Bargaining in *n*-party legislatures over government formation*

Michael Laver, New York University
Scott de Marchi, Duke University
Hande Mutlu, New York University

ABSTRACT

This paper questions results that claim to extend non-cooperative models of bargaining in legislatures from the highly atypical three-party case to a generic *n*-party setting. It identifies problems both with the derivation of theoretical results and the empirical evaluation of these. No empirically robust formateur advantage can be observed in field data on bargaining over government formation. The paper concludes with a modeling agenda of uncontroversial empirical statements about the government formation process and argues that these should form the premises of a more compelling new model of this crucial political process.

* Thanks are due to John Aldrich, Gary Cox, John Ferejohn, Guillaume Fréchette, Macartan Humphreys, Peter Morriss, Kenneth Shepsle and Paul Warwick for comments on earlier drafts of this paper

1. INTRODUCTION

Bargaining in legislatures may concern anything to be decided by a legislature, including the division of private goods, the setting of public policy and/or production of public goods. It is of special importance in parliamentary democracies, where the most important decisions facing legislators are choosing a new government and deciding whether to keep an incumbent government in office. Early theoretical work on this topic developed within the traditions of cooperative game theory and resulted a series of papers on “indices” of legislative bargaining power, traceable to work by Shapley and Shubik and by Banzhaf (Banzhaf 1965; Shapley and Shubik 1954). Felsenthal and Machover provide helpful overview of this theoretical tradition (Felsenthal and Machover 2001). The shift from co-operative to non-cooperative bargaining models began with the now canonical alternating-offers model proposed by Baron and Ferejohn, hereafter BF (Baron and Ferejohn 1989). Morelli proposed an alternative “demand bargaining” model that has also attracted attention, although core results from this have recently been called into question by Montero and Vidal-Puga (Baron and Ferejohn 1989; Montero and Vidal-Puga 2005; Morelli 1999). The defining feature of this approach, indeed of all current non-cooperative models of bargaining in legislatures, is an assumed bargaining protocol that, at any given time point, uses an *exogenous* recognition rule to identify a single agent with the monopoly right to make a proposal. When the bargaining is over government formation in parliamentary democracies, this agent is known as the government *formateur*. The core BF prediction is that, in equilibrium, the *formateur* proposes a minimum winning coalition (MWC) of agents, in which other members receive their continuation values in the bargaining game and the *formateur* retains the balance. The key result is that there will be a disproportionally high payoff to the *formateur*.

Non-cooperative models of bargaining in legislatures typically begin with the tractable but entirely unrealistic assumption of three legislative parties. Since results from such models do not generalize to *n*-party settings, and since we almost never observe a three-party legislature in the real world[1], the search for a model of bargaining in *n*-party legislatures is a very significant intellectual project. In this context, Snyder, Ting and Ansolabehere (hereafter STA) set out to extend the BF approach from its original legislative setting of three parties, none with a majority, to one with an arbitrary number of parties (Ansolabehere et al. 2005; Snyder, Ting, and Ansolabehere 2005). STA follow most previous authors in using “field data” on government formation in parliamentary democracies to test theoretical work on bargaining in legislatures.

What is striking in this context is that government formation is characterized by a strong and robust empirical regularity, Gamson's Law (GL), which contradicts the canonical BF model. GL states that the proportion of cabinet ministries received by each government party, following bargaining over government formation, equals the proportion of seats contributed by that party to the government seat total. GL thus differs from BF-style bargaining models in predicting no *formateur* advantage; it has been tested and retested many times over four decades, and has proved remarkably robust (Browne and Franklin 1973; Browne and Frendreis 1980; Gamson 1961; Laver and Schofield 1998; Warwick and Druckman 2001; Warwick and Druckman 2006). We are left with what Warwick and Druckman (2006) call the “portfolio allocation paradox”.

[1] It is very important to keep in mind that we are talking in this context about *all* political parties represented in a legislature, however small. The conclusion that some small parties are “dummies” that can be ignored by a bargaining model, while other small parties are “pivotal”, must be derived from the model itself and cannot be assumed *ex ante*.

The profession's canonical theory of bargaining in legislatures is contradicted by one of the profession's strongest and most robust empirical laws. Figure 1 illustrates this paradox at its sharpest. It uses the STA replication dataset to plot government parties' portfolio payoffs against their shares of the government's legislative seat total, for a case universe precisely matching conditions assumed by Baron and Ferejohn – that is, for governments forming in legislatures in which only three parties have non-zero voting weight.

