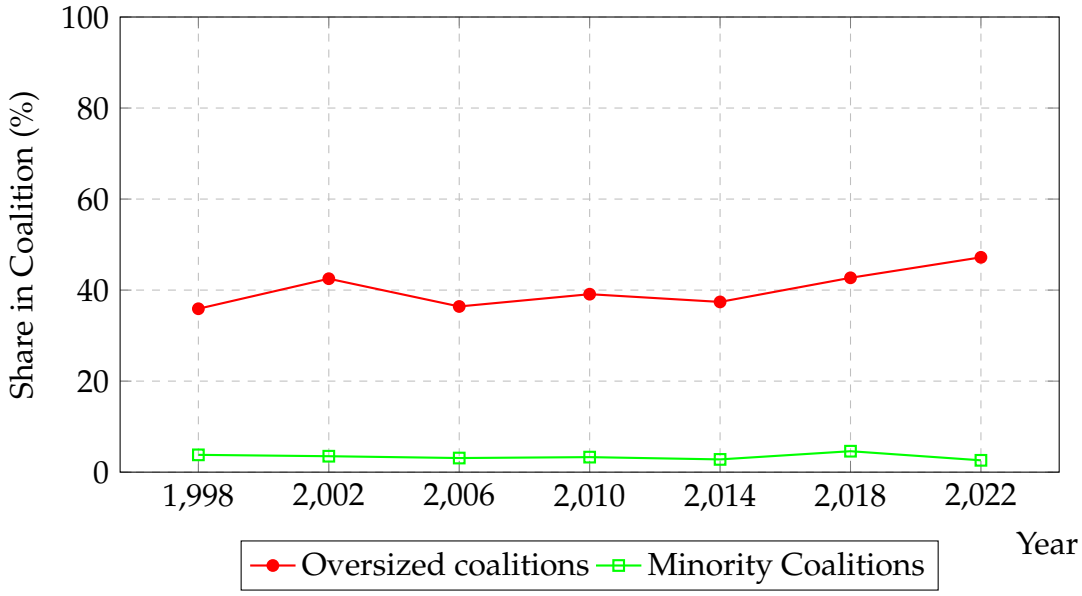

Note: The sample excludes 71 coalitions that *must* be minority coalitions. This occurs, for example, in a municipality with two aldermen positions where the largest and second-largest party have a combined vote share below 50%.

While notable, a 40% aggregate share of oversized coalitions does not directly prove the existence of a broad-coalition norm, as the majority of coalitions remain non-oversized. In Section 5.3, we show that oversized coalitions have been predominantly formed when, without the surplus party, the remaining coalition parties hold a *narrow* majority. Thus, the surplus party lifts the coalition from a narrow to a broad majority.

5.2. **Large parties are often in coalitions.** As demonstrated in Figure 4, larger parties are more common coalition members than smaller ones. In particular, the share

[15]This occurs, for example, if the number of coalition parties is higher than the number of aldermen that the municipality is by law allowed to install.

of the largest party in the coalition is remarkably high (almost 0.9). This naturally raises the question: which largest parties *do not* participate, and are they distinctive? Table A.4 shows that largest parties that remain outside are disproportionately local parties without recent coalition experience, relative to those that join. Table A.5 documents a similar pattern for second parties that are not included. Taken together, these patterns suggest that coalitions generally include large parties, unless they are local and inexperienced.

FIGURE 4. Share of x^{th}-largest parties that participate in coalition

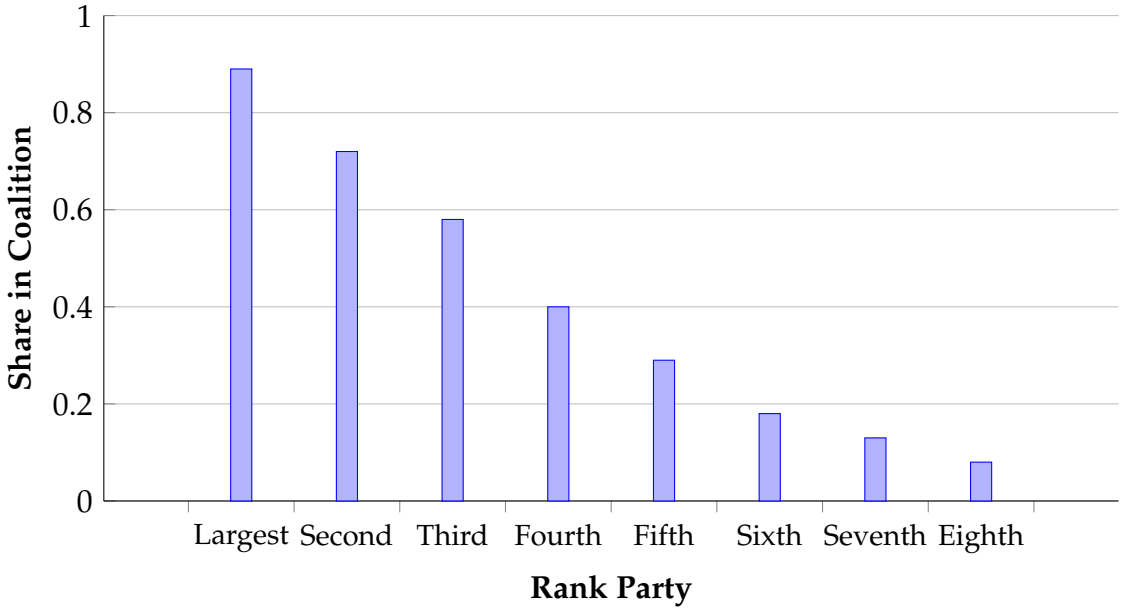

5.3. **A broad-coalition norm versus the minimal winning coalitions hypothesis.** Our descriptive analysis indicates that the largest party and the runner-up often both participate in the coalition. In this and the next section, we restrict our sample to the coalitions in which these two parties jointly form a 'core coalition'. We analyze whether this core block incorporates any additional parties. We exclude municipalities with two alderman positions, as the core cannot expand there. The sample includes 1410 coalitions (53% of total coalitions). Under the minimal winning coalition hypothesis, the core coalition includes another party if it lacks a majority in the council, but does not expand if it holds a (narrow) majority. Under the broad-coalition norm, in contrast, the core coalition also expands when it holds a *narrow* majority.

These hypotheses yield sharp, distinct predictions at the majority threshold, making RDDs well-suited to test them. To test whether the core coalition incorporates a supplementary party, we estimate

$$y_{nt} = \gamma_0 + \tau \cdot 1\{x_{nt} > 0\} + \gamma_1 x_{nt} + \gamma_2 x_{nt} \cdot 1\{x_{nt} > 0\} + \epsilon_{nt}, \tag{3}$$

where the binary outcome variable, y_{nt}, equals 1 if any supplementary party (with rank higher than two) supplies at least one alderman to the core coalition in municipality n during election year t, and 0 otherwise. The running variable, x_{nt}, takes one of two alternative forms: the core coalition's margin to a seat majority (x_{nt}^{seats}) or its margin to a vote share majority (x_{nt}^{votes}). The parameter of interest is τ, which captures the discontinuous effect of securing a legislative majority.

