Cabinet formation and portfolio distribution in European multi-party systems

Josh Cutler
Department of Political Science, Duke University
josh.cutler@duke.edu

Scott de Marchi
Department of Political Science, Duke University
demarchi@duke.edu

Max Gallop
Department of Political Science, Duke University
max.gallop@duke.edu

Florian M. Hollenbach
Department of Political Science, Duke University
florian.hollenbach@duke.edu

Michael Laver
Department of Politics, New York University
michael.laver@nyu.edu

Matthias Orlowski
Berlin Graduate School of Social Sciences, Humboldt-University Berlin
mace84@gmail.com

ABSTRACT

Government formation in multiparty systems is of self-evident substantive importance, and the subject of an enormous theoretical literature. Empirical evaluations of models of government formation tend to separate government formation *per se*, from the distribution of key government payoffs such as cabinet portfolios between members of the government that forms. Models of government formation are necessarily specified *ex ante*, absent any knowledge of the government that forms. Models of the distribution of cabinet portfolios are typically, though not necessarily, specified *ex post*, given knowledge of the identity of some government "*formateur*" or even of the full partisan composition of the eventual cabinet. This disjunction lies at the heart of a notorious contradiction between predictions of the distribution of cabinet portfolios made by canonical models of legislative bargaining, and the robust empirical regularity of proportional portfolio allocations – "Gamson's Law". We resolve this contradiction by specifying and estimating a *joint* model of cabinet formation and portfolio distribution, which for example predicts *ex ante* which parties will receive zero portfolios rather than taking this as given *ex post*. We conclude that canonical models of legislative bargaining do add to our ability to predict government membership, but that portfolio distribution between government members conforms robustly to a proportionality norm … we suggest because portfolio distribution follows the much more difficult process of policy bargaining in the typical government formation process.

INTRODUCTION

Government formation in multiparty systems is of self-evident substantive importance. There is a large and heterodox theoretical literature dealing with this. Of special interest to political scientists, key "outputs" of the government formation process, for example partisan composition of the cabinet and distribution of portfolios between cabinet members, as well as bargaining delays and cabinet durations, are easily and unambiguously observable. This enables a more clear-cut confrontation between theory and data than is typically the case in political science. It has also thrown into sharp relief a puzzle that Warwick and Druckman call the "portfolio allocation paradox", a notorious contradiction between predictions of the distribution of cabinet portfolios made by canonical alternating offers models of legislative bargaining, and the robust empirical regularity of proportional portfolio allocations that has been characterized as "Gamson's Law".[1]

Empirical models of government *formation* are necessarily specified *ex ante*, absent any knowledge of the government that forms.[2] Empirical models of the distribution of cabinet portfolios are typically specified *ex post*, given knowledge of the identity of some putative government "*formateur*" or the full partisan composition of the cabinet.[3] We argue below that this disjunction lies at the heart of the portfolio allocation "paradox", and set out to resolve the paradox by specifying and estimating a joint *ex ante* model of cabinet formation and portfolio distribution. An important feature of such a joint model is that it treats the list of parties *outside* the cabinet as informative, predicting *ex ante* which parties will receive zero portfolios rather

[1] (Warwick and Druckman 2006)

[2] See for example (Martin and Stevenson 2010; Martin and Stevenson 2001)

[3] See for example (Ansolabehere et al. 2005; Snyder et al. 2005; Warwick and Druckman 2001, 2006; Falcó-Gimeno and Indridason 2013)

than taking this as given. Our joint model allows us to conclude that canonical models of legislative bargaining do add to our ability to predict government membership, but that portfolio distribution between government members conforms robustly to a proportionality norm, giving support to recent *ex ante* theoretical models of portfolio distribution.

RESOLVING THE PORTFOLIO ALLOCATION PARADOX

The paradox

The essential features of the portfolio allocation paradox have been extensively discussed by Warwick and Druckman among many others,[4] and we confine ourselves here to its bare bones. A long tradition of empirical research linking legislative seat shares to distributions of cabinet positions originates in work published over 50 years ago by the sociologist William Gamson.[5] This established a strong and non-trivial empirical regularity, "Gamson's Law" (GL). Government parties tend to receive cabinet portfolios in strict proportion to the legislative seats they contribute to the government's aggregate seat total. Gamson did not provide a formal bargaining model that yielded GL as a prediction. Indeed, if party leaders are motivated to maximize "Gamsonian" (proportional-to-seats) payoffs, the government coalition should command the smallest legislative seat total that is also a legislative majority. This is not true empirically, a finding first published over 40 years ago.[6] Nonetheless, GL itself has proved extraordinarily robust to replication, even when different cabinet portfolios are assigned very

[4] (Bassi 2013; Carroll and Cox 2007a; Laver et al. 2011; Warwick and Druckman 2006; Falcó-Gimeno and Indridason 2013)

