

ARTICLE

# Refining Gamson: The Isometric Log-Ratio Transformation and Portfolio Proportionality in Multiparty Governments

Lanny W. Martin<sup>1</sup>  and Georg Vanberg<sup>2</sup>

<sup>1</sup>Department of Social and Political Sciences, Bocconi University, Milan, Italy; <sup>2</sup>Department of Political Science, Duke University, Durham, USA

**Corresponding author:** Lanny W. Martin; Email: [lanny.martin@unibocconi.it](mailto:lanny.martin@unibocconi.it)

(Received 30 May 2025; revised 16 December 2025; accepted 17 December 2025)

## Abstract

One of the most robust empirical findings in political science is that in multiparty democracies cabinet ministries are distributed in rough proportion to parties' legislative seat shares, a pattern known as Gamson's Law. Yet existing research often overlooks the fact that portfolio and seat shares are compositional—mutually dependent parts of a whole. Standard methods treat them as unconstrained, risking bias, misleading uncertainty estimates, and flawed inference. Unfortunately, the most common strategy for handling compositions—the additive log-ratio (ALR) transformation combined with seemingly unrelated regression (SUR)—fails when the number and identity of compositional elements (like parties) vary across cases. We propose the isometric log-ratio (ILR) transformation—new to political science—as an alternative that both respects compositional geometry and adapts to differing compositional structures. Monte Carlo simulations show that ILR sharply outperforms standard approaches, reducing bias, improving coverage probability, and increasing statistical power. While we apply it to portfolio allocation, ILR provides a general solution for modeling compositional outcomes with other potential uses, including in electoral competition, where ALR+SUR has required strong assumptions or ad hoc adjustments. Using this improved methodology, we find that seat–portfolio proportionality is weaker overall than conventionally reported and varies substantially across governments.

**Keywords:** Gamson's law; coalition politics; compositional data

**Edited by:** Daniel J. Hopkins and Brandon M. Stewart

## 1. Introduction

One of the most robust empirical regularities in political science—Gamson's Law—is the finding that in coalition governments parties receive a share of cabinet ministries (portfolios) roughly proportional to their share of legislative seats. Proposed by Gamson (1961) as an intuitive principle for dividing collective rewards, the *seat-proportionality hypothesis* has been a cornerstone of research on multiparty governments for more than 50 years. Browne and Franklin (1973) provided the first systematic test using real-world coalition data; since then, numerous studies in leading journals have reassessed the strength of the seat–portfolio relationship and investigated the conditions under which proportionality is stronger or weaker (e.g., Carroll and Cox 2007; Cox 2021; Cutler *et al.* 2016; Martin and Vanberg 2020; Schofield and Laver 1985; Warwick and Druckman 2006).

This sustained attention is hardly surprising. Cabinet portfolios are central policymaking institutions in parliamentary democracies; the number of ministries a party controls shapes its influence over the legislative agenda and final policy outcomes. Yet empirical research on portfolio allocation faces a fundamental limitation: both the dependent variable (portfolio share) and the key independent variable (seat share) are compositional. Conventional approaches, which typically rely on linear regression,

largely ignore this fact—risking biased coefficients, incorrect uncertainty estimates, and misleading inference.

The pitfalls of applying linear models to compositional data are well known—serious enough, in the context of electoral competition, to motivate a special issue of *Political Analysis* in 2002 (e.g., Honaker, Katz, and King 2002; Jackson 2002; Tomz, Tucker, and Wittenberg 2002). Despite its widespread use, the prevailing solution in that literature—the additive log-ratio (ALR) transformation combined with seemingly unrelated regression (SUR)—is poorly suited to the analysis of portfolio allocation, where the number and identity of coalition parties usually vary across governments. In this study, we propose an alternative: the isometric log-ratio (ILR) transformation introduced by Egozcue *et al.* (2003) in the field of geomathematics. To our knowledge, this is the first application of ILR in political science, though its advantages suggest broad potential for analyzing compositional outcomes beyond portfolio allocation. Indeed, ILR offers a more coherent and flexible foundation for modeling electoral competition as well—particularly in settings where the limitations of ALR+SUR have required researchers to adopt strong assumptions or complex workarounds. In what follows, we explain the limitations of existing approaches to modeling portfolio allocation, assess why ALR+SUR falls short in this context, and show how the ILR framework offers a better solution than standard models.

## 2. Models of Portfolio Allocation: The Unsuitability of Current Approaches

Most studies of quantitative portfolio allocation—the question of “how many” portfolios parties receive—share several features: (1) government parties are the units of analysis; (2) the dependent variable is defined as party portfolio share; (3) one of the covariates is a party’s legislative size relative to its partners; and (4) the analysis entails a simple linear regression of portfolio share on seat share, sometimes with additional covariates like bargaining power, proposer status, or interactions between seat share and contextual features of the bargaining environment. Unfortunately, this empirical strategy overlooks the unique structure of portfolio allocation data and is ill-suited to its analysis.

Both portfolio and seat share are inherently *compositional*: they are proportions representing “parts of a whole” and must sum to one across all parties in a given government. With few exceptions (discussed below), this fundamental property is largely ignored in the literature. Most scholars employ least-squares regression (OLS), which assumes that observations are independent, the dependent variable is unbounded, and the relationship between predictors and the outcome is linear. None of these assumptions hold in the compositional setting. As parts of a whole, individual observations are necessarily interdependent: if one party’s portfolio share increases, at least one other party’s must decrease. In addition, portfolio shares are bounded between 0 and 1, which means that their potential to change is constrained near the limits of the scale, undermining the assumption of linear marginal effects. Collectively, these features imply that portfolio and seat shares lie not in unconstrained Euclidean space but in a unit simplex (Aitchison 1982, 1986).<sup>1</sup> Figure 1 illustrates this geometry in the three-party case, where shares are expressed as barycentric coordinates: the right panel shows a roughly even division, while the left panel depicts a distribution closer to the simplex bounds.

### 2.1. Linear Regression and Its Limitations

Because portfolio shares are constrained in this way, using untransformed shares in a linear regression leads to multiple problems. First, the *bounded nature* of the dependent variable can produce (1) implausible predictions, e.g., portfolio shares below zero or above one; (2) inaccurate uncertainty estimates due to nonnormal and heteroscedastic residuals, especially near the boundaries of the simplex; and (3) biased estimates of covariate effects, since the bounded scale compresses variance and distorts the true functional form. Second, the *non-independence* of observations implies correlated residuals across parties within a government, further undermining valid inference. Third, the *constant-sum constraint* introduces three interrelated problems. One is that the model can produce predicted shares

<sup>1</sup>The unit simplex in  $K$  dimensions,  $S^{K-1}$ , is the set of all points  $x = (x_1, x_2, \dots, x_K)$  in  $\mathbb{R}^K$  such that  $x_i \geq 0$  for all  $i$  and  $\sum x_i = 1$ .

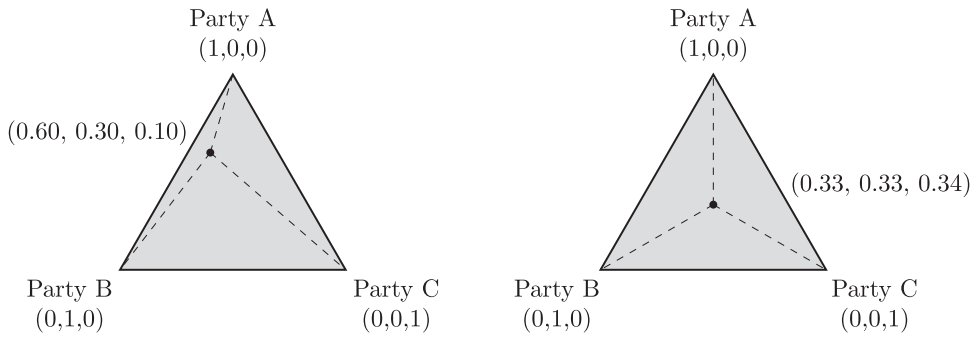


Figure 1. Unit simplex diagrams with different portfolio (seat) shares.

that do not sum to one. A second is the creation of spurious correlations among party shares: increases in one party's share must mechanically be offset by decreases in others, even when no substantive inverse relationship exists (Aitchison 1986; Pawlowsky-Glahn and Egozcue 2006; Pawlowsky-Glahn, Egozcue, and Tolosana-Delgado 2015). A third is the presence of perfect multicollinearity: in any government with  $K$  parties, the  $K$  shares contain only  $K - 1$  linearly independent observations. Including all  $K$  in a regression model renders the design matrix rank-deficient—producing unreliable coefficient estimates, standard errors that are too narrow, and artificially inflated degrees of freedom.