The first running variable, x_{nt}^{seats}, is defined as the combined number of council seats held by parties 1 and 2 minus the number of seats required for a majority. If $x_{nt}^{\text{seats}} < 0$, the core coalition requires an additional party to secure a majority. If $x_{nt}^{\text{seats}} > 0$, an additional party is not necessary for a majority.[16] The drawback of this variable is that council sizes depend on the municipal population, implying that observations near the threshold are more likely to come from smaller municipalities. However, the local polynomial density test does not indicate a discontinuity for this running variable [$p = 0.246$], see also Figure A.4 in the Appendix. Furthermore, the robust clustered standard errors at the municipality level account for the mass points in the running variable.

The second running variable, x_{nt}^{votes}, is defined as the combined vote share of parties 1 and 2 minus 0.5.[17] A drawback of this measure is that vote share majorities do not perfectly align with seat share majorities because of rounding in seat allocation. In 120 observations, the core coalition secures a seat majority without an absolute vote majority, while in 8 observations, the reverse occurs.[18] Consequently, while the

[16] Because the number of seats is odd in all municipalities, a perfect tie ($x_{nt} = 0$) is institutionally impossible.

[17] A local polynomial density test does not indicate a discontinuity for this running variable [$p = 0.824$], see also Figure A.4 in the Appendix.

[18] The reason that the number of observations where $x_{nt}^{seats} > 0$ and $x_{nt}^{votes} < 0$ is much higher than the number of observations where $x_{nt}^{seats} < 0$ and $x_{nt}^{votes} > 0$ is twofold. First, votes cast for parties that fail to meet the electoral threshold do not translate into seats. The weight of these "wasted" votes is redistributed among the parties that did win seats. The largest parties, the core, receive the highest proportional benefit. Second, in municipalities with 19 or more seats, the remaining seats are distributed using the D'Hondt method. This highest-average formula biases the allocation of unassigned seats towards larger parties. Consequently, the two largest parties, the core, frequently capture a disproportionate share of the remaining seats.

continuous nature of the vote margin makes it statistically appealing, estimating a sharp RDD at the 50% vote threshold tests an imperfect proxy.

A fundamental assumption of the RDD is local randomization around the threshold. In our analysis, this translates to the risk of endogenous selection bias. This problem arises when achieving a joint majority affects the probability that the core decides to participate in the coalition together. Although the local polynomial density tests suggest this is not the case, we conduct two additional robustness checks. First, we test for selection bias by estimating the probability that the core participates in the coalition as the dependent variable in an RDD specification. Second, we verify via balance tests whether pre-treatment covariates (the number of parties, the vote share of the largest parties, local/national and incumbency status, and dummies for specific parties that capture ideology) are continuously distributed around the threshold. The results of these tests, presented in tables A.6, A.7, and A.8 in the Appendix, show no statistically significant discontinuities.

Ultimately, our RDD estimations arbitrate between the two theoretical frameworks. The minimal winning coalition hypothesis predicts a sharp, discontinuous drop in the probability of coalition expansion precisely at the majority threshold, and hypothesizes $\tau = -1$. Alternatively, if a broad-coalition norm drives coalition formation, the probability of expansion will remain continuous at the threshold, indicating that $\tau = 0$.

Figure 5 presents the paper's main result. Figure 5a plots the probability that the core coalition expands against the core coalition's surplus of seats. The horizontal line to the left of the threshold indicates the rarity of minority coalitions: all core coalitions that lack a majority expand. Surprisingly, there is no significant discontinuous drop at the threshold, as confirmed by Table 3. To the right of the threshold, the line slopes downward: the larger the majority of the core coalition, the lower the probability that an additional party will join. We interpret Figure 5 as suggestive evidence of a "broad coalition" norm: coalition parties' combined share of council seats should be substantially above one-half. Just right at the threshold, a broad coalition requires an additional coalition party. Further to the right, the core coalition *is* already broad, so it does not need an additional coalition party.

Figure 5b plots the same outcome variable as Figure 5a but against the core coalition vote share surplus. The line to the left of the threshold decreases a bit , as some

core coalitions have vote shares below one-half but hold a majority of seats. The insignificant (see Table 3) drop at the threshold shows that when the core coalition just holds a majority, an additional party joins the coalition more than 90% of the time. Again, right to the threshold, a higher vote share is associated with a lower probability of coalition expansion.

FIGURE 5. The probability that the core coalition expands

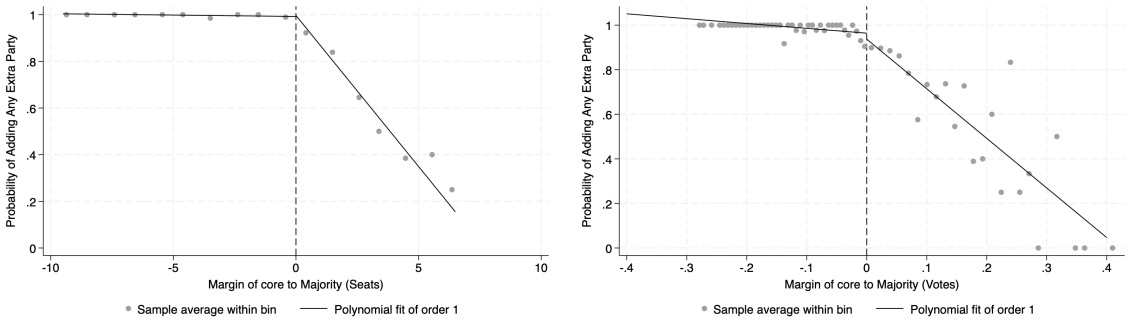

(A) Running variable is seat share core coalition half the council size

(B) Running variable is vote share core coalition minus one half

Note: The sample includes only observations in which the largest and second-largest parties participate in the coalition and the council has more than two aldermen.

TABLE 3. RDD estimates of majority status on the probability of expanding the core coalition

	(1)	(2)	(3)	(4)
RDD Estimate	-0.019	-0.031	0.017	0.022
	(0.039)	(0.041)	(0.055)	(0.055)
Robust *p*-value	0.624	0.645	0.755	0.686
Running Variable	Seats	Seats	Votes	Votes
Bandwidth Type	Symmetric	Asymmetric	Symmetric	Asymmetric
Bandwidth Values	4.183	4.682 (L) / 3.665 (R)	0.061	0.056 (L) / 0.061 (R)
Observations	1410	1410	1410	1410

Note: The sample includes only observations in which the largest party and the runner-up both participate in the coalition and the council has more than two aldermen. Columns (1) and (3) use the baseline symmetric MSE-optimal bandwidth selection (`mserd`). Columns (2) and (4) implement independent MSE-optimal bandwidth selectors for the left (L) and right (R) sides of the cutoff (`msetwo`) to accommodate potential asymmetries in the density or functional form of the data. All coefficients are estimated from a local linear regression with a triangular kernel and a polynomial of order 1 using the `rdrobust` package. Robust bias-corrected standard errors adjusted for clusters (by municipality) are in parentheses. ***, **, and * denote significance at the 1%, 5%, and 10% levels, respectively.

