Measuring portfolio salience using the Bradley–Terry model: An illustration with data from Brazil

Research and Politics
January–March 2019: 1–8
© The Author(s) 2019
DOI: 10.1177/2053168019832089
journals.sagepub.com/home/rap

Cesar Zucco Jr.¹ , Mariana Batista²
and Timothy J. Power³

Abstract

How do political actors value different portfolios? We propose a new approach to measuring portfolio salience by analysing paired comparisons using the Bradley–Terry model. Paired-comparison data are easy to collect using surveys that are user-friendly, rapid, and inexpensive. We implement the approach with serving legislators in Brazil, a particularly difficult case to assess portfolio salience due to the large number of cabinet positions. Our estimates of portfolio values are robust to variations in implementation of the method. Legislators and academics have broadly similar views of the relative worth of cabinet posts. Respondent valuations of portfolios deviate considerably from what would be predicted by objective measures such as budget, policy influence, and opportunities for patronage. Substantively, we show that portfolio salience varies greatly and affects the calculation of formateur advantage and coalescence/proportionality rule measures.

Keywords

Portfolio salience, coalitional presidentialism, measurement

Introduction

Multiparty coalition formation is a longstanding theme in the study of parliamentary democracy and increasingly central in comparative presidentialism. Yet, not all cabinet-status positions are equally important. The long-running research program on “portfolio salience” seeks to assess the relative value of payoffs in cabinet formation.

Measuring portfolio salience is critical to interrogating rival hypotheses about formateur advantage versus proportionality in cabinet formation. While the former predicts that the coalition member recognized as the formateur would attempt to extract a disproportionate payoff (Rubinstein, 1982), empirical analyses tend to support “Gamson’s Law” of proportionality (Browne and Franklin, 1973; Gamson, 1961). This is a classic debate in the parliamentary systems literature (Warwick and Druckman, 2001), but it is especially relevant to presidentialism, wherein the formateur is selected exogenously and ex ante. If a “formateur bonus” is up for grabs, it could imply presidential-party retention of the most valuable cabinet positions. Accounting for portfolio salience might, therefore, reveal a reduction of proportionality (or “coalescence”) in the allocation to coalition partners.

Salience is fundamental to cabinet politics, but its measurement is vexing. It is not obvious whether one should use reputational prestige, policy outputs, or institutional size of the ministry in question. Some portfolios responsible for core policy decisions actually have modest budgets or few personnel, complicating comparisons based on observed characteristics of ministries (Batista, 2017). The revealed preferences of parties may be important, but using this information alone is complicated considering that parties may not always be programmatic or value cabinet positions for policy reasons (Adachi and Watanabe, 2007). Thus, in the wake of seminal work by Laver and Hunt (1992),

¹Brazilian School of Business and Public Administration (EBAPE), Getúlio Vargas Foundation, Rio de Janeiro, RJ, Brazil

²Department of Political Science, Federal University of Pernambuco, Centro de Filosofia e Ciências Humanas, Recife, PE, Brazil

³Department of Politics and International Relations, University of Oxford, St Antony’s College, Oxford, UK

Corresponding author:

Cesar Zucco Jr., Brazilian School of Business and Public Administration (EBAPE), Getúlio Vargas Foundation, Rua Jornalista Orlando Dantas 30, Sala 240, 22231-010, Rio de Janeiro, RJ, Brazil.

Email: cesar.zucco@fgv.br

several variations on expert surveys have been used to gauge portfolio salience. These surveys can be very demanding for respondents, and having experts rate all the extant ministries is often inefficacious. Moreover, there is need of evidence to show that expert judgement coincides with the views of political practitioners.

These earlier efforts inspire the present article, in which we propose the Bradley–Terry (BT) model to analyse pairwise comparisons of cabinet positions. The basic BT model has been applied in many fields, from psychology (to recover “ability” scales from observed pairwise comparisons) to sports journalism (as in the college football rankings popular in the US [New York Times 1979]). Respondents need only choose the more valuable portfolio from a few different pairs of options presented to them.