Almost all studies of Gamson's Law that have acknowledged the compositional nature of portfolio and seat shares have attempted to address this latter problem, rank deficiency, in the same way: by excluding one party from each government. The first to do so were Fréchet, Kagel, and Morelli (2005), who dropped the proposer from each bargaining round in their experimental design. Later studies using observational coalition data—including Carroll and Cox (2007), Cox (2021), and our own work (Martin and Vanberg 2020)—typically dropped one party at random from each government. This strategy resolves only one of the problems described above: it ensures the design matrix is full rank. But it does not address dependency among the remaining observations. Unless the government consists of just two parties, any increase in portfolio share for one included party is still likely to correspond to a decrease for another. Moreover, the strategy does nothing to address the consequences of boundedness and the constant-sum constraint: implausible predictions, heteroscedastic residuals, and spurious correlations that distort the estimated effects of covariates.<sup>2</sup>

Finally, because it ignores the data structure, if there exists *heterogeneity* in slopes across governments, OLS may be unable to detect it. Even in multilevel form, OLS remains vulnerable to distortions from boundedness and the constant-sum constraint, which compress variation across units. As a result, genuine differences in slopes may be attenuated or obscured, leading OLS to struggle in recovering random slope variance present in the data-generating process (DGP). We investigate this limitation directly in our simulations below.

## 2.2. The ALR+SUR Solution

Political scientists have long recognized the limitations of applying OLS to compositional data. In early studies of alternative approaches (e.g., Honaker *et al.* 2002; Jackson 2002; Katz and King 1999;

<sup>2</sup>An alternative approach from Cutler *et al.* (2016) modeled portfolio share using a beta distribution. This helps address boundedness, since the beta is defined on the unit interval and allows for heteroscedasticity. However, it treats each share as independent of the others and ignores the constant-sum constraint—permitting predictions that do not sum to one and failing to capture the inherent interdependence among components, which leaves intact the problems of spurious correlation and invalid inference. Another possible alternative, Dirichlet regression, does respect the simplex structure and models outcomes jointly. But like the ALR+SUR framework discussed below, it requires that the same components (e.g., parties) appear in every observation—a condition rarely met in studies of portfolio allocation.

Tomz *et al.* 2002), the outcome of interest was typically each party's share of the vote across electoral districts. More recent work has applied compositional methods to domains such as public budgeting (e.g., Clayton and Zetterberg 2018; Philips, Rutherford, and Whitten 2016), where the dependent variable represents the allocation of a fixed resource (e.g., spending shares across categories).

To accommodate the boundedness and constant-sum constraint inherent in such data, these studies employ the ALR transformation introduced by Aitchison (1982, 1986). The ALR maps a  $K$ -part composition from the unit simplex  $S^{K-1} \subseteq \mathbb{R}^K$  to  $\mathbb{R}^{K-1}$  by expressing each component relative to a designated reference part. For example, choosing component  $v_n$  as the reference yields the transformed vector:

$$\left[ \ln\left(\frac{v_1}{v_n}\right), \ln\left(\frac{v_2}{v_n}\right), \dots, \ln\left(\frac{v_{n-1}}{v_n}\right) \right].$$

This transformation preserves the relative information in the composition while allowing analysts to apply linear modeling techniques in an unconstrained space.<sup>3</sup>

However, because all log-ratios share a common denominator, the ALR transformation induces strong correlations among the transformed components. This violates the assumption of independent errors across equations and, if unaccounted for, leads to inefficient estimates and invalid standard errors. To address this, nearly all political science applications of ALR have coupled it with SUR, which models the log-ratio equations jointly while accounting for contemporaneous correlation in the residuals. SUR estimation is particularly well-suited to this task: when the covariates differ across equations, it improves efficiency; when the covariates are the same, it still yields more accurate standard errors (Zellner 1962). Without SUR—or some alternative method to account for residual dependence—the use of ALR alone is inadequate. It fails to respect the interdependence of compositional parts and will generally produce misleading inferences.

One might think that the ALR+SUR strategy could easily be extended to portfolio allocation. Yet as Indridason (2015) notes, the approach encounters fundamental limitations in this context. Most obviously, SUR requires a *fixed system of equations* across all observations. But in the context of portfolio allocation, the number of government parties—and hence the number of compositional components—varies across governments. When the number of parties changes, the number of ALR-transformed variables changes accordingly, making it impossible to specify a consistent system of equations for joint estimation. This structural mismatch renders SUR inapplicable, *and therefore makes ALR+SUR effectively unusable*.

Even apart from this fundamental barrier, the ALR transformation itself becomes problematic when applied to varying coalition structures. Because it is defined relative to a single baseline component, ALR assumes that *the same party occupies the denominator role in every observation*. In within-country studies of electoral competition, at least where there exists one party that competes in every district, this assumption is satisfied; researchers can use that party (e.g., the UK Conservative Party) as a constant reference component.

But in portfolio allocation, the identity of governing parties varies both across countries and over time. As a result, no stable reference party exists. This presents a serious complication: when the reference party varies across governments, each log-ratio is anchored to a different baseline. This introduces observation-specific scaling: systematic differences in the chosen baseline are pushed into the error term or confounded with covariates, yielding biased coefficients and misleading inference. In short, ALR+SUR relies on two structural conditions—*invariant baselines* (ALR) and *fixed dimensionality* (SUR)—that do not hold in this setting. Analysts must therefore turn to an alternative that remains valid when the number and identity of components vary.

<sup>3</sup>The ALR transformation implicitly resolves the redundancy of compositional data by removing one part (the reference) from the system of equations. By construction, the log of the reference part is zero ( $\ln(v_n/v_n) = 0$ ), which drops one column from the transformed design matrix and ensures full rank.

### 3. The ILR Transformation

The limitations of the ALR+SUR framework might suggest that log-ratio transformations are poorly suited for analyzing portfolio allocation.<sup>4</sup> However, this impression is misleading. ALR is just one member of a broader family of log-ratio transformations developed to handle the unique constraints of compositional data. Critically, one alternative—the ILR transformation introduced by Egozcue *et al.* (2003)—avoids the central pitfalls of ALR+SUR. It does not require a fixed reference component, can be applied to compositions of varying dimensionality, and produces transformed variables that are uncorrelated and geometrically well-behaved.

Like the ALR, the ILR transformation maps compositions from the constrained unit simplex  $S^{K-1}$  into  $\mathbb{R}^{K-1}$ , thereby allowing standard regression techniques to be applied. But unlike ALR, which expresses each part relative to a single denominator, ILR constructs  $K - 1$  orthogonal coordinates—called *balances*—based on a sequential binary partition (SBP) of the  $K$  compositional parts. Each ILR coordinate compares the geometric mean of one subset of parts to the geometric mean of its complement:

$$ILR_b = \ln \left( \frac{g(S_b)}{g(T_b)} \right),$$

where  $g(S_b)$  and  $g(T_b)$  denote the geometric means of disjoint subsets such that  $S_b \cup T_b = \{1, \dots, K\}$  and  $S_b \cap T_b = \emptyset$ . Typically,  $S_b$  contains a single part and  $T_b$  the remaining  $K - b$  parts.