Panel (B) in Figure 6 shows that the core often forms a coalition with a seat share well above the majority threshold. Coalitions with seat shares of 0.56 to 0.71 are most

common and often consist of three or four parties. Panel (A) displays the number of parties in the coalition by the seat share of the core. It shows that the core often includes additional parties even when it already has a (large) majority. Finally, Figure 6 is consistent with our result that core coalitions with a narrow majority almost always expand. The bar just to the right of the threshold in Panel A shows that it is quite common for the two largest parties together to hold a majority of seats. The bar just to the right of the threshold in Panel A shows that those coalitions rarely form a coalition together without another party.

FIGURE 6. Histogram of the seat share of the core coalition and the full coalition

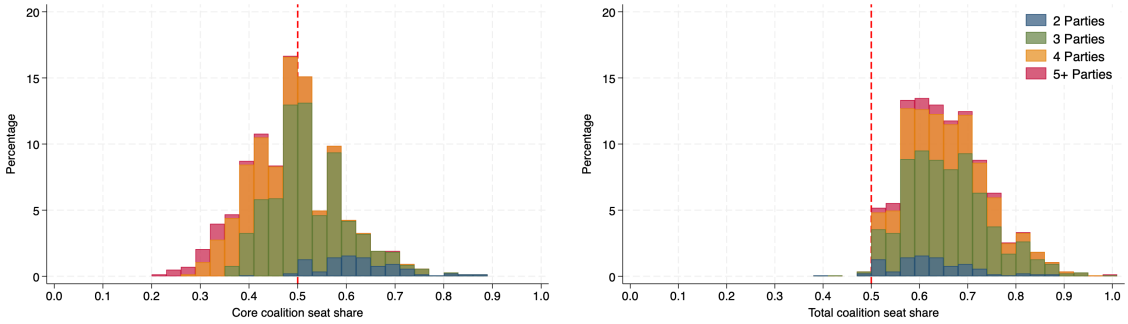

(A) Core (B) Total coalition

5.4. **Large party representation.** In the previous subsection, we found evidence consistent with a norm prescribing broad coalitions with large representation. We will now argue that this broad-coalition norm specifically prescribes the inclusion of large parties that represent a large share of the electorate. Figure 5 provides initial suggestive evidence. To the right of the majority threshold, the probability of incorporating an additional party declines as the size of the core coalition increases. This trend is mechanically linked to the size of the remaining available parties: as the core expands, the total remaining vote and seat shares in the council *must* shrink. As a specific example, among core coalitions holding a majority, a one-percentage-point increase in the core's vote share corresponds to a 0.194 percentage point decrease in the expected vote share of the third-largest party.[19] Similarly, a one-seat increase for the majority core decreases the expected seats of the third party by 0.226.[20] This

[19]Estimated via an OLS regression with municipality fixed-effects and standard errors clustered at the municipality level ($SE = 0.071$, $p < 0.01$).
[20]Estimated via an OLS regression with municipality fixed-effects and standard errors clustered at the municipality level ($SE = 0.060$, $p < 0.001$).

mechanical relationship suggests an underlying mechanism: core coalitions are less likely to include supplementary parties when the available partners are small. To test this more formally, the remainder of this subsection proceeds in two steps. First, we establish that the core coalition incorporates the third-largest party more often than the fourth-largest party. Second, we demonstrate that the probability of incorporating the third-largest party strictly increases with its size.

5.4.1. *Inclusion of the third- versus the fourth-largest party.* To compare the probability that the core coalition includes the third-largest party with the probability that it includes the fourth-largest party, we first estimate Equation (3), where the dependent variable y_{nt} equals 1 if the third-largest party is included in municipality n in election year t and 0 otherwise. Next, we repeat the analysis for the fourth-largest parties. In both RDDs, we use the core coalition's vote share minus one-half as the running variable. Figures 7a and 7b present the RDD plots. Both plots confirm the main results of the previous section. First, there are no big drops at the majority threshold.[21] Second, even if the core coalition holds a (narrow) majority, a considerable share of third- or fourth-largest parties join the coalition. Third, the probability that a party joins the core coalition decreases as the core coalition's vote share increases.

Most importantly, the line right to the threshold in Figure 7a lies fully above the corresponding line in Figure 7b. This suggests that the third-largest party is more likely to join a coalition than the fourth-largest party. To test this more formally, we define $P3_{int}$ as an indicator function equal to 1 if party i in municipality n in election year t is the third-largest party and equal to 0 if it is the fourth-largest party. We estimate the following pooled probability model within a regression discontinuity framework

$$\begin{aligned} y_{int} &= \delta_0 + \delta_1 P3_{int} + \delta_2 \cdot 1\{x_{nt} \geq 0\} + \delta_3 \left(P3_{int} \cdot 1\{x_{nt} \geq 0\}\right) \\ &\quad + \delta_4 x_{nt} + \delta_5 \left(P3_{int} \cdot x_{nt}\right) + \delta_6 (x_{nt} \cdot 1\{x_{nt} \geq 0\}) \\ &\quad + \delta_7 \left(P3_{int} \cdot x_{nt} \cdot 1\{x_{nt} \geq 0\}\right) + \epsilon_{int}, \end{aligned} \tag{4}$$

[21] For the third-largest party, the jump is estimated at 0.135 ($SE = 0.097$) and not significantly different from 0 (robust $p = 0.163$) with the core's vote share margin as the running variable. With the core's seat margin as the running variable, the drop is estimated at -0.013 ($SE = 0.073$, robust $p = 0.856$) for the third party. For the fourth-largest party, the drop is estimated at -0.025 ($SE = 0.095$) and is not significantly different from 0 (robust $p = 0.790$) when the core's vote share margin is used as running variable. When the core's seat share margin is used as the running variable, a jump of 0.050 is estimated for the fourth party ($SE = 0.065$, robust $p = 0.443$)

FIGURE 7

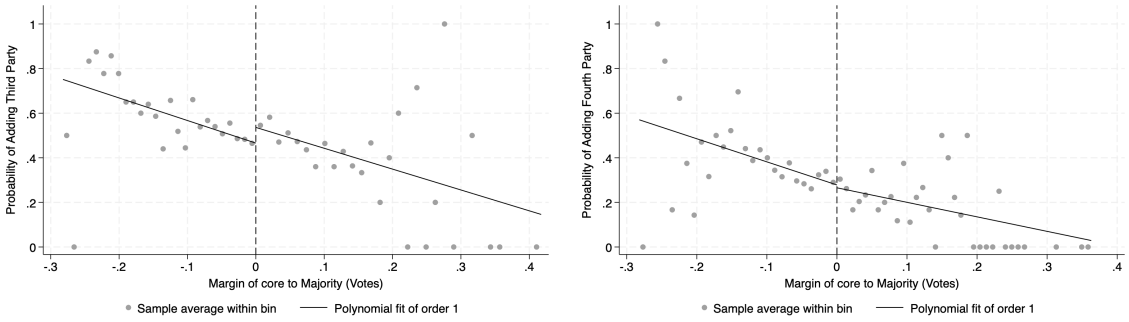

(A) Party 3 (B) Party 4

Note: The running margin is whether the largest and second-largest parties hold a majority. The sample includes only observations in which the largest and second-largest parties participate in the coalition and the council contains more than two aldermen. In Panel A, the outcome variable describes whether Party 3 joins the coalition. In Panel B, the outcome variable describes whether Party 4 joins the coalition.

where y_{int} indicates whether party $i = \{3, 4\}$ is added to the core coalition and x_{nt} is the core's margin to a majority (in seats or in votes). The coefficient δ_0 gives the fourth party's estimated probability of inclusion on the left bandwidth border, and $\delta_0 + \delta_1$ estimates this probability for the third-largest party. The slope of the function on the left of the threshold is given by δ_2 for the fourth party and by $\delta_2 + \delta_3$ for the third. The drops at the majority threshold are given by δ_4 and $\delta_4 + \delta_5$ for the fourth and third parties, respectively. On the right of the threshold, the slope for the fourth-largest party is $\delta_2 + \delta_6$ and for the third party $\delta_2 + \delta_3 + \delta_6 + \delta_7$. We restrict the sample to coalitions in which the largest two parties participate, use robust standard errors clustered on the country level, and use different bandwidths.