We test this simplified approach on new data from original surveys of Brazilian politicians and experts. As a paradigmatic case of “coalitional presidentialism,” Brazil has spawned the largest literature on coalition governments outside of parliamentary democracies (Chaisty et al., 2018). The country boasts a large, heterogeneous set of cabinet-level posts, which would seem to make it impractical to ask respondents to rate all of them—an obstacle that, we show, can be overcome by the BT model.

In sum, our contribution is to apply an established method to an elusive problem. This novel application allows for greatly simplified data collection that helps advance our substantive understanding of portfolio allocation via improved measurement of the formateur advantage (Altman, 2000) and of the proportionality of cabinet shares (Amorim Neto, 2002).

Measuring portfolio salience using the BT model

The BT model, originally proposed in Bradley and Terry (1952), is based on contests between alternatives. Decisions may be based on subjective evaluations of properties of the objects made by referees (e.g. taste) or on objective outcomes, such as the score of a football game. In political science, cognate approaches have appeared sporadically, but direct applications of BT remain rare.¹

For each matchup between two cabinet portfolios, our subjects were asked to choose the more valuable one. In our application, the “items” being rated are “ministries” (M), the referees making the decisions are either legislators or experts, and the underlying scale being recovered is referred to as the “worth” or “value” of each item. We are not interested in examining the impact of specific attributes of each item on this choice; our sole goal is to retrieve the underlying relative value for each item.

The BT model is ideally suited for this because it is built on the probability that one item M_i is chosen over another M_j in a pairwise contest π_{ij} . This probability can then be

expressed as a function of the underlying worth of each item (α_i), such that:

$$\pi_{ij} = \frac{\alpha_i}{\alpha_i + \alpha_j} \quad (1)$$

In its original formulation, this probability of selection is then expressed in terms of exponential functions $\exp(\lambda_i) \equiv \alpha_i$ of the underlying worth, and by employing a logit link function. Thus:

$$\text{logit}(\pi_{ij}) = \lambda_i - \lambda_j \quad (2)$$

In short, λ_i is a transformation of the underlying worth α_i , which determines the probability of one item being chosen over another in pairwise contests, as defined in equation (1).

The BT model employs a logit transformation,² but alternative functions have been used. As others have noted (Luce et al., 1990), this specification meets the strong stochastic transitivity condition, such that when $\pi_{ij} \geq 1/2$ and $\pi_{j,k} \geq 1/2$, then $\pi_{i,k} \geq \max \pi_{ij}, \pi_{jk}$. This is important because items that are not directly compared in any contest can be indirectly compared, as long as the design is connected (that is, no subset of the items exists that has not been compared to another subset of the data) (Bradley, 1984). Additionally, as no subgroup of items exists in which its members always win contests against another subgroup of items, the parameters in this model can be estimated by maximum likelihood with either a sum constraint ($\sum_{i=1}^n \lambda_i = 0$) or, more commonly, a reference object constraint ($\lambda_i = 0$ for one $i \in \{1, \dots, n\}$). For the remaining parameters, standard errors are computed as usual. Items defeating rivals with larger α receive a bigger boost to their own α because information from one contest is implicitly used in other contests. In our design, due to the large number of items being rated, not all possible matchups are observed, implying missing data. By the construction of our survey, these data are missing at random, so they should not affect point estimates directly.

This simple specification we employ is referred to as the “unstructured” BT model, but numerous extensions have been proposed, including allowing for rankings of more than two options (Böckenholt, 2001). These extensions are unnecessary to us, as our goal was simply to retrieve the underlying relative valuations of portfolios, keeping the survey instrument as simple as possible.

The data

We fielded two surveys simultaneously in Brazil in 2017, one with 142 legislators in the National Congress (the

“elite” survey) and one with experts. In both, we presented pairs of ministries that existed in 2014 and 2015, the last two full years of president Dilma Rousseff’s administration.³ A total of 37 ministries were rated,⁴ which implies the existence of 666 pairwise comparisons. Since it was not feasible to have all participants review all possible matchups, we randomized the items presented to each respondent.