[5] (Ansolabehere et al. 2005; Browne and Franklin 1973; Browne and Frendreis 1980; Fréchette et al. 2005b; Gamson 1961; Laver et al. 2011; Schofield and Laver 1985; Warwick and Druckman 2006; Snyder et al. 2005)

[6] (Fréchette et al. 2005a, 2005b; Taylor and Laver 1973)

different empirical weights.[7] On the other side of the paradox we find canonical alternating offers models of legislative bargaining, with an intellectual pedigree traceable to Rubinstein (1982), adapted to legislative bargaining by Baron and Ferejohn (1989). This approach was applied explicitly to portfolio distribution by Ansolabehere et al., who derive two propositions about legislative bargaining over government formation in multi-party systems:[8]

(1) "Elementary microeconomic theory teaches that in competitive situations perfect substitutes have the same price … We show that the noncooperative bargaining model of David P. Baron and John A. Ferejohn (1989) leads naturally to the result that expected payoffs are proportional to voting weights."[9]

(2) "One key prediction of the model … is that the party that is recognized to form a coalition – the *formateur* – will receive a share of the cabinet posts that is much larger than its share of the voting weight."[10]

On the face of things these two propositions seem contradictory: "payoffs are proportional to voting weights" but "the *formateur* gets a payoff much larger than its share of voting weight". The apparent contradiction arises, because the first proposition is stated *ex ante*, absent knowledge of the identity of the *formateur*. The second proposition is stated *ex post*, knowing the identity of the *formateur*. The propositions are therefore not contradictory because they apply in

[7] (Ansolabehere et al. 2005; Browne and Franklin 1973; Browne and Frendreis 1980; Fréchette et al. 2005b; Gamson 1961; Laver et al. 2011; Schofield and Laver 1985; Warwick and Druckman 2006; Snyder et al. 2005)

[8] (Ansolabehere et al. 2005; Snyder et al. 2005)

[9] (Snyder et al 2005: 982)

[10] (Snyder et al 2005: 992)

different settings, but this highlights a critical distinction between *ex ante* and *ex post* models of cabinet portfolio distribution.

Over and above the special position claimed for the *formateur*, Ansolabehere et al. also diverge from typical work on Gamson's Law by arguing that cabinet payoffs should be associated with theoretical voting weights, specifically "minimum integer weights" (MIWs), rather than the "raw" legislative seat shares of each party. Parties' MIWs are derived as follows from their raw seat shares and the winning threshold for passing votes in the legislature. Use the list of raw seat shares to calculate the list of winning coalitions. Now replace the raw seat shares with the list of smallest integers, one per party, generating the same set of winning coalitions. These integers are the parties' MIWs.[11] Ansolabehere et al. therefore make two simultaneous moves away from traditional research on portfolio distribution. They propose a *formateur* advantage; and they propose using MIWs rather than raw seat shares.

Notwithstanding empirical findings published by Ansolabehere et al., analyses by subsequent authors show little if any empirical support for either of the two theoretical propositions stated above. In relation to the first – use MIWs not raw seat shares – Warwick and Druckman's results, replicated by Laver et al., are "clear and strong: cabinet portfolios, in both number and value, are allocated in very close proportion to the seat contributions of cabinet parties, and the bargaining strengths of these parties distort this allocation principle only very slightly (or very occasionally)."[12] In relation to the second proposition, on the *ex post formateur*

[11] For example, in a 5-party 100-seat legislature with a simple majority winning threshold and a raw seat vector of (43, 35, 8, 8, 6), the MIW vector is (3, 1, 1, 1, 0). The smallest party has zero MIW because it is never pivotal; adding its seats never turns a losing coalition into a winning one. The largest party can form a winning coalition with any of the three middle parties; all three of the middle parties must combine to exclude the largest.

[12] Warwick and Druckman (2006) p 659; Laver et al (2011).

effect, Warwick and Druckman (p659) find "the power that ought to come with *formateur* status appears to yield little in terms of portfolios". Laver et al. also show that estimating an *ex post formateur* effect is compounded by the methodological problem that *formateur* party status is coded endogenously in relevant datasets, as the party of the eventual prime minister.[13] Thus Golder reports that "it is not uncommon for the *formateur* to fail to form a coalition on the first or even the second attempt. As an example, it took seven different coalition proposals more than 106 days for a government to form after the 1979 Belgian legislative elections".[14] None of these failed attempts by *formateurs* is analyzed by Snyder, et al., or by Warwick and Druckman, when they estimate the *formateur* advantage, and only one of 250 *formateur* parties in the dataset they use was *not* the party of the eventual prime minister.[15] The prime ministerial portfolio thus appears on both sides of the relevant regressions, as one of the key independent variables (*formateur* status) and as part of the dependent variable (portfolio share). Correcting for this, Laver, et al. find that the empirical *formateur* effect disappears.[16]

The paradox is therefore both simple and striking. As Carroll and Cox say, "all modern bargaining models predict that Gamson's Law should *not* hold".[17] But, as Bassi says,"the most important empirical law in government-formation studies is that coalition partners share cabinet portfolios in proportion to their relative seat shares, which contradicts the predictions of the entire theoretical literature".[18]

---

[13] (Laver et al. 2011)

[14] (Golder 2010) (p8)

[15] Snyder et al., 2005; Warwick and Druckman 2006

[16] Laver et al., 2011

[17] (Carroll and Cox 2007a) p 301.