Table 4 confirms the intuition from the RDD plots. Across all specifications, the baseline coefficient for $P3$ is positive, indicating a preference for the third rather than the fourth-largest party. To the left of the majority threshold, the core parties are roughly 16 to 23 percentage points more likely to include the third party than the fourth. Moreover, achieving a majority does not alter this dynamic discontinuously. The coefficients for the threshold indicator (Majority) and its interaction (P3 $\times$ Majority) are statistically indistinguishable from zero. Hence, the preference for the third party persists independently of whether the core has a majority. Finally, in line with the previous results, if the size of the core coalition increases, the probability that any party is added declines (indicated by the negative coefficient on Margin).

TABLE 4. Effect of being the third versus fourth-largest party on being included by the core.

	(1)	(2)	(3)	(4)
P3	0.209***	0.225***	0.163***	0.186***
	(0.064)	(0.052)	(0.061)	(0.045)
Majority	0.048	0.044	-0.041	-0.016
	(0.050)	(0.042)	(0.055)	(0.042)
P3 × Majority	-0.026	-0.019	0.153	0.087
	(0.092)	(0.076)	(0.097)	(0.071)
Margin	-0.045*	-0.048**	-0.379	-1.081***
	(0.025)	(0.014)	(0.666)	(0.279)
P3 × Margin	0.014	0.026	-0.969	-0.006
	(0.041)	(0.026)	(1.123)	(0.432)
Majority × Margin	-0.023	-0.072	-0.169	0.539
	(0.035)	(0.021)	(1.099)	(0.446)
P3 × Margin × Majority	-0.384	-0.013	-0.217	-0.319
	(0.035)	(0.036)	(1.919)	(0.760)
Constant	0.286***	0.282***	0.297***	0.275***
	(0.035)	(0.029)	(0.036)	(0.027)
Running variable	Seats	Seats	Votes	Votes
Bandwidth	3	5	0.10	0.25
Observations	2236	3603	1904	2787
R^2	0.056	0.065	0.062	0.070

Note: Sample only includes observations in which the largest and second-largest party participates in the coalition and the council contains more than two aldermen. Majority equals 1 if the core coalition has a majority and 0 otherwise. Standard errors adjusted for clusters (by municipality) are in parentheses. ***, **, and * denote significance at the 1%, 5%, and 10% levels, respectively.

At first glance, the core coalition's preference for the third-largest over the fourth-largest party appears to contradict our conclusion in Section 3 that rank does not influence coalition inclusion. However, in Section 3, we isolated the effect of rank by comparing parties with nearly identical vote shares. That analysis demonstrated that ranks do not determine coalition inclusion. In the current analysis, we evaluate *all* third and fourth parties, where the third party may be substantially larger than the

fourth. Therefore, the third party's higher inclusion rate can stem from its larger size rather than its rank.

5.4.2. *The third-largest party is more likely to be included if it is large.* To estimate the effect of the third-largest party's size on its probability of joining the core coalition, we specify a pooled linear probability model within a Regression Discontinuity framework. We estimate this model using different local bandwidths around the core's vote share majority threshold. The baseline specification is

$$y_{nt} = \zeta_0 + \zeta_1 z_{nt} + \zeta_2 \cdot 1\{x_{nt} \geq 0\} + \zeta_3 x_{nt} + \zeta_4 x_{nt} \cdot 1\{x_{nt} \geq 0\} + \epsilon_{nt}, \tag{5}$$

where the dependent variable, y_{nt}, is an indicator equal to 1 if the third-largest party is included in the core coalition and 0 otherwise. We define the running variable, x_{nt}, as the combined vote share of the two largest parties minus 0.5. The indicator function $1\{x_{nt} \geq 0\}$ equals 1 when the two largest parties hold a combined majority. To isolate the effect of the third party's size, we introduce z_{nt}. Depending on the specification, we measure z_{nt} either as a continuous variable representing the third party's total vote share or as a dummy variable equal to 1 if this vote share exceeds 0.16.[22]

Table 5 presents the results. In all specifications, the core is more likely to include the third-largest party when it is large. This suggests that the broad-coalition norm specifically promotes the inclusion of large parties. Moreover, we again reject the minimum-coalition hypothesis: if anything, the third party is *more* likely to be included when the vote shares of the largest and second-largest parties just cross the majority threshold.

5.5. **Repercussions for norm deviation.** Bargaining models of coalition formation suggest that a broad-coalition norm is not always in a party's short-run interest. In the short-run, a party may have incentives to cash in on its electoral gains. Standard repeated games show that if such short-run incentives exist, a norm can, under conditions, be sustained by punishing parties that deviate from it. In this section, we explore if parties that deviate from the broad-coalition norm are punished.

The broad-coalition norm consists of two parts. First, a coalition should be broad, and second, big parties should be included in the coalition. Ideally, we would test the

[22]This is approximately the median vote share of the third-largest party in the relevant sample.

TABLE 5. Effects of the size of the third party and the majority margin of the core on the probability that the third party is included in the coalition.

	(1)	(2)	(3)	(4)
Third large	0.134***	0.108***	2.334***	2.275***
	(0.044)	(0.035)	(0.722)	(0.546)
Majority	0.132	0.112*	0.135	0.111*
	(0.085)	(0.059)	(0.085)	(0.070)
Margin	-3.406*	-1.552**	-3.573*	-1.748**
	(1.961)	(0.703)	(1.964)	(0.698)
Majority × margin	2.544	-0.297	2.384	-0.218
	(3.016)	(1.195)	(3.017)	(1.187)
Constant	0.343***	0.395***	0.027	0.072
	(0.061)	(0.045)	(0.131)	(0.100)
Third largest	Dummy	Dummy	Continuous	Continuous
Bandwidth	0.05	0.10	0.05	0.10
Observations	536	1052	536	952
R^2	0.025	0.020	0.027	0.027

Note: Sample only includes observations in which the largest and second-largest parties participate in the coalition and where the council contains more than two aldermen. If *Margin* is a dummy variable, it equals 1 if the vote share of the third party is larger than 0.16 and 0 otherwise. Standard errors adjusted for clusters (by municipality) are in parentheses. ***, **, and * denote significance at the 1%, 5%, and 10% levels, respectively.

first part of the broad-coalition norm with a fuzzy RDD where we exploit the core's margin to a majority as an instrument on whether or not it expands. However, as also visualized in Figure 5, the discontinuous effect at the majority threshold is weak. This lack of noncompliance aligns with a *strong* broad-coalition norm. If the electoral costs of noncompliance are high, parties simply do not violate the broad-coalition norm. We find similar patterns of strong compliance with the broad-coalition norm when investigating core coalitions that contain the largest and third-largest party but exclude the second-largest party, see Figure A.5 in the Appendix. Hence, we cannot investigate whether coalition parties are punished for forming a narrow coalition.