The expert survey was conducted online with 283 members of the Brazilian Political Science Association (ABCP). It solicited demographic data and asked whether respondents considered themselves “experts” in executive–legislative relations or other subfields. In pairwise comparisons, respondents were asked to choose the ministry they thought a typical politician would prefer to obtain for his or her party.⁵ Each respondent answered an initial battery of eight comparisons, and most opted to review three additional pairs. Item non-response was minimal, so we treat it as completely random.

Because the BT model uses the transitivity assumption to leverage information from multiple comparisons, we undertook to minimize the number of “lopsided” matchups. Thus, we stratified the portfolios into four tiers based on our own *a priori* expectations about their reputational value. In eight of the 11 matchups, the two competing portfolios were drawn from a subset of two adjacent tiers, and six of these involved the highest-valued tier. Thus, while portfolios were still randomly drawn, there are more “close” comparisons in our dataset than would have obtained with simple randomization, and more contests involving high-profile ministries. The expert survey generated data from 3101 comparisons; 275 of the possible pairwise 666 contests were never observed, while four contests were observed more than 20 times.

The elite survey was embedded in the Brazilian Legislative Survey (BLS), so space was limited to four pairwise comparisons per respondent. The BLS was implemented using both an online version and a paper questionnaire. In the electronic version, we implemented a randomization scheme almost identical to the one described for the expert survey. On the hard-copy instrument, we created five different sets of four randomly drawn pairwise comparisons, and these were randomly distributed to parliamentarians.⁶ Legislators were asked which ministry they would prefer their party held. Some 62 responded online and 80 on paper. The elite data encompass 560 pairwise comparisons, with 221 different matchups occurring at least once.

In both datasets, all 37 portfolios were rated (at least 102 times by experts and at least seven times by practitioners). Items that appear more often have smaller confidence intervals, and estimates from the expert dataset are more precise than those from the elite dataset. Substantively, the elite survey is our preferred source of data, as it reflects the perceptions of individuals who operate directly in the arena of government formation. However, below we present both

surveys, not only to gauge robustness and to check the convergence between elite and expert valuation of portfolios, but also to assess the feasibility of applying the BT method to this and other political science measurement tasks.

Survey instruments with pairwise comparisons are impressively user-friendly. Several ABCP colleagues reported to us that the survey was “fun,” which may explain why 273 of the 278 participants opted to continue on past the first eight pairwise comparisons. Median time to completion for the entire expert survey was just over 4 min.

The values of ministerial positions

Our main results were generated by applying the BT model to data from the elite survey. In Table 1, cabinet-level positions are ranked from highest to lowest according to the point estimates of the value of each portfolio. The first column reports our pre-stratification of ministries into four tiers. The following columns report our estimates of α_i , which are provided relative to the baseline category (set *a posteriori* to be the highest-ranked ministry), whose value is arbitrarily fixed at zero. We also present the statistical significance of differences in value between each portfolio and this baseline.

Our initial intuitions were not far off from legislators’ opinions. However, there were some important differences, as exemplified by the *Ministério das Cidades* (Urban Affairs). This portfolio was assigned to the second tier by us but was viewed by legislators as the most valuable ministry. Interestingly, among the top-ranked portfolios, we observe two that are important sources of pork and patronage (Urban Affairs and Education), two that are ostensibly “core policy” ministries with few material goods to distribute but considerable policy influence (Finance and Planning), and one that is the quintessential political coordination portfolio (*Casa Civil*, the presidential chief of staff). Salience incorporates various influences not easily captured by a common set of observable characteristics, thus reinforcing the importance of a perceptual measure such as the one we propose.

However, it is also the case that the estimates are noisy. The top-ranked portfolio, for instance, cannot be statistically distinguished from the subsequent five. It is not feasible to report *p*-values for all pairwise comparisons, so we developed and reported bootstrapped confidence intervals for the rank-order of each portfolio.⁷ Uncertainty, however, can be resolved with slightly larger datasets. Consider the rightmost columns in Table 1, which report estimates obtained with data from the larger expert survey. The group of ministries that cannot be distinguished from the most valuable one is now reduced to three, and the average standard error is 0.26 compared to 1.13 in the elite survey. Similarly, the maximum width of the rank-order confidence interval is 9 (compared to 27 in the elite survey).