[18] (Bassi 2013) p778

Resolving the paradox

Two theoretical questions underlie this paradox, as we have seen. The first is the claim that there is some *formateur* in a privileged bargaining position. The second is the claim that parties' portfolio payoffs respond to their theoretical voting weights (MIWs) rather than their raw seat shares. We argue above against predicting portfolio payoffs *ex post*, after Nature has told us the identity of the randomly chosen *formateur* but before the government has formed, given the problem of identifying the list of exogenously selected *formateurs* in a way that is not endogenous to the bargaining outcome, compounded by the lack of any measurable effect once this problem is addressed. We therefore address the paradox by predicting portfolio allocation *ex ante*, absent any knowledge of the *formateur*, and developing a joint *ex ante* model of government formation and portfolio distribution.

This leads us to the second question, of whether raw seat shares or theoretical voting weights better predict parties' observed portfolio shares. Absent knowledge of the *formateur*, the canonical bargaining models referred to above predict that:

- payoffs will be proportional to theoretical voting weights;
- only minimal winning coalitions will form; (minority cabinets face a majority opposition with both the incentive and ability to capture all portfolios; surplus coalitions contain members who are not needed yet consume some portfolios);
- parties with zero MIW will therefore not be part of any government coalition.

Unlike the canonical bargaining models, the GL prediction that portfolio distribution will reflect raw seats shares lacks deep theoretical underpinnings, though several scholars have recently addressed this. Anna Bassi assumes that *formateurs* are not picked exogenously, but emerge *endogenously* from negotiations between party leaders. Her model predicts no *formateur*

advantage (in equilibrium, party leaders bargain this away), and Gamsonian payoffs.[19] The latter arise because Bassi (p784) assumes any proposed payoff distribution must satisfy the full set of party legislators, not just a single party leader. Carroll and Cox model bargaining over government formation that begins *before* an election, not after this. Their argument is that committing to Gamsonian payoffs in pre-election deals provides incentives for all parties to the deal to expend maximum electoral effort on increasing their legislative seat total, increasing the probability that both they and their coalition partners will take office. Falco-Gimenó and Indridason argue that Gamsonian portfolio distributions act as natural "focal points" in a difficult and complex bargaining environment, involving complicated policy negotiations and many other matters besides, thereby taking one piece of complexity off the bargaining table.[20]

Given theoretical arguments on both sides of the debate, it is helpful to resolve empirically the question of whether parties' portfolio payoffs tend to respond to raw seat shares or their theoretical voting weights. Building on the argument of Falco-Gimenó and Indridason, it is also helpful to investigate whether Gamsonsian proportional payoffs become more likely as the bargaining environment becomes more complex. We address both of these questions below.

A joint *a priori* model of government formation and portfolio distribution

Until now, all empirical work on portfolio payoffs has predicted these *given a particular coalition of parties that has already formed*.[21] In other words, existing empirical work has modeled portfolio distribution *ex post*, conditional on government formation and regardless of whether the government that formed was predicted by any model. From a theoretical perspective,

[19] Bassi 2013

[20] (Falcó-Gimeno and Indridason 2013)

[21] (Ansolabehere et al. 2005; Snyder et al. 2005; Warwick and Druckman 2006; Laver et al. 2011; Carroll and Cox 2007b; Falcó-Gimeno and Indridason 2013)

this is despite the fact that bargaining models predict portfolio distribution as an integral part of a model of government formation, so that (in)ability to predict government membership is highly relevant empirically. In practical terms, *ex post* modeling of payoff distributions results in deletion from the analysis of every party receiving no cabinet portfolio. This happens even when the theoretical model under investigation predicts that, on average, many of these parties should receive a positive bargaining payoff and should with some positive probability be members of the government. Selecting on the dependent variable in this way ignores the possibility that there may be complexities in the bargaining process that systematically select different types of party, contrary to the claim that parties should receive portfolio payoff solely in proportion to their bargaining strength.