Compliance to the second part of the norm, coalitions should include large parties, is strong but not perfect. As shown in Figure 4, approximately 89% of the largest

parties and 72% of the second-largest parties participate in the coalition. In the remainder of this subsection, we investigate whether coalition parties face electoral repercussions when they exclude large parties. Specifically, we test whether coalition parties are less likely to participate in the subsequent coalition when they exclude a party with a vote share above 20%.

To this end, we estimate

$$\begin{aligned} Participation_{int+1} &= \alpha_0 + \alpha_1 \cdot Participation_{int} + \alpha_2 \cdot LargeOut_{nt} \\ &+ \alpha_3 \cdot Participation_{int} \cdot LargeOut_{nt} + \epsilon_{int}, \end{aligned} \tag{6}$$

where $Participation_{int}$ (and $Participation_{int+1}$) equals 1 if party i in municipality n participates in the coalition in election year t (and $t+1$) and 0 otherwise. Moreover, $LargeOut_{nt}$ is equal to 1 if there is a party with a vote share above 20% in municipality n in election year t that does not participate in the coalition. The parameter of interest is α_3: if it is negative, coalition parties are punished for excluding a large party.

TABLE 6. Effects of coalition participation and large party exclusion on future coalition participation across different model specifications.

	(1)	(2)	(3)	(4)
Participation	0.384***	0.225***	0.395***	0.392***
	(0.014)	(0.018)	(0.012)	(0.013)
Large Out	0.148***	0.100***	0.117***	0.153***
	(0.015)	(0.017)	(0.020)	(0.021)
Participation × Large Out	-0.234***	-0.178**	-0.257***	-0.294***
	(0.027)	(0.029)	(0.034)	(0.034)
Constant	0.302***		0.325***	0.320***
	(0.007)		(0.007)	(0.007)
Model Specification	OLS	FE	IV (local)	IV (incumbent)
First stage F			2072.08	1650.63
Observations	11597	11158	11597	8923
R^2	0.118	0.233	0.104	0.099

Note: Standard errors adjusted for clusters (by municipality) are in parentheses. Large out equals 1 if there is a party with more than 20% vote share that is not included in the coalition and 0 otherwise. Model (1) is a linear regression (OLS). Model (2) includes municipality-party and year fixed effects. In models (3) and (4), we instrument the exclusion of a large party with the large party's local/national and incumbency status, respectively. ***, **, and * denote significance at the 1%, 5%, and 10% levels, respectively.

Table 6 presents the results. All four specifications show that coalition parties are less likely to participate in the subsequent coalition if they exclude a large party. We see the model in Column (2), which includes municipality-party and year fixed effects, as the most reliable estimate. In this specification, the probability that a coalition party participates again in the subsequent coalition decreases by 17.8 percentage points when it excludes a large party, compared to coalition parties that include all large parties.

Tables A.4 and A.5 in the Appendix show that there are two strong predictors for the largest and second-largest parties' coalition participation. First, largest and second-largest local parties are less likely to be included in the coalition than largest and second-largest parties that also participate in national politics. Second, largest and second-largest parties are more likely to participate in the coalition if they also participated in the previous period. We exploit this in the models presented in columns (3) and (4) of Table 6, where we instrument participation by those two factors. The results are in line with the other two models. However, the exclusion restriction states that the instrument should not influence the outcome variable directly. That is, an excluded party's local/incumbency status must not impact the coalition parties' inclusion in the subsequent elections. This restriction is unlikely to hold. Therefore, the results in columns (3) and (4) can, just as the other two columns, not be interpreted as causal estimates.

5.6. **The broad-coalition norm versus stability.** While our results reject the minimal coalition hypothesis, they do not provide direct proof of a broad-coalition norm in Dutch local politics. Parties may also form oversized coalitions to insure against legislative instability. However, both our descriptive and causal findings indicate that stability concerns do not exclusively drive coalition formation.

Descriptive evidence from single-party majorities provides an initial counterargument. Oversized coalitions frequently emerge even when a single party holds an absolute majority of council seats and strictly requires no partners. In 85% of the 40 cases where the largest party secures a unilateral majority, it still chooses to form a coalition with another party, most often the runner-up. Furthermore, in 7.5% of these single-party majority cases, the largest party incorporates *two* superfluous partners to form a three-party executive.

Our regression discontinuity estimates provide a second counterargument against a pure stability mechanism. If stability were the sole objective, a core coalition lacking a majority would optimize its legislative buffer by incorporating a small partner necessary to comfortably cross the threshold. This strategy would maximize legislative security while minimizing the loss of executive power and allocated aldermen. Instead, our results show a preference for the third-largest party over the fourth-largest. Moreover, Table 5 demonstrates that the probability of incorporating the third party strictly increases with its size. Core parties actively select partners that broaden the coalition's aggregate electoral base, rather than minimizing the bargaining cost of stability.

Finally, the persistence of coalition expansion well beyond the majority threshold reinforces the concept of a broad-coalition norm over a stability motive. If core parties just sought a buffer to ensure legislative survival, coalition expansion would cease once a safe margin was achieved. Figures 5 and 6 contradict this expectation. Even when the core coalition already has a 60% majority, it still includes an additional party in approximately in more than 50% of observed cases. Moreover, Figure 6 shows that if the core has a majority, it often even includes multiple additional parties. This indicates that parties are not simply seeking a safe buffer. Rather, this persistent expansion is consistent with an underlying norm of representing a broad share of the electorate.

6. Discussion

We used rank-based election outcomes to test bargaining-power coalition-formation models à la Baron and Ferejohn. In the context of Dutch municipalities, we find no support for these models. Moreover, we present evidence against the minimal winning hypothesis. Instead, using vote-share discontinuities, we find evidence consistent with a broad-coalition norm prescribing that the total vote share of coalition parties should be substantially above 50%. This norm also prescribes that large parties should participate in the coalition and suggest that dynamic penalties for violating this norm apply.

A fundamental function of democracy is to align policy with voter preferences. A broad-coalition norm may help the government to be responsive to the preferences of a large majority of the electorate. Another possible advantage of a broad coalition

is that it shifts incentives from distributive policies to common good policies. However, it is worth emphasizing that, in the Dutch context, coalitions that encompass the entire council are rare. Councils have opposition parties that should hold the coalition accountable [cf.Persson et al. (1997)]. An interesting topic for future research is to examine whether broad coalitions pursue different kinds of policies than minimal coalitions.

In a way, our result runs counter to existing theoretical models of coalition formation. We do not claim that the Dutch broad-coalition norm generalizes mechanically. Instead, we show that coalition bargaining equilibria can differ sharply across institutional environments, and that local governments provide enough close-election variation to identify which equilibrium is operating.