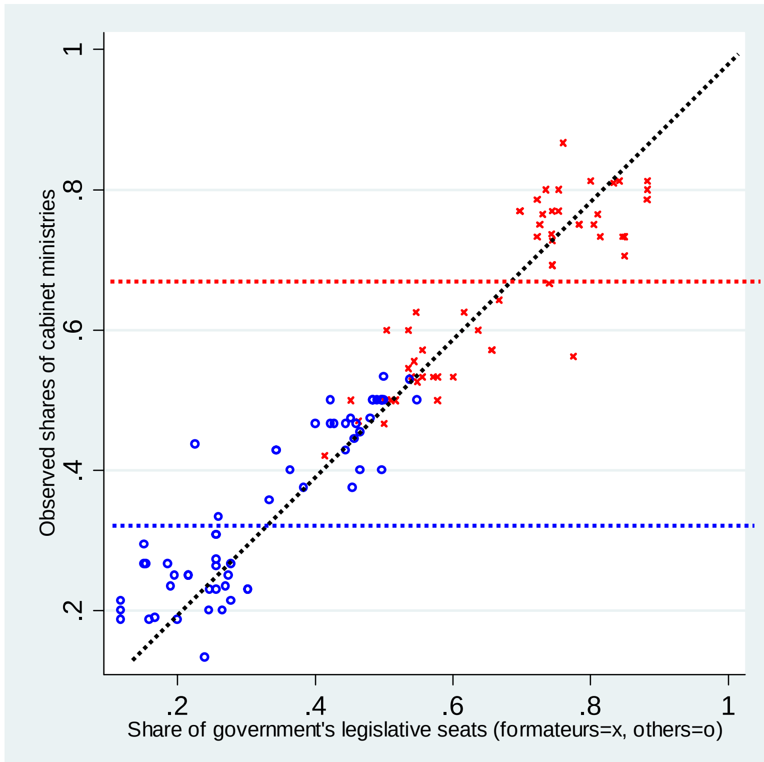

Figure 1: The relationship between seat shares and portfolio payoffs to government parties in "three-party" legislatures

GL states that parties' portfolio payoffs are proportional to their share of the government's legislative seat total – and thus that they will lie on the 45-degree line in Figure 1. BF predicts that *formateur* parties (plotted as "x") will receive two-thirds of the portfolio payoffs in this setting and other parties (plotted as "o") will receive one third. BF thus predicts portfolio payoffs along the upper horizontal line for *formateur* parties and along the lower horizontal line for other parties. Figure 1 tells a crystal-clear story; BF is empirically false. Observed portfolio payoffs tend strongly to lie on the 45-degree Gamson line, not on the horizontal BF lines. Large parties get more. *Formateur* parties tend to get more because they tend to be large but, *directly contra* BF, *formateur and non-formateur parties of the same size get the same payoff.*

The first puzzle we address in this paper concerns how, given Figure 1, scholars might infer empirical support for BF-style bargaining over government formation – that "there is a strong, significant, formateur advantage … consistent with proposal-based bargaining models" (Ansolabehere et al. 2005: 561). Our conclusion is that all empirical support for BF-style bargaining disappears once account is taken of the fact that the crucial dependent variable, *formateur* status, is endogenously coded in the data and thus appears on both sides of the relevant regressions. Given this lack of empirical support, our second puzzle is to find the potential flaws in the *n*-party extension of the BF model. Our conclusion is that STA's core propositions are consistent with many different things, including both BF and GL, as the result of what we argue to be a flawed proof strategy deployed to deal with the awkward problem, in an *n*-party setting, of "non-homogenous" weighted voting games. Nonetheless, this does not explain the clear empirical failure of the core BF model prediction documented in Figure 1, which shows the lack of any observable *formateur* advantage. We argue this follows from a fundamentally flawed modeling assumption about the government formation process in parliamentary democracies. Since this leaves us with no model of bargaining in legislatures that is both theoretically rigorous and supported in field data, we conclude that this important topic must be reconsidered from the bottom up. We begin this reconsideration with a review of a set of statements about bargaining over government formation that we take to be self-evident and uncontroversial premises for a new model. Before doing any of this, we review the current state of the literature on BF bargaining in *n*-party legislatures.

2. BARON-FEREJOHN BARGAINING IN *n*-PARTY LEGISLATURES

The Baron-Ferejohn foundation

The original Baron-Ferejohn (1989) paper specifies and solves a precise model of legislative bargaining with particular parameter settings. Subsequent models claiming BF ancestry retain core features of the original model but use different parameter settings, generating a diverse family of "BF-style" bargaining models. A BF-style legislative bargaining game is characterized by an exogenously determined and fixed set of perfectly disciplined political parties, each led by a single rational agent who bargains on the party's behalf and consumes all payoffs from government formation. There is an exogenous mechanism that determines a vector $\boldsymbol{L}$ specifying the numbers of legislators, l_i controlled by each party i. There is a quota, Q_l, determined under exogenous rules of legislative procedure, defining the number of legislators required to pass any proposal. There is a random recognition mechanism, parameterized by a vector $\boldsymbol{P}$ of exogenously determined common knowledge recognition probabilities, p_i, for each party. At each stage in the bargaining process, a single agent is selected by the recognition mechanism to have the monopoly right to make one proposal to other agents. Agents not recognized by this exogenous mechanism may not make any proposal. All proposals are immediately voted on without debate or amendment. All legislators vote and, if votes in favor equal or exceed Q_l, the proposal is instantly implemented by an unmodeled automaton.[2] The game then ends and all payoffs are consumed. If a proposal does not pass the winning threshold, another monopoly proposer is selected by the recognition mechanism. Time horizons and discount rates vary between models, but in all cases the reversion point is zero for all parties if no agreement is reached; there is no *status quo* allocation.