We therefore eschew *ex post* prediction of portfolio distributions and specify an *ex ante* statistical approach that includes all parties involved in government formation, whether or not they join the cabinet and receive a payoff, and evaluates joint predictions of government membership and payoff allocation between government members. This is analogous to an approach developed recently by Chiba, et al. which develops a statistical model to evaluate joint predictions of government formation and government duration.[22]

STATISTICAL MODEL

Our unit of analysis is a political party in a government formation situation. Our dependent variable is the party's observed share of cabinet portfolios. This ranges in theory from 0 to 1, and has a point mass at zero in empirical data. A party with zero cabinet portfolios is by all conventional definitions considered not to be a member of the government; a party with non-zero portfolios is considered to be a government member; a party with all the portfolios constitutes a

[22] (Chiba et al. 2014)

single party cabinet. Thus, our dependent variable is far from normally distributed. While we present results of OLS regressions in supplementary materials, to facilitate comparison with previously published work, our core method involves maximum likelihood models based on a mixed continuous–discrete distribution. The dependent variable is a proportion and therefore has a limited range. This suggests the choice of a beta distribution, which is flexible enough to suit the problem (see Brehm and Gates on this topic).[23] As a continuous function, however, it is not appropriate for modeling large numbers of values that are at a single point, and the dependent variable has a very large number of zeros, which arise whenever a given party is not included in the government. This implies using a mixed distribution. We adopt this approach, following work by Ospina and Ferrari,[24] and model these data as a mixture between a beta distribution and a degenerate distribution in 0. By using a zero-inflated beta model such as this, we directly model both the likelihood a party enters a coalition or not, and then, upon entry, its payoff in cabinet seats. Hence, the likelihood function we maximize is:

$$L(y;\nu,\mu,\sigma) = \prod_{i=1}^{n} {\nu_i}^{1-I(y_i)} (1-\nu_i)^{I(y_i)} \prod_{i:y_i\in(0,1)} \frac{\Gamma(\sigma_i)}{\Gamma(\mu_i\sigma_i)\Gamma((1-\mu_i)\sigma_i)} {y_i}^{\mu_i\sigma_i-1}(1-y_i)^{(1-\mu_i)\sigma_i-1},$$

… where:

$$\nu_i = \log\left(\frac{\beta_{\nu0}+\beta_{\nu1}\cdot miw\ share_i+\beta_{\nu2}\cdot seat\ share_i}{1-(\beta_{\nu0}+\beta_{\nu1}\cdot miw\ share_i+\beta_{\nu2}\cdot seat\ share_i)}\right),$$

$$\mu_i = \log\left(\frac{\beta_{\mu0}+\beta_{\mu1}\cdot miw\ share_i+\beta_{\mu2}\cdot seat\ share_i}{1-(\beta_{\mu0}+\beta_{\mu1}\cdot miw\ share_i+\beta_{\mu2}\cdot seat\ share_i)}\right),$$

$$\sigma_i = \log(\beta_{\sigma0}+\beta_{\sigma1}\cdot miw\ share_i+\beta_{\sigma2}\cdot\ seat\ share_i)$$

Link functions were chosen to constrain parameters to the unit interval (in the case of ν_i and μ_i) and to be strictly positive in the case of the precision parameter σ_i. The parameter ν_i has a direct

[23] Brehm and Gates, 1993.

[24] Ospina and Ferrari, 2012. We could employ a more complicated model that includes one's (i.e., governments where a single party forms the governments, despite the fact that this party does not have a majority of seats in the parliament), but there are very few observations where this occurs.

and relevant substantive interpretation: it gives the likelihood of a zero observation, hence the *inverse* likelihood of cabinet membership. The parameter μ_i is the mean for the beta distribution and shows the strength of the relationship between the independent variables and the party portfolio shares; σ_i is the precision of the beta distribution. We thereby achieve our core objective of deriving distinct joint estimates of the probability of cabinet membership, v_i, and the distribution of cabinet portfolios, μ_i.

Weighting cases

Since our unit of analysis is a political party in a government formation situation, we might ask whether all observations in the data should, as in all previous empirical analyses, have equal weight. Different countries have different numbers of parties and some have many more governments than others. Italy, for example, had more governments than most other European countries in the post-war era and also tends to have more parties than most other countries. The dataset we describe below covers 16 countries, but Italy accounts for about 22% of the observations. Italy, Belgium, Denmark, and Finland together account for about 55% of all observations. Empirical analyses that ignore this may be biased, and we don't want our "cross-national" analysis of portfolio distribution to be mostly about Italy, Belgium, Denmark, and Finland. In order to control for the different weights of parties in the estimation due to the differences in party system sizes and frequency of elections, we calculated a probability for each party to be selected into a subsample of our data. This selection probability is inversely related to the number of parties in parliament. It is calculated as:

$$\pi_i = \frac{1}{n_k} \cdot \frac{1}{K}$$

… where n_k is the number of parties in a given cabinet k and K is the total number of cabinets in the truncated data set. Based on the selection probabilities, 80% of observations (1134 parties) in the truncated data set are randomly selected in a subsample.[25] By providing results based on multiple, randomly drawn samples, we mitigate the possibility that one nation exercises undue influence on the results.[26] In many respects, our sampling strategy functions like the bootstrap. By creating empirical distributions from the data and generating multiple samples (with replacement), we hedge against any single nation (or subset of cabinets) having undue influence on our results.[27]

Estimation approach

As indicated in the previous section, our estimation method involves creating multiple training sets based on the principle of the empirical bootstrap. We use these training sets to test three different models. The first includes all cabinets and is the main test of whether theoretical voting weights or raw seat shares best predict outcomes, including both government membership and

[25] We use a threshold of 80% for the training set (and accordingly, 20% for the test set) to insure that the test samples are large enough to make reasonable inferences. Splitting the sample in this way is accepted practice in the machine learning literature, though there are more complex ways to accomplish the same goals (see, for example, Bishop and Nasrabadi, 2006).