For example, in a five-party government with parties A through E, one valid SBP produces:

$$ILR_1 = \ln \left( \frac{A}{\sqrt[4]{B \cdot C \cdot D \cdot E}} \right), \quad ILR_2 = \ln \left( \frac{B}{\sqrt[3]{C \cdot D \cdot E}} \right),$$

$$ILR_3 = \ln \left( \frac{C}{\sqrt{D \cdot E}} \right), \quad ILR_4 = \ln \left( \frac{D}{E} \right).$$

Each balance contrasts one party with the geometric mean of a successively smaller group of others. The ordering of parts is arbitrary, but any valid SBP produces an orthogonal basis in the simplex (Egozcue *et al.* 2003, 290–292), ensuring that each balance captures unique, uncorrelated information.<sup>5</sup>

As with ALR, analysts must specify a basis when constructing ILR coordinates. In ALR, this means choosing a reference category; in ILR, defining the SBP. These choices affect the labeling and

<sup>4</sup>Indeed, given the limitations of OLS and ALR+SUR, some researchers have shifted away from directly testing Gamson's Law, effectively setting aside the enterprise of evaluating seat–portfolio proportionality. For example, Indridason (2015) turned to a fixed-effects Poisson specification, which treats portfolios as counts. As he notes, however, such models “come at a cost when the objective is to evaluate Gamson's Law” (19): coefficients capture multiplicative effects on counts rather than linear effects on shares. This makes it impossible to test the proportionality hypothesis (slope = 1, intercept = 0) or to evaluate contextual theories—such as those offered by Carroll and Cox (2007), Cox (2021), and many others—about when proportionality should be stronger or weaker. By contrast, the ILR framework introduced here treats portfolios explicitly as compositional outcomes, permitting direct tests of proportionality and its contextual variation.

<sup>5</sup>An orthogonal basis is a set of linearly independent, uncorrelated coordinates. In the ILR context, this ensures that each balance captures a unique contrast between subsets of components, and no coordinate is a linear combination of the others. While the ILR-transformed variables are not statistically independent in a strict sense, their orthogonality eliminates linear dependencies, enabling the use of standard regression techniques without modeling residual covariance (we thank Vera Pawlowsky-Glahn for drawing our attention to this point). Optionally, each ILR balance can be scaled by a factor of  $\sqrt{\frac{s-t}{s+t}}$ , where  $s$  and  $t$  are the number of components in the numerator and denominator subsets, respectively. This produces an orthonormal basis (standardized orthogonal vectors) rather than merely an orthogonal one. We omit this weighting because our focus is on interpreting model coefficients and comparative effects, not on geometric distances or variance decomposition. The unweighted balances remain orthogonal, satisfy the requirements for linear modeling, and offer more straightforward interpretation.

interpretation of coefficients but not the underlying results: model fit and predictions are invariant to either choice. In ALR, coefficients from one baseline can be obtained from another through simple subtraction; in ILR, the same principle applies, though implemented through orthogonal transformations.<sup>6</sup> While orthogonal rotation is more cumbersome than ALR's subtraction, it does not change the mathematics of the model. Thus, in both approaches, baseline or partition choices matter only for interpretation, not inference.

Most crucially for portfolio allocation, the ILR transformation solves the dimensionality problem. Unlike ALR+SUR—which requires the same set of components and a fixed denominator across observations—ILR constructs  $K_g - 1$  balances from the parts actually present in each government  $g$ . If one government includes five parties and another only three, ILR simply constructs four balances for the former and two for the latter. As party systems change across cabinets, ILR thus permits pooling without imputation, collapsing parties, or estimating separate models by composition. By contrast, earlier strategies—e.g., imputing shares for non-competing parties, fitting distinct ALR+SUR systems for each pattern of party absence, or collapsing minor parties into “Other” (Honaker *et al.* 2002; Jackson 2002; Katz and King 1999; Tomz *et al.* 2002)—impose strong assumptions, fragment analyses, and undermine comparability. In short, ILR provides a single coherent framework that accommodates variation in both the number and identity of parties while preserving the simplex constraints and the joint nature of the composition.<sup>7</sup>

#### 4. Monte Carlo Simulations

We evaluate the ILR approach via Monte Carlo simulations, benchmarking it against standard OLS estimators. Because our later empirical application uses multilevel models, we include both single- and multilevel specifications for the simulations. ILR is estimated only in multilevel form and compared to four OLS alternatives: *OLS-LOO*, *OLS-LOO-CSE*, and their multilevel counterparts *OLS-ML* and *OLS-ML-CSE*. All OLS variants use the “leave-one-out” (LOO) strategy; CSE models cluster-robust standard errors at the government level. We focus on two parameters: the mean slope linking seats to portfolios ( $\bar{\beta}$ ) and the random-slope variance ( $\tau$ ).

##### 4.1. Design

In each replication, we simulate seat shares for the parties in each government as valid compositions on the simplex, and generate portfolio shares from:

$$z_{jg} = \beta_g s_{jg} + \varepsilon_{jg}, \quad \beta_g \sim N(\bar{\beta}, \tau^2), \quad \varepsilon_{jg} \sim N(0, \sigma^2),$$

where  $s_{jg}$  and  $z_{jg}$  are ILR coordinates of seats and portfolios. The overall intercept is fixed at zero. We treat governments as consisting of different parties across cases, which reflects the empirical reality of cabinet formation and immediately rules out ALR, which requires the same reference denominator across all observations.

<sup>6</sup>Formally, the ILR of  $x$  is  $z = V^\top \ln(x)$ , where  $V$  is the contrast matrix determined by the SBP. For two partitions with matrices  $V$  and  $V^*$ , the associated coordinates satisfy  $z^* = (V^*)^\top \ln(x) = (V^*)^\top Vz \equiv Rz$ , where  $R$  is an orthogonal rotation matrix. The same rotation links regression coefficients: if  $\hat{\beta}$  are estimated under one partition, then  $\hat{\beta}^* = R\hat{\beta}$  under another. This is directly analogous to ALR, where coefficients under one reference are linear transformations of those under another. In Appendix 1 of the Supplementary Material, we provide a more detailed elaboration and a numerical illustration.

<sup>7</sup>Beyond portfolio allocation, ILR can be applied to electoral competition with uneven party entry, pooled cross-national studies with differing party systems, and public budgets with shifting categories. The simulations below apply to any setting with compositional outcomes whose dimensionality varies across units or over time.

We fix the number of parties at  $K = 3$ , the median cabinet size in our data, and set the other parameters as follows:

- (1) A mean slope in ILR space of  $\bar{\beta} = 0.70$ . On the original simplex, this corresponds approximately to 4:3 proportionality between portfolios and seats, or an effective slope of about 0.74.<sup>8</sup> All hypothesis tests reported below are conducted in simplex space, treating 0.74 as the true value.
- (2) A random slope SD of  $\tau = 0.13$ , defined in ILR space, with the random intercept fixed at zero.
- (3) A residual SD of  $\sigma = 0.36$ .

We evaluate estimator performance across ten bins of a variable we call the *boundary proximity index* (BPI), which is defined at the government level as,

$$BPI_g = \frac{\sum_{j=1}^K p_{jg}^2 - \frac{1}{K}}{1 - \frac{1}{K}},$$

where  $p_{jg}$  is party  $j$ 's portfolio share in government  $g$ . Values near zero indicate balanced allocations at the simplex center, while values near one reflect concentration near the boundary. We use this index because treating compositional outcomes as unconstrained is most problematic at high BPI, where OLS estimators should perform worse. Each BPI bin contains 500 replications of 300 simulated three-party governments (i.e., 150,000 per bin). Bins go from low to high concentration: lower bins collect governments near the simplex center, higher bins those near the boundary.