From a theoretical perspective, a norm functions as an equilibrium outcome. Because models of social norms frequently generate multiple equilibria (Fujiwara and Sanz, 2020), countries can adhere to different norms. Furthermore, institutional differences often act as focal points that select which equilibrium is realized. In Spain, the statutory rules dictate that the largest party provides the mayor if parties fail to reach an agreement. The formal system therefore establishes the largest party as the default option, fostering a rank-based norm. In the Netherlands, the national government appoints the mayor. Negotiations *must* continue until a coalition is installed. Without a default that favors the largest party, Dutch municipal negotiations lack a rank-based focal point, allowing the broad-coalition norm to emerge as the equilibrium outcome. Finally, in Norwegian municipalities, being large ensures inclusion in the coalition. Gravdahl (1998) argues that this institutional feature shapes a consensus norm.

Finally, norms about who should govern affect not only coalition formation, but also likely voting behavior. If ordinal rank matters, as in Spain, voters face incentives to vote strategically for front runners. In contrast, in the Dutch system, where absolute size and broad representation drive coalition inclusion, strategic voting on front runners makes less sense. This may help to explain the highly fragmented political landscape in the Netherlands.

REFERENCES

Ansolabehere, S., Snyder Jr, J. M., Strauss, A. B., and Ting, M. M. (2005). Voting weights and formateur advantages in the formation of coalition governments. *American Journal of Political Science*, 49(3):550–563.

Austen-Smith, D. and Banks, J. (1990). Stable governments and the allocation of policy portfolios. *American Political Science Review*, 84(3):891–906.

Axelrod, R. (1970). *Conflict if Interest*. Chicago: Markham.

Azari, J. R. (2014). *Delivering the people's message: The changing politics of the presidential mandate*. Cornell University Press.

Baron, D. P. (1991). A spatial bargaining theory of government formation in parliamentary systems. *American Political Science Review*, 85(1):137–164.

Baron, D. P. and Diermeier, D. (2001). Elections, governments, and parliaments in proportional representation systems. *The Quarterly Journal of Economics*, 116(3):933–967.

Baron, D. P. and Ferejohn, J. A. (1989). Bargaining in legislatures. *American Political Science Review*, 83(4):1181–1206.

Bassi, A. (2017). Policy preferences in coalition formation and the stability of minority and surplus governments. *The Journal of Politics*, 79(1):250–268.

Battaglini, M. (2021). Coalition formation in legislative bargaining. *Journal of Political Economy*, 129(11):3206–3258.

Calonico, S., Cattaneo, M. D., and Titiunik, R. (2014). Robust data-driven inference in the regression-discontinuity design. *The Stata Journal*, 14(4):909–946.

Carozzi, F., Cipullo, D., and Repetto, L. (2022). Political fragmentation and government stability: evidence from local governments in spain. *American Economic Journal: Applied Economics*, 14(2):23–50.

Carozzi, F., Cipullo, D., and Repetto, L. (2024). Powers that be? political alignment, government formation, and government stability. *Journal of Public Economics*, 230:105017.

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

Clinton, W. J. (1992). Address accepting the presidential nomination at the democratic national convention in new york. Accessed: 2025-06-26.

Diermeier, D., Eraslan, H., and Merlo, A. (2003). A structural model of government formation. *Econometrica*, 71(1):27–70.

Diermeier, D. and Merlo, A. (2000). Government turnover in parliamentary democracies. *Journal of Economic Theory*, 94(1):46–79.

Diermeier, D. and Merlo, A. (2004). An empirical investigation of coalitional bargaining procedures. *Journal of Public Economics*, 88(3-4):783–797.

Fujiwara, T. and Sanz, C. (2020). Rank effects in bargaining: evidence from government formation. *The Review of Economic Studies*, 87(3):1261–1295.

Gravdahl, H. P. (1998). Consensual coalitions? coalition formation in norwegian municipalities. *Scandinavian Political Studies*, 21(4):307–323.

Kendall, C. and Matsusaka, J. G. (2025). The common good and voter polarization. *Journal of the European Economic Association*, page jvaf059.

Kiesraad (2022). Verkiezingsuitslagen gemeenteraden 2022 [election results municipal councils 2022]. Databank Verkiezingsuitslagen.

Laver, M., Benoit, K., and Garry, J. (2003). Extracting policy positions from political texts using words as data. *American political science review*, 97(2):311–331.

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

Morelli, M. (1999). Demand competition and policy compromise in legislative bargaining. *American Political Science Review*, 93(4):809–820.

Mueller, D. C. (2003). *Public choice III*. Cambridge University Press.

Nannicini, T., Stella, A., Tabellini, G., and Troiano, U. (2013). Social capital and political accountability. *American Economic Journal: Economic Policy*, 5(2):222–250.

OECD and UCLG (2023). Subnational governments around the world: Structure and finance (country profile: Netherlands). Technical report, World Observatory on Subnational Government Finance and Investment.

Otjes, S., Nagtzaam, M., and Van Well, R. (2023). Scrutiny and policymaking in local councils: How parties use council tools. *Local Government Studies*, 49(5):1110–1134.

Persson, T., Roland, G., and Tabellini, G. (1997). Separation of powers and political accountability. *The Quarterly Journal of Economics*, 112(4):1163–1202.

Riker, W. H. (1962). *The theory of political coalitions*. Yale University Press.

Schofield, N. and Laver, M. (1985). Bargaining theory and portfolio payoffs in european coalition governments 1945–83. *British Journal of Political Science*, 15(2):143–164.

Von Neumann, J. and Morgenstern, O. (1947). Theory of games and economic behavior, 2nd rev.

APPENDIX

FIGURE A.1. Histogram Running Variable: Difference between vote shares of the largest and second-largest party

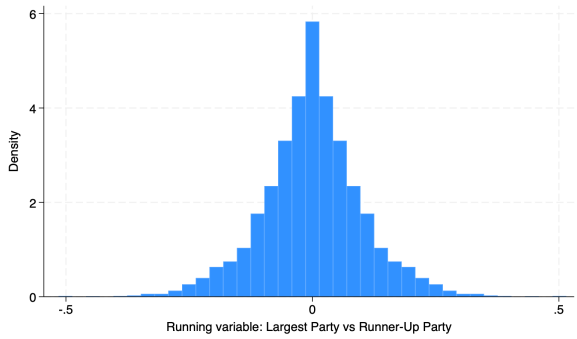

FIGURE A.2. Histogram Running Variable: Difference between vote shares of the runner-up and third-largest party

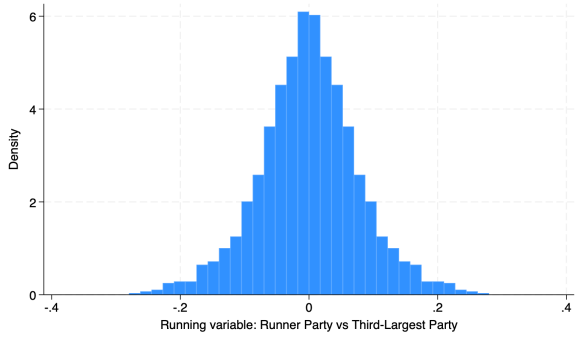

TABLE A.1. Robustness Check: RDD estimates correcting for mechanical dependence (Random Selection Test)

	(1)	(2)
	P1 vs. P2	P2 vs. P3
RDD Estimate	0.006	0.015
	(0.042)	(0.123)
Robust *p*-value	0.878	0.903
Running Variable	Margin 1 & 2	Margin 2 & 3
Observations	2672	797
Bandwidth	0.081	0.048

Note: To account for the mechanical dependence (zero-sum nature) of coalition inclusion within municipalities, the estimation sample is restricted to one randomly selected party per election to eliminate double-counting. Column (1) tests the inclusion probability of the largest party versus the runner-up. Column (2) tests the runner-up versus the third party, conditional on the largest party participating in the coalition and the first and third party holding a joint majority. The coefficients are conventional point estimates from a local linear regression using the `rdrobust` package with a symmetric MSE-optimal bandwidth. Robust bias-corrected standard errors adjusted for clusters (by municipality) are in parentheses. ***, **, and * denote significance at the 1%, 5%, and 10% levels, respectively.