[26] To provide comparability with previously published results that do not weight cases in this way, we provide unweighted ("flat pi"), boostrapped OLS versions of all key estimates in the online supplementary materials (Appendix 3, Tables A3.1 and A3.2).

[27] We also tried a more aggressive strategy that rebalanced the relative proportions of the various nations in the training sets to make them more even. This modified approach did not produce substantively different results but is available on request.

payoff distribution. Addressing Falco-Gimenó and Indridason's argument that Gamsonian payoffs are a response by party leaders to bargaining complexity, we test two further models using these training data, separating bargaining settings according to their complexity. We divided "simple" from "complex" settings according to the difficulty of calculating the full set of 2^n - 1 coalitions in an *n*-party system, and therefore of exploring all coalition possibilities, using a natural inflection point in the time taken by our computational algorithm to calculate MIWs. This led us to specify simple settings as those with fewer than eight pivotal parties and complex settings as those with eight or more pivotal parties.[28] For the full population of cases, as well as those reflecting simple and complex settings, we draw 1000 subsamples from the data set, split each subsample into training and test sets and fit the model. Instead of evaluating our model performance using the training samples, we report results based only on the 20% out-of-sample observations in the test set.[29] This is a further hedge against the possibility that results could be overfitted to non-systematic elements of the original sample.

DATA

Most of our data derive from the *Parliament and Government Composition Database* (Parlgov), constructed by Döring, Manow and collaborators.[30] This includes election results and government formation data for all EU members as well as many OECD countries from 1945 onwards. Our dataset includes data from 1945 until the most recent available Parlgov data on

[28] A pivotal party P is one with non-zero MIW. For robustness, we examined models setting the "complexity" threshold at 7 or 9 pivotal parties and the results were similar.

[29] Given the vanishingly small probability that all members of any cabinet would be selected in all samples, this sampling strategy also addresses the methodological issue that the set of party portfolio proportions within any given cabinet generate compositional data.

[30] http://www.parlgov.org

cabinet seats for 16 parliamentary democracies in Western Europe.[31] While we rely on Parlgov for data on elections, seat share, as well as the party of the prime minister, this source does not include data on the distribution of cabinet seats between parties and we collected these data from a reliable online database.[32] Summary statistics are contained in online supplementary materials (Appendix 2).

Theoretical voting weights

A key independent variable in our analysis concerns parties' MIWs, which Ansolabehere et al. use to implement their claim that theoretical voting weights rather than raw seat shares are what inform government formation and subsequent portfolio allocation. There is a focused technical literature on MIWs, much of it outside mainstream political science.[33] Montero has shown, for example, that *ex post* predicted payoffs under the canonical alternating offers bargaining protocol are proportional to agents' MIWs.[34] The vector of MIWs can be surprisingly difficult to calculate, especially in systems with more than a few parties. Computationally, as Strauss shows, this is because the problem of coalition enumeration is NP-hard, which of course generates problems for real politicians as well as for political scientists.[35] While Snyder, et al. address this by programming a calculator for computing MIWs,[36] the "perfect substitutes have the same price" argument they deploy highlights a theoretical distinction drawn by Freixas and Kurz

[31] These countries are: Austria, Australia, Belgium, Denmark, Finland, Germany, Great Britain, Iceland, Ireland, Italy, Luxembourg, Netherlands, Norway, Portugal, Spain, Sweden.

[32] The data on the number of cabinet seats were collected from http://www.kolumbus.fi/taglarsson/dokumentit/governm2.htm.

[33] (Freixas and Kurz 2011; Freixas and Molinero 2009a, 2009b; Montero 2002, 2006)

[34] (Montero 2006)

[35] (Strauss 2003)

[36] (Ansolabehere et al. 2005; Snyder et al. 2005)

between "simple" MIWs, and "MIWs preserving types". At issue is whether mutually substitutable parties should be given equal voting weights, as Snyder et al. argue they should. If so, we should use "MIWs preserving types", which are constrained so that parties that are perfect substitutes have the same voting weights.[37] Following Freixas and Kurz, therefore, we replace MIWs with MIWs preserving types.[38] Solving for these is an application of linear programming.[39] We programmed and verified our own algorithm and, assuming a simple majority winning quota, calculated MIWs preserving types from party seat totals reported in Parlgov.