Table 1. Estimated values of portfolios, elite and expert surveys.

Tier	Portfolio	Legislators		Experts	
		Worth	Rank CI	Worth	Rank CI
		(α)	Low–high	(α)	Low–high
2	Urban Affairs (<i>Cidades</i>)	0.00	1–6	–0.45*	2–9
1	Planning	–0.15	1–5	–0.22	1–6
1	Finance (<i>Fazenda</i>)	–0.24	1–5	0.00	1–6
1	Presidential chief of staff (<i>Casa Civil</i>)	–0.61	1–10	–0.03	1–4
1	Education	–1.25	2–8	–0.37*	3–8
1	Mines and Energy	–1.40	3–11	–0.38*	3–8
1	Health	–1.69	3–14	–0.01	1–4
2	Science and Technology	–2.17**	5–17	–2.14***	18–25
1	National Integration (Regional Dev.)	–2.17**	6–18	–1.09***	9–15
1	Agriculture	–2.17**	5–18	–0.99***	8–14
1	Labor	–2.33**	4–22	–1.44***	12–19
1	Foreign Affairs	–2.40**	6–21	–2.30***	19–28
1	Justice	–2.45**	7–20	–1.08***	9–15
1	Social Security and Pensions	–2.61**	7–24	–1.35***	11–18
2	Industry and Foreign Trade	–2.71***	7–21	–0.72***	5–12
1	Communications	–2.75**	7–25	–1.54***	13–19
2	Institutional Relations (Leg. Affairs)	–2.90***	10–23	–1.38***	11–18
2	Social Development	–2.90**	6–27	–1.47***	12–19
2	Sec. Gen. of Presidency (Civil Society)	–3.06***	8–26	–0.82***	6–13
4	Small Business Administration	–3.08**	4–31	–3.01***	27–34
2	Defense	–3.22***	12–26	–2.47***	21–29
4	Strategic Affairs	–3.33***	10–28	–2.19***	18–27
1	Transportation	–3.64***	13–32	–0.68***	5–11
3	Environment	–3.67***	17–28	–2.42***	20–28
2	Port Authority	–3.82***	15–32	–2.35***	19–28
4	Comptroller General	–3.92***	11–35	–2.23***	19–27
3	Agrarian Reform	–3.98***	19–30	–1.76***	15–20
3	Tourism	–4.35***	22–33	–2.80***	26–32
4	Human Rights	–4.36***	18–34	–3.39***	31–35
3	Culture	–4.69***	26–34	–2.99***	28–33
3	Sports	–4.82***	24–36	–2.54***	22–29
3	Fisheries	–5.06***	24–37	–3.05***	28–34
4	Social Communication	–5.11***	26–37	–2.49***	20–29
4	Advancement of Women	–5.20***	26–37	–4.00***	35–37
4	National Security Advisor	–5.44***	28–37	–3.41***	31–35
3	Civil Aviation Authority	–5.53***	29–37	–3.22***	30–35
4	Advancement of Racial Equality	–5.57***	28–37	–4.21***	35–37

BT estimates are in column α and the p -values of *** ($<.01$), ** ($<.05$), and * ($.10$) correspond to comparisons relative to the baseline (highest-value) portfolio, which is boldfaced. The other columns report the 90% confidence intervals (CIs) for the rank order outcomes of each portfolio, computed by a parametric bootstrap.

We expected some differences between assessments made by politicians and political scientists, as some portfolios are esteemed by elites for reasons not always apparent to external observers. For example, Urban Affairs controls one of the largest social housing programs in the world (*Minha Casa Minha Vida*). This program is important not simply because of its raw budget, but also because it liaises with over 5000 mayors and connects them to construction

firms. Mines and Energy controls many concessions and is responsible for regulation. While these functions may not appear in objective indicators of importance, they are critically important to interest groups.