FIGURE A.3. The effect of being the runner-up rather than the third-largest party on coalition inclusion, for elections where the first and third parties are ideologically closer than the first and second.

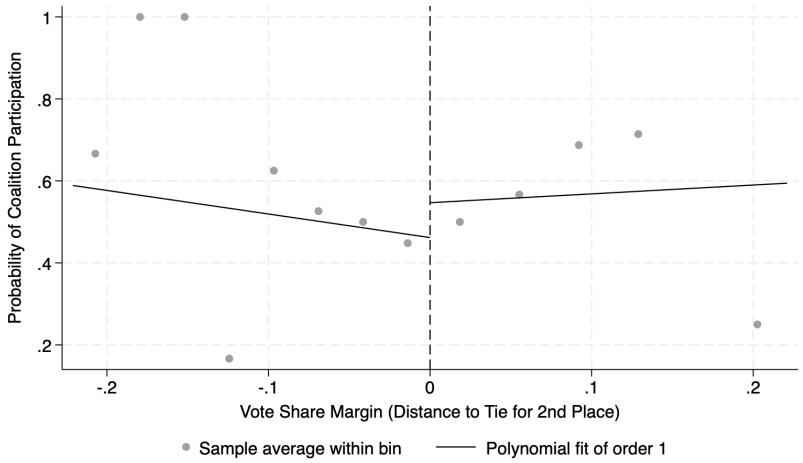

Note: The sample only includes observations where the first party participates in the coalition, the first and third party have a majority, and the first party is ideologically closer to the third party than to the second party. The ideological difference is based on the principal component of the general and economic left-right scale.

TABLE A.2. RDD estimates of being the runner-up rather than the third-largest party on coalition inclusion, for elections where the first and third parties are ideologically closer than the first and second

	(1)	(2)	(3)
	PCA	General	Economic
RDD Estimate	0.376	0.383	0.342
	(0.334)	(0.332)	(0.338)
Robust p-value	0.261	0.248	0.313
Observations	186	200	196
Bandwidth	0.036	0.035	0.036

Note: The sample only includes observations where the first party participates in the coalition, the first and third party have a majority in votes, and the first party is ideologically closer to the third party than to the second party. The ideological difference is based on the general left-right dimension in column (2), the economic left-right dimension in column (3), and the principal component of these two dimensions in column (1). Each column reports the coefficient from a local linear regression estimated using the `rdrobust` package with a triangular kernel, polynomial of order 1, and MSE-optimal bandwidth. Robust bias-corrected standard errors adjusted for clusters (by municipality) are in parentheses. ***, **, and * denote significance at the 1%, 5%, and 10% levels, respectively.

TABLE A.3. RDD estimates on bargaining predictions separate for different council sizes

	(1)	(2)	(3)	(4)
RDD Estimate	-0.039	0.019	0.113	0.042
	(0.037)	(0.055)	(0.121)	(0.103)
Robust p-value	0.725	0.588	0.351	0.684
Observations	3112	2232	600	994
Bandwidth	0.067	0.093	0.056	0.062
Council size	Large	Small	Large	Small
Hypothesis	1 vs 2	1 vs 2	2 vs 3	2 vs 3

Note: Each column reports the coefficient from a local linear regression estimated using the `rdrobust` package with a triangular kernel, polynomial of order 1, and MSE-optimal bandwidth. Robust bias-corrected standard errors adjusted for clusters (by municipality) are in parentheses. The analyses in columns (1) and (3) are on municipalities with 19 or more council seats and in columns (2) and (4) on municipalities with fewer than 19 council seats. Columns (1) and (2) test the prediction that the first party is more likely to participate in the coalition than the second party. Columns (3) and (4) test the prediction that the third party is more likely to participate in the coalition than the second party. ***, **, and * denote significance at the 1%, 5%, and 10% levels, respectively.

TABLE A.4. Differences between largest parties that participate in the coalition and largest parties that do not participate in a coalition

	Full	Largest in	Largest out	t
National	0.706	0.647	0.442	-6.918***
Left-right pca	0.000	0.396	0.336	-0.484
Extreme ideology	0.305	0.230	0.217	-0.272
Municipality seats	21.555	21.822	19.412	-4.917***
In Coalition last term	0.418	0.712	0.513	-6.122***

Note: The column with 'full' gives the mean in the full sample. The column 'Largest in' gives the mean for largest parties who participate in the coalition. The column 'Largest out' gives the mean for the largest parties that do not participate in the coalition. The two-sided t-test is performed on the difference between 'Largest in' and 'Largest out'. 'National' equals 1 if a party is also active in national elections and 0 otherwise. 'Extreme ideology' equals 1 if the parties Left-right ideology pca differs more than one standard deviation from the average pca and 0 otherwise. 'Coalition last year' equals 1 if the party was part of the previous coalition and 0 otherwise. ***, **, and * denote significance at the 1%, 5%, and 10% levels, respectively.

TABLE A.5. Differences between runner-up parties that participate in the coalition and runner-up parties that do not participate in a coalition

	Full	Second in	Second out	t
National	0.706	0.743	0.594	-7.604***
Left-right pca	0.000	0.308	0.273	-0.370
Extreme ideology	0.305	0.285	0.281	-0.124
Municipality seats	21.555	22.009	20.380	-4.748***
In coalition last term	0.418	0.664	0.505	-6.768***

Note: The column with 'full' gives the mean in the full sample. The column 'Second in' gives the mean for runner-up parties who participate in the coalition. The column 'Second out' gives the mean for runner-up parties that do not participate in the coalition. The two-sided t-test is performed on the difference between 'Second in' and 'Second out'. 'National' equals 1 if a party is also active in national elections and 0 otherwise. 'Extreme ideology' equals 1 if the parties Left-right ideology pca differs more than one standard deviation from the average pca and 0 otherwise. 'Coalition last year' equals 1 if the party was part of the previous coalition and 0 otherwise. ***, , and * denote significance at the 1%, 5%, and 10% levels, respectively.