Single party majorities and minority governments

When a single political party wins a legislative majority, models of government formation invariably converge on the prediction that it will form a single party government and award itself all cabinet portfolios. This is almost invariably true empirically and is not an open theoretical question. While we might think of plausible models that treat single party majority governments as coalitions of party factions, this is neither the focus of the literature with which we are dealing

[37] (Freixas and Kurz 2011) . Freixas and Kurz note this problem is especially acute with "non-homogenous" voting games and can occur with as few as eight parties. "Non-homogenous" voting games are those for which all MWCs do not have the same aggregate MIW. For example in the legislature (8: 4,3,3,2,2), expressed in MIWs, the coalitions (4,2,2), (4,3,2) and (4,3,3) are all MWCs, but each has a different aggregate MIW. Theoretical complications arise because a party with MIW 2 in this example can substitute in an MWC for a party with weight 3. Laver et al. (2011) find that non-homogenous voting games arise in about one-third of real legislatures in the Snyder et al. replication dataset, so this is a non-trivial issue.

[38] For stylistic reasons, we refer to these weights as MIWs in what follows. The Freixas and Kurz definition should, in our opinion, be universally adopted.

39 See Appendix 1 in online supplementary materials

nor our focus in this paper. We therefore exclude from our analysis, as uninformative for our purposes, settings in which a single party won a legislative majority.

Minority cabinets pose a more difficult challenge for a joint model of government formation and portfolio distribution. Predicting the *formation* of minority governments remains a wide open theoretical question for models of government formation. But models focusing on *portfolio distribution* predict no minority governments, as we have noted, since a minority government faces a majority opposition both willing and able to consume all portfolios. To include minority governments in our analysis is thus to include a set of cases where we know the models of portfolio distribution we are evaluating make failed predictions, and it is not at all clear what portfolio distributions these models predict in cabinets they deem out of equilibrium. To exclude minority cabinets, however, amounts to selecting on the dependent variable for a model that jointly predicts government formation and portfolio distribution, and to include an indicator variable for minority governments on the right hand side of the relevant regressions violates our objective of *ex ante* prediction. Since our goal of *ex ante* prediction is paramount and we do not wish to select on the independent variable, we opt for the lesser of two evils and include all cases of minority coalitions in our analysis.[40]

RESULTS

Estimation results are reported in Tables 1 – 3 below; these include all observations except those where a single party forms the government.[41] The first parameter of interest is v showing the

[40] Results for models that exclude minority governments are available on request. In broad terms, and unsurprisingly, excluding minority governments improves overall model fit, though only by a small margin. The effects for the variables of interest – MIWs and raw seat shares – remain the same.

[41] As noted, we exclude cases where one party receives a simple majority of the seats in the parliament. A much smaller subset of exclusions (75 observations) concerns observations where a single minority party

relationship between the independent variables and a party being out of office (having no portfolio). The second parameter is μ, showing the relationship between the independent variables and the strictly positive proportion of portfolios allocated to a party. Taking all cases together (Table 1), both the MIW and raw legislative seat share of a party are correlated with its inclusion in government.

Table 1: Distribution of 1000 parameter estimates and root mean squared error: out of sample predictions for all cabinets

		Mean	Std. Dev.
ν (cabinet membership)	Intercept	1.95	.04
	MIW	-5.75	.45
	Raw Weights	-2.94	.35
μ (portfolio share)	Intercept	-2.14	.02
	MIW	1.32	.14
	Raw Weights	5.43	.13
σ	Intercept	3.33	.06
	MIW	0.15	.50
	Raw Weights	-2.51	.43
	RMSE	0.16	.003

The ν coefficients show a negative relationship between exclusion from office for both MIWs and raw seat shares; higher MIWs are more strongly connected with cabinet membership than higher raw seat shares. However, and consistent with Warwick and Druckman's findings, the μ coefficients in Table 1 show that raw seats share are far better predictors than MIWs of the share of cabinet portfolios a party receives, conditional on it being in the government at all.

Our analysis therefore yields an unambiguous substantive finding about legislative bargaining. Independent of raw legislative seat shares, parties are more likely to get into the

receives all portfolios, since these do not conform to the assumptions of a zero-inflated beta model. We do however include such cases in the, substantively similar, results estimated using OLS regressions and reported in Appendix 3 of the online supplementary materials (Tables 5 and 6).

cabinet if they have a higher theoretical voting weight, measured using MIWs. Theoretical bargaining models do in this sense add value to predictions about cabinet *membership*. Once a party is in the government, however, its *payoff*, measured as its share of cabinet portfolios, is predicted by its raw legislative seat share and not at all by its theoretical bargaining weight. In this precise sense, the empirical robustness of Gamson's Law has again been vindicated. While raw seat shares robustly predict portfolio distribution, they do not do this because they are associated with theoretical voting weights.

Comparing Tables 2 and 3, we see that bargaining complexity has a strong effect. In the more complex settings described in Table 3, the ν coefficients show that MIWs no longer even predict cabinet membership, which is much better predicted by raw legislative seat shares. Table 2 shows the reverse for less complex settings, bearing in mind these are the typical settings for post-war Western Europe; *membership* of the government is much better predicted by MIWs, controlling for raw legislative seat shares. In both types of setting, however, the μ coefficients show that it is raw legislative seat shares, and not MIWs, that predict portfolio payoffs.