Comparison of the two datasets reveals many more similarities than differences. The linear correlation coefficient between the two sets of estimates is $r = 0.87$ and the estimate of the Spearman rank correlation is $\rho = 0.88$. This

Figure 1. Comparison of elite and expert opinions.

rank correlation is only slightly weaker for the lower- and upper-bounds of the rank confidence intervals ($\rho = 0.83$ and 0.87 , respectively). The rank confidence intervals in the two sets of estimates overlap for all portfolios save two.

Interestingly, Small Business is once again the most noisily estimated cabinet-level post, and Comptroller General is not far behind. This suggests that much of the uncertainty around these estimates reflects real irresolution among respondents and is not an artifact of the model or dataset.

Figure 1 reinforces our finding of broad similarity between the two surveys. There is agreement about the most valuable ministries, although legislators discriminate within this subset more strongly than do academics. Both sets of referees also identify a sizeable set of bottom-feeding portfolios. From a universe of 37, the scatterplot identifies only four real outliers: Foreign Affairs, Small Business, Transportation, and Science and Technology. Only for the latter two portfolios do the rank confidence intervals not overlap. Expert opinion largely conforms to practitioner judgment.

Estimates versus observable measures

Does our approach generate any information we could not obtain observationally? If we could predict salience solely with objective indicators, surveys would be unnecessary. To examine this question, we compiled data on some intuitively important characteristics of Brazilian ministries in 2015, the last year in which all 37 of our portfolios coexisted.

We obtained each ministry's effective total and investment budget,⁸ as well as the total number of political appointments available within each ministry,⁹ and the total number of legislative bills authored or co-signed by each ministry in 2015 (our measure of policy influence). Table 2 shows the correlations between the expert- and elite-based BT estimates and each objective indicator. The correlations are strongest when it comes to political appointees, suggesting that a key dimension of portfolio salience is patronage potential. The correlations are slightly weaker for policy influence and, surprisingly, much weaker for the indicators of budgetary impact.

In fact, even if we combine all indicators into a linear predictor by regressing the BT estimates on all observed characteristics, the linear correlation between the predicted worth and our estimates does not rise past 0.6 (implying an R^2 of just under 0.4). This suggests that both experts and politicians have information beyond what objective

indicators can reveal, and that the desirability of portfolios depends on attributes that vary (qualitatively) across units. Therefore, even if we correctly identify a menu of characteristics that matter, they are unlikely to matter equally for all portfolios. Deductively defining a model to account for this heterogeneity is much harder than simply collecting perceptual data directly from politicians and experts. We suspect that perceptual data already incorporate this elusive cross-unit heterogeneity, and this folk wisdom then informs a series of simple binary choices made in the surveys.

Table 2. Linear associations between BT estimates and observable indicators for 2015.

Variable	Elites	Experts
Appointees	0.52***	0.54***
Policy influence	0.50***	0.43***
Investment budget	0.29*	0.38**
Log total budget	0.33**	0.36**
Predicted worth	0.60***	0.61***

Entries are correlation coefficients and respective p -values (*** < 0.01, ** < .05, and * < .10) between the BT estimates of portfolio worth (drawn from Table 1) and each of the other variables. Appointees is the raw number of discretionary DAS positions (levels 1–6) allocated to the ministry (source: Ministry of Planning); Policy Influence is the number of legislative bills signed by the ministry (source: Library of the Presidency); Total Budget and Investment Budget are the value in Brazilian reais of the authorized total and investment budget of the ministry (source: National Secretary of the Treasury). Predicted Worth scores are the predicted portfolio estimates drawn from a regression model of our BT measure of worth on the four indicators, estimated separately for each survey. See supplemental material for other definitions and sources.

Substantive implications: formateur advantage and proportionality

We return now to how these estimates can be used to inform key debates in the coalition politics literature. Do the formateur bonus and coalescence remain the same once we account for salience?¹⁰ In order to address this question, we computed two variants of monthly indicators of both variables, one unweighted version (that treats each ministry equally) and one in which we weight each ministry based on our BT point estimates.