#### 4.2. Metrics

Following Hopkins *et al.* (2024), we assess estimator performance for  $\hat{\beta}$  and  $\hat{\tau}$  using bias, root mean square error (RMSE), coverage probability, and statistical power.<sup>9</sup>

#### 4.3. Results for $\hat{\beta}$

Figures 2 and 3 summarize the findings for  $\hat{\beta}$  when  $K = 3$ .<sup>10</sup> Figure 2 shows that ILR-ML accurately recovers the true slope across the BPI range: bias remains near zero and RMSE is minimal. By contrast, OLS variants underestimate when  $BPI < 0.10$  and overestimate once  $BPI > 0.20$ , with bias rising toward the boundaries.<sup>11</sup> At  $BPI > 0.50$ , bias is about 0.13, which implies portfolio-to-seat proportionality of roughly 87%, much higher than the true value of 74%.<sup>12</sup> RMSE follows the same pattern, underscoring how OLS error compounds under concentration. The sign flip in OLS bias reflects the geometry of the ILR-to-simplex transformation: near the simplex center (low BPI), the relationship between seats and portfolios is slightly concave in share space, leading OLS to underestimate the slope; as compositions become more concentrated (higher BPI), the mapping turns convex, and OLS increasingly overestimates the slope. When governments of different sizes are pooled (Appendix 3 of the Supplementary Material), this asymmetry disappears and *OLS bias is positive throughout*, pushing slopes upward even at low BPI.

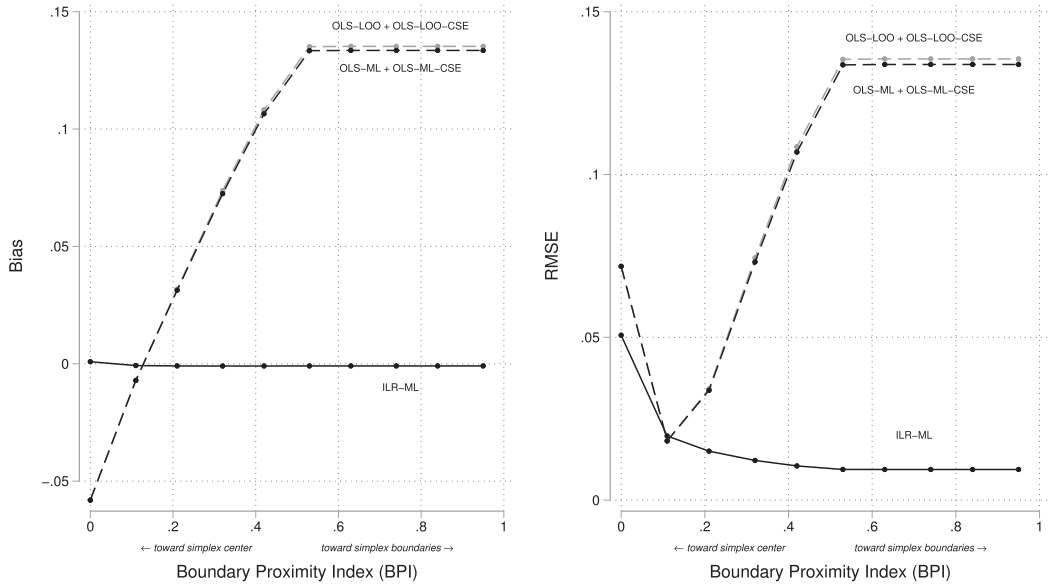
<sup>8</sup>The mapping between ILR and share-space slopes is not exact, but this normalization is natural: perfect proportionality corresponds to a unit slope in ILR space ( $\beta = 1$ ), so exponentiating compares the implied multiplicative change in balances under our  $\beta$  to that under  $\beta = 1$ . Using this scale,  $e^{0.70}/e^1 \approx 0.74$  is a good approximation to the simplex-scale benchmark in our designs.

<sup>9</sup>In Appendix 2 of the Supplementary Material, we report two additional metrics for  $\hat{\beta}$ : standard deviation (SD) and overconfidence.

<sup>10</sup>In Appendix 3 of the Supplementary Material, we vary  $K$  in two ways: by fixing  $K = 2, 4, 5, 6$  (the other cabinet sizes in our sample) and by mixing cabinet sizes (2/3/4 and 2/6). Allowing  $K$  to vary changes the pattern of OLS bias (see below) but not the superiority of ILR-ML.


<sup>11</sup>For bias and RMSE—point-estimate metrics—CSE and non-CSE coincide.

<sup>12</sup>In the assumed three-party case, a BPI of 0.50 corresponds to one party receiving 80% of portfolios and the other two 10% each.



Each point represents the average value of the metric across 500 replications of 300 three-party governments per BPI bin.

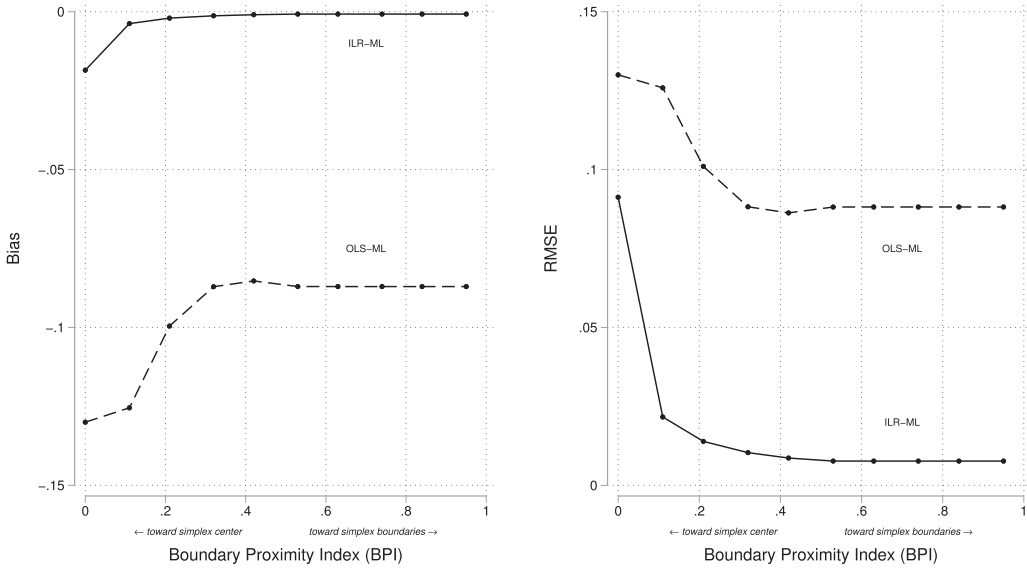
Figure 2. Bias and RMSE for  $\hat{\beta}$  across BPI.



Each point represents the average value of the metric across 500 replications of 300 three-party governments per BPI bin.

Figure 3. Coverage probability and power ( $H_0 : \beta \geq 0.80$ ) for  $\hat{\beta}$  across BPI.

In Figure 3, we see that coverage probability—the proportion of simulation replications in which the estimator’s confidence interval contains the true value—mirrors this pattern. ILR-ML maintains nominal 95% coverage across the BPI range. By contrast, the OLS estimators collapse almost immediately: once  $BPI > 0.20$ , their coverage plunges toward zero, meaning their 95% confidence intervals almost never contain the true slope.



Each point represents the average value of the metric across 500 replications of 300 three-party governments per BPI bin.

Figure 4. Bias and RMSE for  $\hat{\tau}$  across BPI.