Table 2: Distribution of 1000 parameter estimates and root mean squared error out-of-sample predictions for cabinets with at most 8 parties with non-zero MIWs

		Mean	Std. Dev.
ν	Intercept	2.31	.08
	MIW	-8.14	.58
	Raw Weights	-1.42	.45
μ	Intercept	-1.83	.06
	MIW	0.73	.32
	Raw Weights	4.91	.19
σ	Intercept	2.70	.18
	MIW	4.01	1.14
	Raw Weights	-4.41	.71
	RMSE	.18	.004

Table 3: Distribution of 1000 parameter estimates and root mean squared error out-of-sample predictions for cabinets with more than 8 parties with non-zero MIWs

		Mean	Std. Dev.
ν	Intercept	1.89	.05
	MIW	-4.19	.80
	Raw Weights	-5.13	.63
μ	Intercept	-2.25	.02
	MIW	-0.02	.21
	Raw Weights	7.17	.20
σ	Intercept	3.46	.07
	MIW	-1.80	.80
	Raw Weights	0.35	.69
	RMSE	.116	.004

To facilitate comparison with previously published results, we present estimates derived from OLS regressions in the online supplementary materials (Appendix 3), with the *caveat* these models are mis-specified for reasons we discuss above. The OLS estimates comport with those derived from zero inflated beta models presented here. Theoretical voting weights predict government membership, especially in the lower complexity settings with eight or fewer pivotal parties. But it is raw legislative seats, not MIWs, which predict distributions of cabinet portfolios within a given coalition.

CONCLUSIONS

Our zero-inflated beta model jointly estimates membership of the government and, conditional on this, the distribution of cabinet portfolios between government members. The results are unambiguous. First, controlling for raw legislative seat shares, the theoretical voting weights used by non-cooperative models of legislative bargaining do predict *membership* of the government emerging from bargaining between parties, especially in settings with fewer than eight pivotal parties that are typical in postwar Western Europe. Indeed Table 1 shows that

theoretical voting weights predict government membership much better than raw seat shares in this common setting. In a nutshell, controlling for raw seat shares, *parties with higher theoretical voting weights are better able to force their way into government coalitions.* This is a theoretically significant empirical finding. Second, whatever the setting and conditional on the government that forms, *raw legislative seat shares, not theoretical voting weights, predict the distribution of cabinet portfolios between government parties.* The latter finding is, in effect, a further replication of the robust Gamson's Law results using a more sophisticated and appropriate statistical apparatus. Third, Tables 2 and 3 show that, while MIWs predict cabinet membership in simple bargaining settings, they are much less effective than raw seat shares at predicting cabinet membership in more complex settings.

We draw the following conclusions from these findings. First, our findings imply that the scientific community should accept and exploit the robust empirical regularity that raw seat shares and not theoretical voting weights predict the distribution of cabinet portfolios in European governments. The theoretical task is not to make this empirical result go away but to find a rigorous theoretical model of government formation that explains it. Second, and newly emerging from our joint modeling of government membership and portfolio distribution, the good news for legislative bargaining models is that theoretical voting weights, more precisely MIWs preserving types, do help to explain government membership.

We see no necessary contradiction between these findings. One way to reconcile them is to infer that the distribution of cabinet portfolios is neither the only, nor even the most important, matter that concerns party leaders bargaining over the formation of coalition cabinets. Golder's analysis of bargaining delays over government formation in western Europe shows that, in post-electoral settings, these delays range from 86 days on average in the Netherlands though 61 days

in Belgium to less than three days in Norway.[42] One recent Belgian government took more than a year to form. Informal press reports of what is actually happening during these often lengthy bargaining periods suggest that the party leaders are *not* spending their time bargaining over the distribution of cabinet portfolios, but rather are trying to negotiate a joint policy program for government that reconciles conflicting party manifestos. This is detailed and time consuming work, and it appears from these reports that party leaders typically turn to the distribution of cabinet portfolios only after a joint policy program has been agreed. For us to go further would require setting out a model of bargaining over government formation that takes account of both portfolio distribution and the need for a joint policy program. This is clearly not our task in this paper though we are exploring it in ongoing work. Our point here, simply put, is that our twin empirical findings can be reconciled if portfolio distribution is not the only payoff that concerns party leaders when they bargain over government formation. If it were, then we would expect parties' theoretical voting weights to predict portfolio distribution, but they do not. Rather, while these voting weights have a significant bearing on who gets into government, the distribution of portfolios between the parties in government very systematically conforms to a norm of proportionality with raw legislative seat shares.

[42] Golder 2010, 8

REFERENCES

Ansolabehere, Stephen, James M. Snyder, Aaron B. Strauss, and Michael M. Ting. 2005. "Voting weights and formateur advantages in the formation of coalition governents." *American Journal of Political Science* 49 (3):550-63.