This comparison requires a non-trivial manipulation of the data because the BT-worth parameters are estimated on an arbitrary scale. Hence, to provide them with a substantive interpretation we rescale the weights so that the most valuable portfolio is worth 10 times the least valuable one.¹¹

Figure 2 shows the extent to which weighting the cabinet positions affects our indicators of interest. The panel on the

Figure 2. Raw and weighted coalescence and formateur advantage (1995–2016).

Figure shows raw and unweighted indicators as well as the 90% confidence intervals about the raw figures. Confidence intervals were generated by a non-parametric bootstrap of the weights (see supplemental material for details).

right reports the formateur advantage. In a few short periods between 1995 and 2016, reweighting actually generates figures that are statistically different from the unweighted values.¹² Unsurprisingly, the salience-corrected values of the formateur advantage are almost always higher than the unweighted values, a fact that is likely not lost on legislators as they monitor the president's self-awarded bonus.

Weighted coalescence, in contrast, does deviate from the unweighted version significantly between 2007 and 2013 and at the tail end of the period. This suggests that salience correction does matter, which stands in contrast to results for parliamentary systems (Warwick and Druckman, 2001). Proportionality of portfolio allocation is viewed through a political lens that is absent in standard treatments of coalescence.

Conclusion

We propose a portfolio salience measure estimated by applying the BT model to pairwise comparisons of ministries undertaken by relevant informants (either experts or politicians). A major advantage of this approach is that the relevant data can be collected through simple, brief surveys that present respondents with a series of head-to-head matchups between concrete alternatives.

We show that legislators' ranking of cabinet posts deviates from what would be predicted by observable characteristics of ministries, revealing "private information" that is not otherwise easily quantifiable. We also show that experts reach very similar conclusions, which means that expert surveys implementing our approach—which tend to be easier to field than surveys with elites—might generate estimates that are "good enough" for most research purposes. In both pools of informants, the "wisdom of the crowds" seems to trump observation of objective indicators. Lastly, we show that portfolio allocation patterns can appear significantly different when adjusted for salience, which is a novelty relative to the parliamentary systems literature.

Acknowledgements

The authors are grateful to the executive council of the Brazilian Political Science Association for assistance with the expert survey. Octavio Amorim Neto provided helpful comments on an earlier draft.

Declaration of conflicting interests

The author(s) declared no potential conflicts of interest with respect to the research, authorship, and/or publication of this article.

Funding

The author(s) disclosed receipt of the following financial support for the research, authorship, and/or publication of this article: The 8th wave of the Brazilian Legislative Survey was supported by the Brazilian School of Business and Public Administration, Getúlio Vargas Foundation.

Supplemental materials

The supplemental files are available at <http://journals.sagepub.com/doi/suppl/10.1177/2053168019832089>

The replication files are available at <https://dataverse.harvard.edu/dataset.xhtml?persistentId=doi:10.7910/DVN/HJZSIM>

Notes

1. An important recent contribution is Loewen et al. (2012), which employed BT to analyze data from a survey experiment designed to assess which types of arguments are most persuasive in a referendum.
2. We use the BradleyTerry2 package for *R* (Turner and Firth, 2012).
3. Rousseff was elected president in 2010 and 2014, and impeached in mid-2016.
4. The list of 37 portfolios incorporated ministries and "secretariats" with cabinet status. We refer to all posts as ministries or portfolios.
5. See supplemental materials for question wording.
6. In the supplemental materials we show that results are unaffected by the imperfect randomization of survey items.
7. We drew 1000 simulations for each ministry's position from a multivariate normal distribution with mean and variance-covariance matrix set to their empirical values, and report each ministry's 5th and 95th percentile rank across all simulations. For the baseline ministry, which is not estimated, confidence intervals are based on pairwise *t*-tests.
8. Figures are for the liquidated budget, which can differ significantly from what was earmarked.
9. Defined as DAS positions (*Direção e Assessoramento Superior*, or High Level Directorships and Advisory Posts), levels 1 through 6, selected at the discretion of ministers.
10. We measure formateur advantage as the difference between cabinet share and legislative seat share for the president's party and the coalescence rate as $1 - 1/2 \sum_{i=1}^n (|S_i - M_i|)$ where S_i is the seat share with which the party i contributes to the coalition in the lower house and M_i is the proportion of cabinet seats allocated to the party i (Amorim Neto, 2002).
11. Druckman and Warwick (2005) and Druckman and Roberts (2008) adopted a scheme in which this ratio was 4.5 and 4.1, respectively. Given that the number of ministries is larger in Brazil and that observable differences are gigantic (e.g. the largest budget is 40,404 times the smallest budget), we chose a higher multiple.
12. See supplemental materials for details regarding the bootstrapped confidence intervals.