The right panel turns to power. Conventional nulls of  $\beta = 0$  or  $\beta = 1$  are uninformative—biased OLS estimates lie far from either, so rejection is trivial. Instead we ask whether models can rule out 5:4 proportionality ( $\beta \geq 0.80$ ), the standard benchmark in the portfolio allocation literature. We define power as the probability that the 95% confidence interval lies entirely below 0.80, i.e.,  $P(\hat{\beta}_U < 0.80)$ , where  $\hat{\beta}_U$  is the 95% upper bound. ILR-ML shows high power across most of the BPI range, whereas OLS estimators rarely reject overly proportional allocations, consistent with their upward bias.<sup>13</sup>

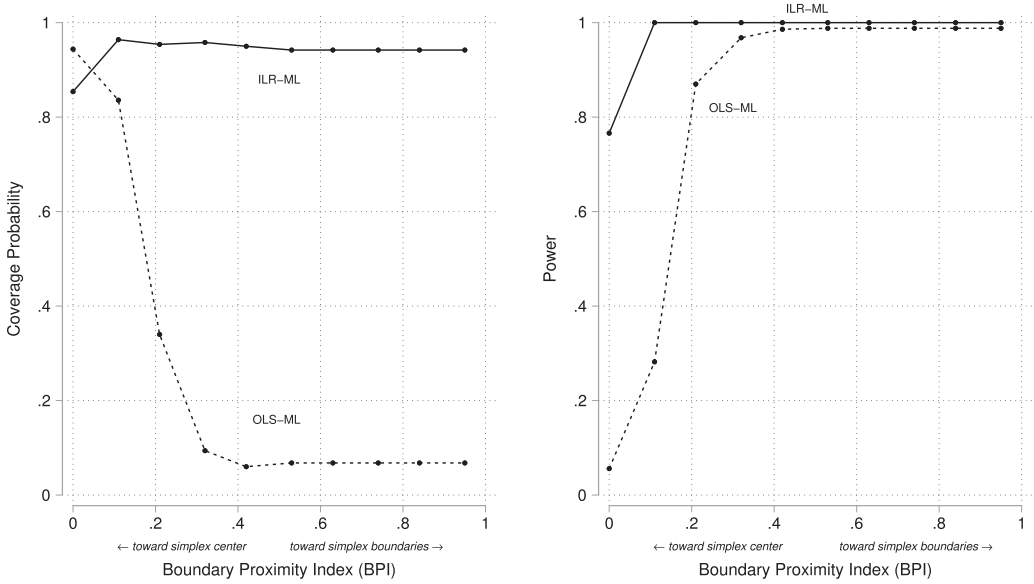
#### 4.4. Results for $\hat{\tau}$

Because variance components are only identifiable in multilevel models, results for  $\hat{\tau}$  compare ILR-ML and OLS-ML only (i.e., there is no OLS-LOO analog).<sup>14</sup> As shown in Figure 4, ILR-ML successfully recovers slope heterogeneity except at the simplex center, where balanced portfolios provide little signal and mild downward bias appears. RMSE tells the same story, peaking in the center where heterogeneity is hardest to identify, but falling sharply toward the boundaries. OLS-ML, by contrast, consistently estimates  $\hat{\tau}$  at or near zero (the assumed true value of  $\tau$  minus the estimated bias), producing what looks like low RMSE only because it is systematically wrong. Put differently, ILR-ML recovers slope heterogeneity except at the simplex center, where balanced governments naturally mask variance; OLS-ML misses it everywhere.

The coverage probability and power results from Figure 5 confirm this divergence. ILR-ML maintains coverage close to the 95% nominal level, with some erosion at the simplex center where slope heterogeneity is hardest to identify. OLS-ML, by contrast, shows essentially no coverage. Because its variance estimates concentrate at or near zero, the Wald-type confidence intervals are anchored on

<sup>13</sup>The exception for ILR-ML's high power is the first BPI bin: when allocations are nearly balanced, the true 0.74 is close to 0.80. This dip in power is not a deficiency of ILR but an expected property of one-sided margin tests: when the benchmark lies close to the true value, even well-calibrated estimators will often fail to reject (Lehmann and Romano 2005, 65).

<sup>14</sup>There is also no CSE variant for  $\hat{\tau}$ , since cluster-robust adjustments are not defined for variance components. Standard errors for  $\hat{\tau}$  are obtained from the mixed-effects likelihood using the observed Fisher information, which is an approximation when variance parameters lie near the zero boundary (Stram and Lee 1994).



Each point represents the average value of the metric across 500 replications of 300 three-party governments per BPI bin.

Figure 5. Coverage probability and power ( $H_0 : \tau = 0$ ) for  $\hat{\tau}$  across BPI.

the boundary and rarely include the true  $\tau = 0.13$ . This systematic failure is not a matter of random fluctuation, but a structural consequence of misspecified estimation: intervals that collapse at the boundary cannot provide valid coverage for a positive variance component.<sup>15</sup>

The power test in this case is  $H_0 : \tau = 0$ . ILR-ML behaves as expected: power is substantial even at the simplex center and quickly approaches one as  $BPI$  increases. OLS-ML rises much more slowly: at low  $BPI$ ,  $\hat{\tau}$  piles up at the boundary ( $\approx 0$ ), so the test rarely rejects; once  $BPI$  is high enough that OLS returns positive  $\hat{\tau}$ , rejection becomes almost automatic. This pattern reflects both misspecification in share space that drives  $\hat{\tau}$  to the boundary and the irregular asymptotics of Wald tests for a variance parameter on the boundary (Stram and Lee 1994). Importantly, the coexistence of high apparent power with near-zero coverage implies that OLS rejections do not indicate reliable detection of slope heterogeneity, but rather artifacts of boundary estimation in a misspecified model.

4.5. Summary

The Monte Carlo evidence is clear. Across BPI values, ILR-ML recovers both the mean proportionality slope and slope heterogeneity, with unbiased estimates, low RMSE, near-nominal coverage, and high power. OLS estimators, by contrast, show systematic bias, higher RMSE, confidence intervals that rarely contain the true slope, and low power to reject overly proportional allocations; in multilevel form, they also fail to recover the variance component  $\tau$ . These results indicate that, among the approaches considered, ILR-ML provides reliable inference for compositional outcomes on the simplex.

5. Refining Gamson’s Law: Portfolio Allocation in the ILR Framework

Having demonstrated the advantages of the ILR transformation in simulations, we now turn to real-world data to evaluate Gamson’s Law. We draw on the dataset from our previous study (Martin and Vanberg 2020), which covers 910 government parties across 308 governments in 16 parliamentary

<sup>15</sup>See Stram and Lee (1994) on boundary issues in variance component estimation.

**Table 1.** Effect of party seat share on share of ministerial posts (single-level OLS–LOO).

Explanatory variables	Estimates
Seat share	0.790 (0.011)
Intercept	0.068 (0.004)
$R^2$	0.89

Note: Displayed are OLS coefficient estimates, with conventional standard errors in parentheses. Dependent variable: Share of cabinet ministries. *N*: 602 coalition parties (after 308 parties, one per government, are randomly dropped from the analysis).

**Table 2.** Effect of party seat share on share of ministerial posts (multilevel OLS–LOO).

Explanatory variables	Estimates
Seat share	0.790 (0.012)
Intercept	0.068 (0.004)
$\sigma^2$ (Seat share)	0.001 (0.001)
$\sigma^2$ (Residual)	0.004 (0.000)
Log-likelihood	783.74

Note: Displayed are maximum likelihood coefficient estimates from a multi-level linear regression, with conventional standard errors in parentheses. The random slope is estimated at the level of the government, and is assumed to be normally distributed. Dependent variable: Share of cabinet ministries. *N*: 602 coalition parties (after 308 parties, one per government, are randomly dropped from the analysis). The estimated variance of the seat-share slope is effectively zero, and a likelihood-ratio test confirms no improvement over a fixed-slope model ( $\chi^2 = 1.16$ ,  $p > 0.14$ ).

democracies.<sup>16</sup> Before presenting the ILR-based analysis, we replicate our earlier findings in Tables 1 and 2.<sup>17</sup>

In Table 1, we show the results from the standard (single-level) OLS–LOO model. The estimates confirm what most researchers of portfolio allocation and Gamson's Law will recognize: an intercept near 0.07, a coefficient on *Seat Share* close to 0.80, and an  $R^2$  of approximately 0.90.<sup>18</sup> These results indicate that the relationship between seat share and portfolio share falls short of exact 1:1 proportionality: the intercept differs significantly from zero, and the slope is statistically less than one. Substantively, they suggest a modest numerical “bonus” for small parties and a corresponding “penalty” for large ones—

<sup>16</sup>The countries and years covered are detailed in Appendix 4 of the Supplementary Material.

<sup>17</sup>We also estimated models with cluster-robust standard errors by government, but the results are virtually indistinguishable from those with conventional errors, and so we report only the latter.