Bassi, Anna. 2013. "A model of endogenous government formation." *American Journal of Political Science*.

Brehm, John, and Scott Gates. 1993. "Donut shops and speed traps: evaluating models of supervision on police behavior." *American Journal of Political Science*:555-81.

Browne, Eric, and Mark Franklin. 1973. "Aspects of Coalition Payoffs in European Parliamentary Democracies." *American Political Science Review* 67:453-69.

Browne, Eric, and John Frendreis. 1980. "Allocating Coalition Payoffs by Conventional Norm: An Assessment of the Evidence from Cabinet Coalition Situations." *American Journal of Political Science* 24 753-68.

Carroll, Royce, and Gary W Cox. 2007a. "The Logic of Gamson's Law: Pre-election Coalitions and Portfolio Allocations." *American Journal of Political Science* 51 (2):300-13.

Carroll, Royce, and Gary W. Cox. 2007b. "The logic of Gamson's Law: pre-election coalitions and portfolio allocations." *American Journal of Political Science* 51 (2):300-13.

Chiba, D., L. Martin, and R. Stevenson. 2014. "A copula approach to the problem of selection bias in models of government survival." University of Essex.

Diermeier, Daniel, and Peter Van Roozendaal. 1998. "The duration of cabinet formation processes in western multi-party democracies." *British Journal of Political Science* 28:609-26

Falcó-Gimeno, Albert, and Indridi H Indridason. 2013. "Uncertainty, complexity, and Gamson's Law: Comparing coalition formation in Western Europe." *West European Politics* 36 (1):221-47.

Fréchette, Guillaume, John H. Kagel, and Massimo Morelli. 2005a. "Behavioral identification in coalitional bargaining: an experimental analysis of demand bargaining and alternating offers." *Econometrica* 73 (6):1893-937.

Fréchette, Guillaume, John H. Kagel, and Massimo Morelli. 2005b. "Gamson's law versus non-cooperative bargaining theory." *Games and Economic Behavior* 51:365-90.

Freixas, J., and S. Kurz. 2011. "On minimal integer representations of weighted games." *arXiv preprint arXiv:1103.0868*.

Freixas, Josep, and Xavier Molinero. 2009a. "On the existence of a minimum integer representation for weighted voting systems." *Annals of Operations Research* 166 (1):243-60.

Freixas, Josep, and Xavier Molinero. 2009b. "Weighted games without a unique minimal representation in integers." *Optimization Methods and Software* 25 (2):203-15.

Gamson, William A. 1961. "A Theory of Coalition Formation." *American Sociological Review* 26:373-82.

Golder, Sona N. 2010. "Bargaining delays in the government formation process." *Comparative Political Studies* 43 (1):3–32.

Laver, Michael, Scott de Marchi, and Hande Mutlu. 2011. "Negotiation in legislatures over government formation." *Public Choice* 147 (3):285-304.

Martin, Lanny W., and Randolph Stevenson. 2010. "The conditional impact of incumbency on government formation." *American Political Science Review* 104:503-18.

Martin, Lanny W., and Randolph T. Stevenson. 2001. "Government Formation in Parliamentary Democracies." *American Journal of Political Science* 45 (1):33-50.

Martin, Lanny W., and Georg Vanberg. 2003. "Wasting Time? The Impact of Ideology and Size on Delay in Coalition Formation." *British Journal of Political Science* 33 (02):323-32.

Montero, Maria. 2002. "Non-cooperative bargaining in apex games and the kernel." *Games and Economic Behavior* 41 (2):309-21.

Montero, Maria. 2006. "Noncooperative foundations of the nucleolus in majority games." *Games and Economic Behavior* 54:380–97.

Ospina, Raydonal, and Silvia LP Ferrari. 2012. "A general class of zero-or-one inflated beta regression models." *Computational Statistics & Data Analysis* 56 (6):1609-23.

Schofield, Norman, and Michael Laver. 1985. " Bargaining theory and portfolio payoffs in European coalition governments." *British Journal of Political Science* 15:51-72.

Snyder, James M., Michael Ting, and Stephen Ansolabehere. 2005. "Legislative bargaining under weighted voting." *American Economic Review* 95 (4):981-1004.

Strauss, Aaron. 2003. Applying Integer Programming Techniques to Find Minimum Integer Weights of Voting Games, Department of Electrical Engineering and Computer Science, Massachusetts Institute of Technology, Cambridge Mass.

Taylor, Michael, and Michael Laver. 1973. "Government coalitions in Western Europe." *European Journal of Political Research* 1:205-48.

Warwick, Paul V., and James N. Druckman. 2001. "Portfolio salience and the proportionality of payoffs in coalition governments." *British Journal of Political Science* 31:627-49.

Warwick, Paul V., and James N. Druckman. 2006. "The portfolio allocation paradox: an investigation into the nature of a very strong but puzzling relationship." *European Journal of Political Research* 45:635-65.