ORCID iD

Cesar Zucco <https://orcid.org/0000-0001-8263-9585>

Carnegie Corporation of New York Grant

This publication was made possible (in part) by a grant from the Carnegie Corporation of New York. The statements made and views expressed are solely the responsibility of the author.

References

- Adachi T and Watanabe Y (2007) Ministerial weights and government formation: Estimation using a bargaining model. *Journal of Law, Economics & Organization* 24(1): 95–119.
- Altman D (2000) The politics of coalition formation and survival in multiparty presidential democracies: The case of Uruguay, 1989–1999. *Party Politics* 6(3): 259–283.
- Amorim Neto O (2002) Presidential cabinets, electoral cycles, and coalition discipline in Brazil. In: Morgenstern S and Nacif B (eds) *Legislative Politics in Latin America*. Cambridge: Cambridge University Press, pp.48–78.
- Batista M (2017) Taking portfolio differences seriously: A composite measure based on policy, office, and budget in Brazil. *Brazilian Political Science Review* 11(1): 1–28.
- Böckenholt U (2001) Hierarchical modeling of paired comparison data. *Psychological Methods* 6(1): 49–64.
- Bradley RA (1984) Paired comparisons: Some basic procedures and examples. In: Krishnaiah PR and Sen PK (eds) *Handbook of Statistics*. Vol. 4. Amsterdam: Elsevier, pp.299–326.
- Bradley R and Terry ME (1952) Rank analysis of incomplete block designs, I. The method of paired comparisons. *Biometrika* 39(3): 324–345.
- Browne EC and Franklin MN (1973) Aspects of coalition payoffs in European parliamentary democracies. *American Political Science Review* 67(2): 453–469.
- Chaisty P, Cheeseman N and Power TJ (2018) *Coalitional Presidentialism in Comparative Perspective: Minority Presidents in Multiparty Systems*. Oxford: Oxford University Press.
- Druckman JN and Roberts A (2008) Measuring portfolio salience in Eastern European parliamentary democracies. *European Journal of Political Research* 47(1): 101–134.
- Druckman JN and Warwick PV (2005) The missing piece: Measuring portfolio salience in Western European parliamentary democracies. *European Journal of Political Research* 44(1): 17–42.
- Gamson WA (1961) A theory of coalition formation. *American Sociological Review* 373–382.
- Laver M and Hunt WB (1992) *Policy and Party Competition*. New York: Routledge.
- Loewen PJ, Rubenson D and Spirling A (2012) Testing the power of arguments in referendums: A Bradley–Terry approach. *Electoral Studies* 31(1): 212–221.
- Luce D, Krantz D, Suppes P, et al. (eds) (1990) *Foundations of Measurement: Representation, Axiomatization, and Invariance*. Vol. III. New York: Academic Press.
- New York Times* (1979) Times' computer ranks football Top 20. *New York Times*, 19 October, A25.
- Rubinstein A (1982) Perfect equilibrium in a bargaining model. *Econometrica: Journal of the Econometric Society* 50(1): 97–109.
- Turner H and Firth D (2012) Bradley-Terry models in R: The BradleyTerry2 package. *Journal of Statistical Software* 48(9): 1–21.
- Warwick PV and Druckman JN (2001) Portfolio salience and the proportionality of payoffs in coalition governments. *British Journal of Political Science* 31(4): 627–649.