<sup>18</sup>See, e.g., the study by Warwick and Druckman (2006, Table 2, Model 1), which reported an intercept of 0.069, a coefficient of 0.793, and an  $R^2$  of 0.89.

an asymmetry long recognized in the coalition literature (Browne and Franklin 1973; Warwick and Druckman 2006).

We also estimated in our previous study a multilevel (mixed-effects) version of the OLS–LOO model (OLS–ML), which allows the effect of *Seat Share* to vary across governments. Those findings are reported in Table 2.<sup>19</sup> As the results show, the estimated variance of the seat-share slope is close to zero, the likelihood-ratio test fails to reject the null of no heterogeneity ( $p > 0.14$ ), and the point estimates are nearly identical to those from the single-level model. In short, the multilevel analysis provides no evidence of variation in the effect of seat share across governments, indicating that the simpler single-level specification in Table 1 fits the data equally well.

But as our simulations make clear, there are strong reasons to doubt the reliability of either specification. Both the single-level and multilevel OLS approaches are structurally misspecified for compositional outcomes: they impose linearity on shares that are bounded and subject to a constant-sum constraint, producing biased point estimates and miscalibrated uncertainty. They also generate predictions that fail to respect compositional coherence.<sup>20</sup> These deficiencies cast doubt on the apparent stability of the seat–portfolio relationship. What looks like a uniform, law-like association may instead be an artifact of model misspecification. In the analysis that follows, we re-examine the question using the ILR framework, which respects the geometry of compositional data and provides a sounder basis for testing Gamson’s Law.

Because our goal is to examine not only the average relationship between seats and portfolios but also the possibility that this effect varies across governments, we estimate a multilevel ILR model, reported in Table 3.<sup>21</sup> The dependent variable is the *ILR-Transformed Portfolio Share* received by a government party, with the central independent variable defined analogously as the *ILR-Transformed Seat Share* (*ILR-TSS*). Because seat shares are themselves compositional, the ILR transformation ensures that changes in a party’s legislative strength are interpreted in *relative* rather than absolute terms. A one-unit increase in the ILR-TSS represents a proportional gain in that party’s portfolio share relative to the geometric mean of its coalition partners’ shares, not an increase holding others constant.<sup>22</sup>

To construct these ILR coordinates, we apply an SBP strategy that begins with the smallest party in each government and proceeds in ascending order of seat share. For each party—except the largest—we compare its seat or portfolio share to the geometric mean of the remaining, larger parties, thereby generating a series of balances that contrast smaller parties with blocs of larger coalition partners. The ILR transformation is carried out separately within each government, using only the parties actually present and ordering them consistently by size. This SBP choice is made solely for interpretive clarity; as shown in Appendix 1 of the Supplementary Material, alternative SBPs are linked by orthogonal rotations and yield identical model fit and predictions when mapped back to the simplex.

<sup>19</sup>In Martin and Vanberg (2020), we estimated a specification with both a random slope and a random intercept with unstructured covariance (see Appendix C of that study). Repeating that analysis, we see that a likelihood-ratio test favors the simpler random-slope specification ( $LR \chi^2(2) = 5.21, p > 0.07$ ); AIC values likewise indicate no improvement from adding a random intercept. We therefore present the more parsimonious model here.

<sup>20</sup>While no predictions from the two models above fall outside the  $(0, 1)$  range for any individual party, the predicted shares across all parties in a government rarely sum to exactly one. In 46% of governments, the total predicted portfolio shares fall below 0.95; in another 31%, they exceed 1.05. One might attempt to correct this by normalizing predictions post-estimation (e.g., rescaling them proportionally so they sum to one). But this approach is problematic: the normalization factor varies across governments, effectively implying different slopes and intercepts from case to case, and it decouples the adjusted predictions from the estimated coefficients and their standard errors.

<sup>21</sup>For parsimony, the model includes a random slope for seat share but no random intercept for governments. In an ancillary specification (not shown), we add a random intercept, but an AIC comparison with the model in Table 3 reveals no improvement in fit. Following recommended practices in the multilevel modeling literature (Bates *et al.* 2018), we therefore exclude it.

<sup>22</sup>This interpretation parallels Adolph’s (2013) discussion of compositional covariates, where changes must be understood as ratio-preserving counterfactuals. The ILR formulation satisfies this requirement by expressing the predictor as a log-ratio contrast between one component of the composition and the geometric mean of the remaining parts.

**Table 3.** Effect of party seat share on share of ministerial posts (multilevel ILR).

Explanatory variables	Estimates
ILR-transformed seat share (ILR-TSS)	0.712 (0.024)
Intercept	-0.001 (0.026)
$\sigma^2$ (ILR-TSS)	0.017 (0.005)
$\sigma^2$ (Residual)	0.132 (0.009)
Log-likelihood	-289.22

*Note:* Displayed are maximum likelihood coefficient estimates from a multilevel linear regression, with conventional standard errors in parentheses. The random slope is estimated at the level of the government, and is assumed to be normally distributed. Dependent variable: Isometric log-ratio (ILR) transformation of numerical share of cabinet ministries. *N*: 602 balances, reflecting 910 government coalition parties, in 308 governments. A likelihood ratio test of the multilevel ILR model against a single-level ILR alternative leads us to reject the null hypothesis of no difference between them ( $\chi^2 = 21.82, p < 0.001$ ).

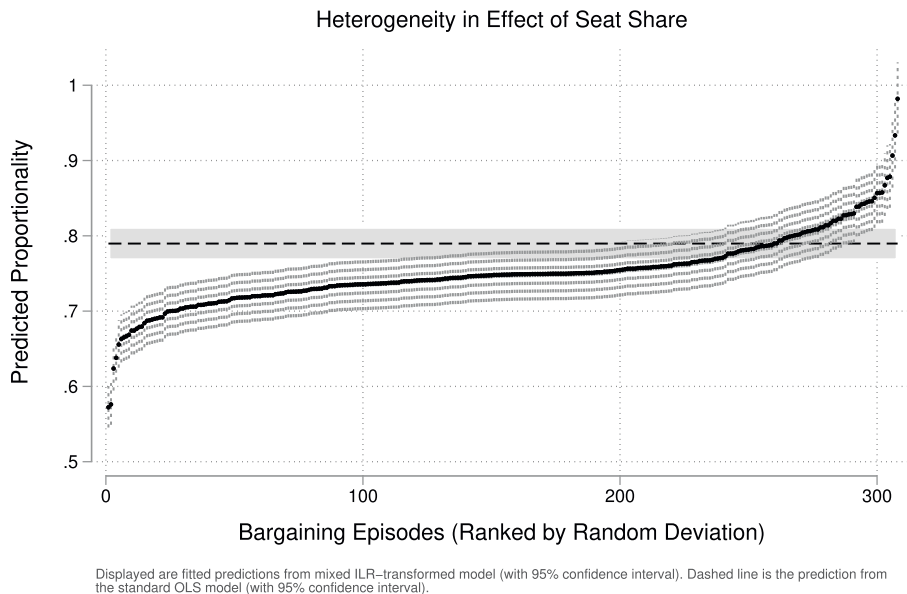
Because ILR-transformed variables represent log-ratio balances rather than individual proportions, the parameter estimates in Table 3 are not directly interpretable as effects on a single party's portfolio share. ILR coordinates reflect the log-ratio of one party's value to the geometric mean of the remaining components in the composition; and under our partitioning strategy, each balance compares a party's seat share to the geometric mean of the seat shares of all larger coalition partners in the same government. Accordingly, a positive coefficient on ILR-TSS implies that as a party's seat share increases relative to this geometric mean, its expected portfolio share also increases.