FIGURE A.4. Histogram Running Variable: the core's margin to a majority

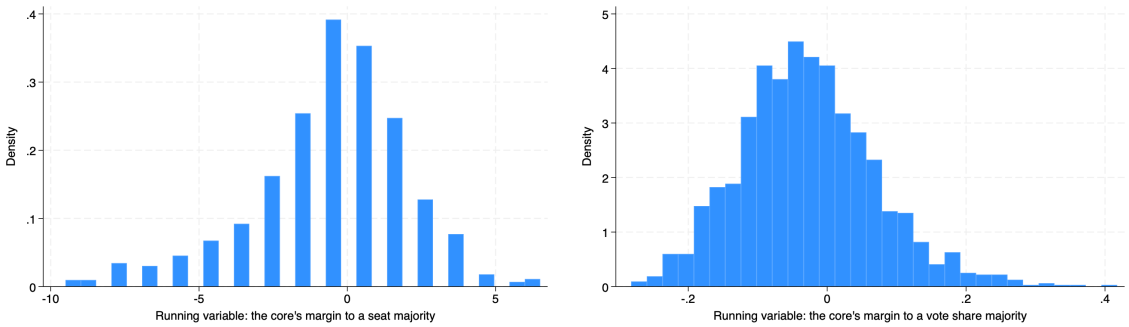

(A) (B)

Note: In Panel A, the running variable is the core's seat margin. In Panel B, the running variable is the core's vote share margin.

TABLE A.6. Selection test: RDD estimates of the core's majority on the core's joint participation in the coalition

	(1)	(2)
RDD Estimate	0.015	-0.067
	(0.053)	(0.069)
Robust p-value	0.774	0.330
Running variable	Seats	Votes
Observations	2226	2226
Bandwidth	4.751	0.079

Note: The coefficients are estimated from a local linear regression using the **rdrobust** package with a triangular kernel, a polynomial of order 1, and an MSE-optimal bandwidth. Robust bias-corrected standard errors adjusted for clusters (by municipality) are in parentheses. ***, **, and * denote significance at the 1%, 5%, and 10% levels, respectively.

TABLE A.7. Balance test: RDD estimates of the core's majority in seats on covariates

Covariate	RDD Estimate	Standard Error	Robust p-value	Observations	Bandwidth
Number of parties	0.084	(0.067)	0.214	1411	4.617
Vote share first	0.002	(0.005)	0.690	1411	3.452
Vote share second	-0.001	(0.005)	0.758	1411	4.238
Vote share third	-0.001	(0.005)	0.888	1411	4.187
Incumbent	-0.039	(0.083)	0.636	1087	4.518
Local	-0.017	(0.078)	0.826	1411	4.105
PvdA	0.036	(0.069)	0.599	1411	4.490
CDA	-0.048	(0.075)	0.526	1411	5.053
VVD	0.001	(0.072)	0.988	1411	5.102
D66	-0.015	(0.024)	0.532	1411	4.551
GL	-0.025	(0.025)	0.316	1411	3.414
SP	-0.000	(0.023)	0.998	1411	5.299
CU	-0.020	(0.037)	0.588	1411	4.615

Note: The sample only includes observations in which the largest party and the runner-up both participate in the coalition. Each row reports the coefficient from a separate local linear regression estimated using the `rdrobust` package with a triangular kernel, polynomial of order 1, and MSE-optimal bandwidth. Robust bias-corrected standard errors adjusted for clusters (by municipality) are in parentheses. Number of parties indicates the total number of parties that participated in the elections. Vote share first, second, and third measure the vote share of the largest, second-largest, and third-largest party, respectively. Incumbent, local, and the party variables equal 1 if the core includes at least one party with this characteristic and 0 otherwise. None equals 1 if all the parties above are not part of the core. PvdA is a left-leaning social-democratic party, CDA a Christian democratic party, VVD a right-leaning liberal party, D66 a left-leaning liberal party, GL a green party, SP a socialist party, and CU a Christian social party. ***, **, and * denote significance at the 1%, 5%, and 10% levels, respectively.

TABLE A.8. Balance test: RDD estimates of the core's majority in vote share on covariates

Covariate	RDD Estimate	Standard Error	Robust p-value	Observations	Bandwidth
Number of parties	-0.017	(0.084)	0.837	1411	0.068
Vote share first	-0.003	(0.005)	0.561	1411	0.088
Vote share second	0.003	(0.005)	0.561	1410	0.088
Vote share third	0.003	(0.005)	0.561	1411	0.088
Incumbent	-0.002	(0.098)	0.984	1087	0.091
Local	-0.051	(0.088)	0.562	1411	0.087
PvdA	-0.017	(0.087)	0.847	1411	0.080
CDA	0.054	(0.099)	0.586	1411	0.068
VVD	0.010	(0.091)	0.917	1411	0.073
D66	0.005	(0.029)	0.863	1411	0.054
GL	-0.035*	(0.021)	0.099	1411	0.059
SP	-0.025	(0.026)	0.333	1411	0.047
CU	-0.012	(0.022)	0.587	1411	0.051

Note: The sample only includes observations in which the largest party and the runner-up both participate in the coalition. Each row reports the coefficient from a separate local linear regression estimated using the `rdrobust` package with a triangular kernel, polynomial of order 1, and MSE-optimal bandwidth. Robust bias-corrected standard errors adjusted for clusters (by municipality) are in parentheses. Number of parties indicates the total number of parties that participated in the elections. Vote share first, second, and third measure the vote share of the largest, second-largest, and third-largest party, respectively. Incumbent, local, and the party variables equal 1 if the core includes at least one party with this characteristic and 0 otherwise. None equals 1 if all the parties above are not part of the core. PvdA is a left-leaning social-democratic party, CDA a Christian democratic party, VVD a right-leaning liberal party, D66 a left-leaning liberal party, GL a green party, SP a socialist party, and CU a Christian social party.

FIGURE A.5. The probability that the core (the largest party and the third-largest party) expands

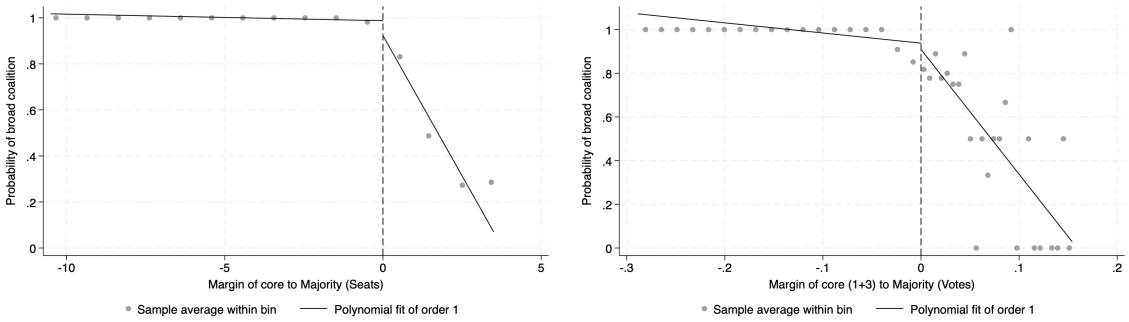

(A) Running variable is seat share core coalition half the council size

(B) Running variable is vote share core coalition minus one half

Note: The sample includes only observations in which the largest and third-largest parties participate in the coalition, the second-largest party does not participate in the coalition, and the council has more than two aldermen.