In this context, the model intercept represents the expected portfolio balance when the logged seat-share balance is zero—that is, when a party's seat share is exactly equal to the geometric mean of the seat shares of its larger partners. The estimated intercept is close to zero and not statistically distinguishable from it ( $p \approx 0.98$ ). Substantively, this suggests that when a party is perfectly balanced in legislative size relative to its larger partners, it receives a proportionate share of cabinet portfolios. There is no systematic upward or downward shift at this point of balance.<sup>23</sup>

Turning to the average slope, the coefficient on ILR-TSS is positive and highly significant, but its magnitude—0.71—is well below one, indicating that the relationship between seat share and portfolio share is less than perfectly proportional. To facilitate comparison with the OLS estimate in share space, we approximate the corresponding proportionality by taking the ratio of exponentiated coefficients,  $e^{0.71}/e^1 \approx 0.75$ . This suggests that a proportional 4% increase in seat share corresponds, on average, to only about a 3% increase in portfolio share. The effect is notably smaller—and statistically different ( $p < 0.001$ )—than the 5:4 proportionality estimated using OLS on untransformed data. In short, while the association between seats and portfolios remains strong, it is weaker than previously thought.

This conclusion also revises our earlier finding in Martin and Vanberg (2020)—based on an OLS model using untransformed shares—that the effect of seat share on portfolio share is homogeneous across governments. That earlier result now appears to reflect a limitation of the OLS framework rather

<sup>23</sup>Importantly, this does not imply the absence of a small-party bonus. That disproportionality is captured in the slope. A slope of one would imply strict proportionality: as a party's seat share increases relative to its partners, its portfolio share increases in lockstep. A slope less than one implies diminishing returns, favoring smaller parties. For example, if a party has one-third the seat share of its larger partners, its ILR-transformed balance is  $\log(1/3) \approx -1.10$ . A slope of 0.71 yields an expected portfolio balance of -0.78, which (taking the inverse log) back-transforms to a portfolio share about 0.46 times that of its partners—substantially greater than its relative seat share of 0.33.



**Figure 6.** Heterogeneity in effect of seat share.

than a feature of the underlying data. As shown in our Monte Carlo simulations, linear models applied to raw compositional outcomes impose a misspecified error structure that can mask true heterogeneity, either through biased point estimates or distorted uncertainty estimates. By contrast, modeling the data in ILR space respects its compositional structure and reveals variation in the seat-to-portfolio relationship that is substantively meaningful but previously hidden.

The multilevel ILR model indeed uncovers such heterogeneity. The variance component of ILR-TSS indicates that the strength of the seat-to-portfolio relationship varies across governments. Using best linear unbiased predictions (BLUPs) of the random slope deviations, we calculate predicted seat share coefficients for each of the 308 governments. Figure 6 displays these coefficients (with 95% confidence intervals), ordered by magnitude. To facilitate direct comparison with the fixed coefficient estimate from the conventional OLS model (Table 1), we again transform the ILR-based slope estimates back to the original share metric by exponentiating each value and dividing by the exponentiated value of one (corresponding to perfect proportionality). This places them on the same scale as the OLS coefficient and allows us to assess how proportionality varies across governments. As a benchmark, the figure overlays the OLS slope and its confidence interval. While this OLS estimate treats the slope as uniform across cases, the ILR estimates show otherwise: 148 of the 308 government-specific slopes fall significantly below the OLS estimate ( $p < 0.05$ ), while 17 fall significantly above it. The wide dispersion underscores the value of the ILR-based multilevel approach in capturing variation that conventional methods obscure.<sup>24</sup>

Beyond uncovering heterogeneity, the ILR model also fits the data far better than the conventional OLS model—a point made clear by comparing residual patterns between the two. If both models effectively capture the underlying dynamics, we should observe no systematic relationship between residuals and the (untransformed) variable of interest, *Portfolio Share*. To assess this, we calculate

<sup>24</sup> Appendix 5 of the Supplementary Material graphs the estimated heterogeneity in the effect of seat share by country and government. Mirroring the results of earlier country-level tests of Gamson's Law (e.g., Schofield and Laver 1985), we find that the degree of proportionality varies across countries. Our multilevel approach, however, goes further by showing that there is significant variation not only between countries but also across governments within the same country. This pattern suggests that differences in portfolio allocation are driven less by stable institutional features and more by contextual factors specific to each episode of coalition bargaining.

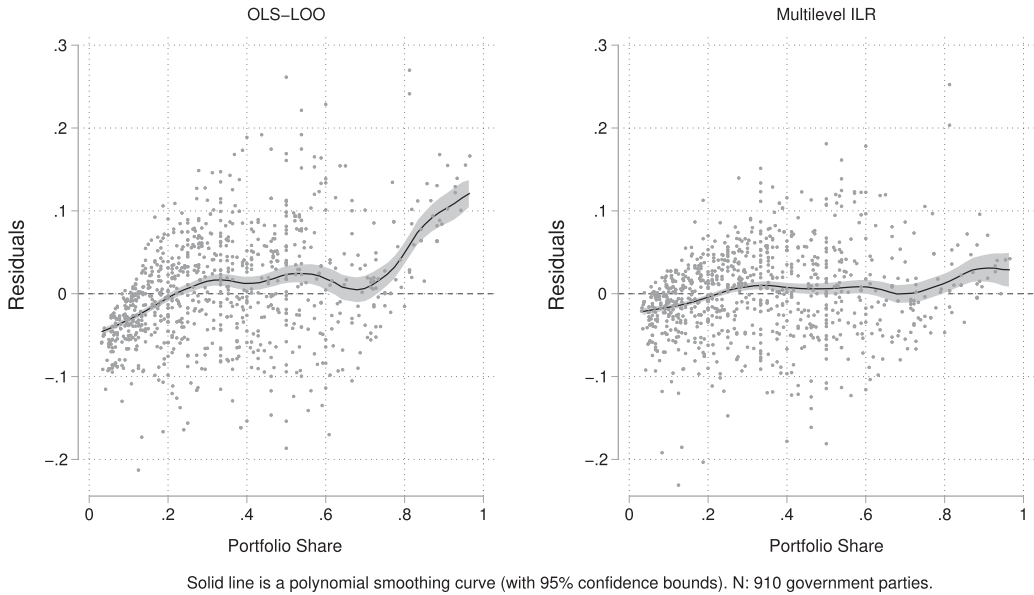


Figure 7. Residual patterns: Conventional OLS versus multilevel ILR.

predicted values from each model. For OLS, this is straightforward. For the ILR model, predictions are made in log-ratio space and then back-transformed to the compositional scale.<sup>25</sup> Residuals are then calculated as the difference between observed and predicted portfolio shares.

Figure 7 provides a visual comparison of residuals from the conventional OLS model and the ILR-based multilevel model. The contrast is striking. In the OLS model, residuals exhibit pronounced curvature—the mean reflected in the polynomial smoothing curve is rarely centered at zero—with systematic overprediction at the low end of the distribution and underprediction at the high end.<sup>26</sup> This pattern reflects the inherent limitations of applying linear regression to compositional outcomes: the model imposes a linear structure on data constrained to a simplex, leading to distortion—especially near the boundaries. By contrast, residuals from the ILR model are much more evenly distributed around zero, with little curvature and substantially reduced magnitude. While some misprediction persists at the extremes, it is markedly diminished—particularly at the upper end of the distribution, where the OLS model performs worst. The visual contrast is matched by a substantial difference in fit: the residual sum of squares in the OLS model is 56% higher than in the ILR model. Together, these results reinforce the argument that OLS imposes a misspecified structure on portfolio data, obscuring underlying patterns. The ILR framework not only reveals heterogeneity in effects but also yields more accurate and statistically coherent estimates across the board.

## 6. Discussion and Conclusion

Our analyses indicate that the apparent stability of the seat–portfolio link rests on methods not appropriate for modeling compositional outcomes. Because portfolio and seat shares are parts of a whole, their interdependence and bounded support violate key assumptions of conventional linear

<sup>25</sup>See Appendix 6 of the Supplementary Material for the ILR back-transformation procedure.

<sup>26</sup>At first glance, the OLS fit line in Figure 7 appears lower for small parties, which might suggest that OLS predicts a larger small-party bonus than the ILR model. This impression is local: OLS overpredicts portfolios for small parties near the lower tail, but its overall slope is steeper than the ILR slope. In aggregate, OLS therefore implies stronger proportionality—and thus a smaller small-party advantage—than the ILR model, which better captures the curvature of the relationship in the simplex.

models. The Monte Carlo evidence shows these violations are not benign: across a wide range of conditions, OLS estimators produce systematic bias, low power, and sharply degraded coverage once allocations move away from the simplex center. Even in multilevel form, OLS struggles to recover genuine random-slope variation. In short, part of what earlier studies treated as a “law-like” regularity reflects model misspecification rather than true stability. By contrast, the ILR transformation estimated in a multilevel framework aligns the model with the geometry of the data. In the simulations, ILR–ML accurately recovers the mean proportionality parameter, with high power and near-nominal coverage throughout the simplex. It also identifies random-slope variance that OLS largely obscures.

Applied to real-world coalition data, the ILR model revises the standard picture. The average relationship between seat share and portfolio share is weaker than the canonical 5:4 estimate, closer to 4:3. This implies that small parties receive larger portfolio rewards relative to their legislative strength than previously recognized. Just as important, the effect of seat share is not constant across governments: ILR–ML uncovers wide dispersion in slopes, indicating that proportionality is contingent on features of specific bargaining episodes. Some governments approach near-perfect proportionality, while others deviate sharply from it. This pattern accords with our residual comparison, where the ILR model eliminates much of the curvature and magnitude seen under OLS. Substantively, these findings suggest that Gamson’s “law” is better viewed as a central tendency that strengthens or weakens with context. The ILR framework provides a way to investigate those contexts directly.

Methodologically, the implications extend well beyond portfolio allocation. Many political outcomes are compositional and differ in which components are observed across cases: party vote shares when not all parties compete in every district, media attention when the set of salient policy issues changes over time or across countries, or public spending shares when budget categories are redefined or reorganized. In such settings, ALR+SUR is often *infeasible* because it requires a fixed set of components and a stable reference part across all observations. The ILR transformation avoids these constraints by adapting to the observed composition within each unit while preserving the geometry of the simplex. As researchers model outcomes that represent divisions of finite resources, ILR offers a practical way to reduce bias and achieve more reliable inference.

**Acknowledgements.** We wish to thank participants in the Dynamic Pie workshop (Texas A&M University) for helpful comments on a previous version of the manuscript.

**Funding statement.** This research received no particular funding.

**Data Availability Statement.** Replication code for this article has been published in the Political Analysis Harvard Dataverse at <https://doi.org/10.7910/DVN/MZAALH> (Martin 2025).

**Competing Interests Statement.** The authors declare no competing interests.

**Supplementary Material.** For supplementary material accompanying this paper, please visit <https://doi.org/10.1017/pan.2026.10036>.

## References

- Adolph, C. 2013. *Bankers, Bureaucrats, and Central Bank Politics: The Myth of Neutrality*. Cambridge: Cambridge University Press.
- Aitchison, J. 1982. “The Statistical Analysis of Compositional Data.” *Journal of the Royal Statistical Society: Series B (Methodological)* 44 (2): 139–177.
- Aitchison, J. 1986. *The Statistical Analysis of Compositional Data*. Caldwell, NJ: Blackburn Press.
- Bates, D., R. Kliegl, S. Vasishth, and H. Baayen. 2018. “Parsimonious Mixed Models.” arXiv (Statistics–Methodology) preprint [arXiv:1506.04967](https://arxiv.org/abs/1506.04967). Accessed 25 April 2023.
- Browne, E. C., and M. N. Franklin. 1973. “Aspects of Coalition Payoffs in European Parliamentary Democracies.” *American Political Science Review* 67 (2): 453–469.
- Carroll, R., and G. W. Cox. 2007. “The Logic of Gamson’s Law: Pre-Election Coalitions and Portfolio Allocations.” *American Journal of Political Science* 51 (2): 300–313.
- Clayton, A., and P. Zetterberg. 2018. “Quota Shocks: Electoral Gender Quotas and Government Spending Priorities Worldwide.” *Journal of Politics* 80 (3): 916–932.

- Cox, G. W. 2021. "Nonunitary Parties, Government Formation, and Gamson's Law." *American Political Science Review* 115 (3): 917–930.
- Cutler, J., S. de Marchi, M. Gallop, F. M. Hollenbach, M. Laver, and M. Orłowski. 2016. "Cabinet Formation and Portfolio Distribution in European Multiparty Systems." *British Journal of Political Science* 46 (1): 31–43.
- Egozcue, J. J., V. Pawłowsky-Glahn, G. Mateu-Figueras, and C. Barceló-Vidal. 2003. "Isometric Logratio Transformations for Compositional Data Analysis." *Mathematical Geology* 35 (3): 279–300.
- Fréchette, G. R., J. H. Kagel, and M. Morelli. 2005. "Gamson's Law Versus Non-cooperative Bargaining Theory." *Games and Economic Behavior* 51 (2): 365–390.
- Gamson, W. A. 1961. "A Theory of Coalition Formation." *American Sociological Review* 26 (3): 373–382.
- Honaker, J., J. N. Katz, and G. King. 2002. "A Fast, Easy, and Efficient Estimator for Multiparty Electoral Data." *Political Analysis* 10 (1): 84–100.
- Hopkins, V., A. Kagalwala, A. Q. Philips, M. Pickup, and G. D. Whitten. 2024. "How Do we Know What we Know? Learning from Monte Carlo Simulations." *Journal of Politics* 86 (1): 36–53.
- Indridason, I. H. 2015. Live for Today, Hope for Tomorrow? Rethinking Gamson's Law. Working paper: University of California, Riverside. <https://indridason.politicaldata.org/workingpapers/Live%20for%20today.pdf>. Accessed 21 July 2023.
- Jackson, J. E. 2002. "A Seemingly Unrelated Regression Model for Analyzing Multiparty Elections." *Political Analysis* 10 (1): 49–65.
- Katz, J. N., and G. King. 1999. "A Statistical Model for Multiparty Electoral Data." *American Political Science Review* 93 (1): 15–32.
- Lehmann, E. L., and J. P. Romano. 2005. *Testing Statistical Hypotheses*. 3rd ed. New York: Springer.
- Martin, L. W. 2025. "Replication Data for: Refining Gamson: The Isometric Log-Ratio Transformation and Portfolio Proportionality in Multiparty Governments." Harvard Dataverse. <https://doi.org/10.7910/DVN/MZAALH>. Harvard Dataverse.
- Martin, L. W., and G. Vanberg. 2020. "What you See Is Not Always What you Get: Bargaining before an Audience under Multiparty Government." *American Political Science Review* 114 (4): 1138–1154.
- Pawłowsky-Glahn, V., and J. J. Egozcue. 2006. "Compositional Data and their Analysis: An Introduction." *Geological Society of London Special Publications* 264 (1): 1–10.
- Pawłowsky-Glahn, V., J. J. Egozcue, and R. Tolosana-Delgado. 2015. *Modeling and Analysis of Compositional Data*. Chichester: John Wiley & Sons.
- Philips, A. Q., A. Rutherford, and G. D. Whitten. 2016. "Dynamic Pie: A Strategy for Modeling Trade-Offs in Compositional Variables over Time." *American Journal of Political Science* 60 (1): 268–283.
- Schofield, N., and M. Laver. 1985. "Bargaining Theory and Portfolio Payoffs in European Coalition Governments 1945–83." *British Journal of Political Science* 15 (2): 143–164.
- Stram, D. O., and J. W. Lee. 1994. "Variance Components Testing in the Longitudinal Mixed Effects Model." *Biometrics* 50 (4): 1171–1177.
- Tomz, M., J. A. Tucker, and J. Wittenberg. 2002. "An Easy and Accurate Regression Model for Multiparty Electoral Data." *Political Analysis* 10 (1): 66–83.
- Warwick, P. V., and J. N. Druckman. 2006. "The Portfolio Allocation Paradox: An Investigation into the Nature of a Very Strong but Puzzling Relationship." *European Journal of Political Research* 45 (4): 635–665.
- Zellner, A. 1962. "An Efficient Method of Estimating Seemingly Unrelated Regressions and Tests for Aggregation Bias." *Journal of the American Statistical Association* 57 (298): 348–368.