Common features of sub-game perfect stationary equilibria of published BF-style bargaining models include: the first agent recognized by the random mechanism makes an equilibrium offer to some other agent(s); this offer is accepted; the combined voting weights of the proposer and

[2] In real parliamentary democracies, constitutional rules specify precise procedures for implementing agreements over government formation that typically involve a sequence of parliamentary votes, for house Speaker, Prime Minister, cabinet ministers, junior ministers, etc., as well as actions by the Head of State. In federal systems, more complicated arrangements may prevail.

recipient(s) of this offer exceed the winning threshold; no offer is made to an agent whose voting weight is not required to pass the winning threshold; each other agent receiving an offer receives only his/her reservation price (continuation value in the iterated game) while the proposer gets the balance of the payoff, resulting in a bonus for the proposer.[3] If the game has a real time line and agents discount future payoffs, then any bonus to the first proposer is enhanced.[4] Many implications of BF-style models – that only minimum winning coalitions (MWCs) will form, that the first *formateur*'s offer is always accepted – are testable, though in practice few have been investigated in empirical work to date, which has focused on portfolio payoffs to *formateur* and other parties. It is well-known, for example, that the majority of governments that form in such settings are *not* MWCs, directly contradicting a core BF model prediction. Indeed in STA's own replication dataset, only 143 of the 329 governments analyzed are MWCs.

STA'S *n*-party extension of three-party BF-style bargaining

In setting out to extend the BF model to *n*-party politics, STA apply a basic insight of microeconomic theory:

> Elementary microeconomic theory teaches that in competitive situations perfect substitutes have the same price. In a political setting in which votes might be traded or transferred in the formation of coalitions, one might expect the same logic to apply. If a player has *k* votes, then that player should command a price for those votes equal to the total price of *k* players that each have one vote. In terms of expected payoffs, the player with *k* votes should expect to have a payoff *k* times as great as the payoff expected by a player with one vote (STA, p. 982).

Building on this argument, the core result generated by STA can be found in their Propositions 2 and 3. This states that agents' continuation values in an *n*-party BF-style bargaining game are proportional to their voting weights.[5] The consequence of these results, assuming BF-style bargaining, is that the *formateur*, once selected by the random recognition mechanism, offers coalition partners in some MWC their continuation values, retaining the surplus and resulting once more in a *formateur* advantage. If correct, these results would be very significant indeed, because they would extend the canonical non-cooperative bargaining model to a much more general *n*-party context.

STA test their model predictions using field data on portfolio allocations in coalition cabinets. Regressing portfolio allocation shares on voting weight shares and adding a dummy variable for observed *formateur* status, they infer empirical support for their model from the fact that the coefficient for the formateur dummy is positive and statistically significant, providing "strong evidence that the parties chosen to form a coalition typically receive more than their voting weight" (STA: 994).

---

[3] The Morelli "demand bargaining" model does not imply this.

[4] The original BF model includes an assumption of time preference with a common discount rate; the STA extension to *n*-party systems assumes no time discounting.

[5] This proposition is a bit more complex than that stated here. STA (p 986) claim that "the price a type-*t* coalition partner can command equals that player's continuation value … divided by his or her share of the voting weight in the *r*th replication". The intuition, however, is the same and we discuss the crucial role of STA "replication" approach in the following section. STA's Proposition 4 modifies this conclusion somewhat for certain corner solutions when recognition probabilities are equal for all agents, but equilibrium continuation values remain monotonic in voting weights.

We now show in section 3, however, that the statistical inferences to be drawn from the STA replication dataset, in stark contrast to the authors' own claims, provide little or no evidence in favor of BF-style bargaining models. In section 4, we will return to STA's propositions 2 and 3 (i.e., that continuation values are proportional to voting weights), focusing on flaws in their proofs.

3. EMPIRICAL PROBLEMS WITH BF-STYLE MODELS

It is hard to escape the conclusion that BF-style bargaining models are empirically false in the context of government formation. Figure 1 flatly contradicts STA model predictions for the restricted case universe in which only three parties have non-zero voting weight. Turning to the more general *n*-party case, Figure 2 uses the STA replication dataset to compare the precise predictions made by the STA model (horizontal axis) with observed portfolio payoffs (vertical axis). All cases analyzed by STA in their empirical work are plotted and the 45-degree line shows where perfect model predictions would lie. Observed portfolio payoffs to *formateur* parties (plotted as "x") are distinguished from payoffs to other parties (plotted as "o"). Once more the pattern could not be clearer. STA's *n*-party extension of the BF model systematically over-predicts payoffs for *formateur* parties and systematically under-predicts payoffs for all other parties. Figure 2 shows that something is clearly wrong with the STA model predictions.

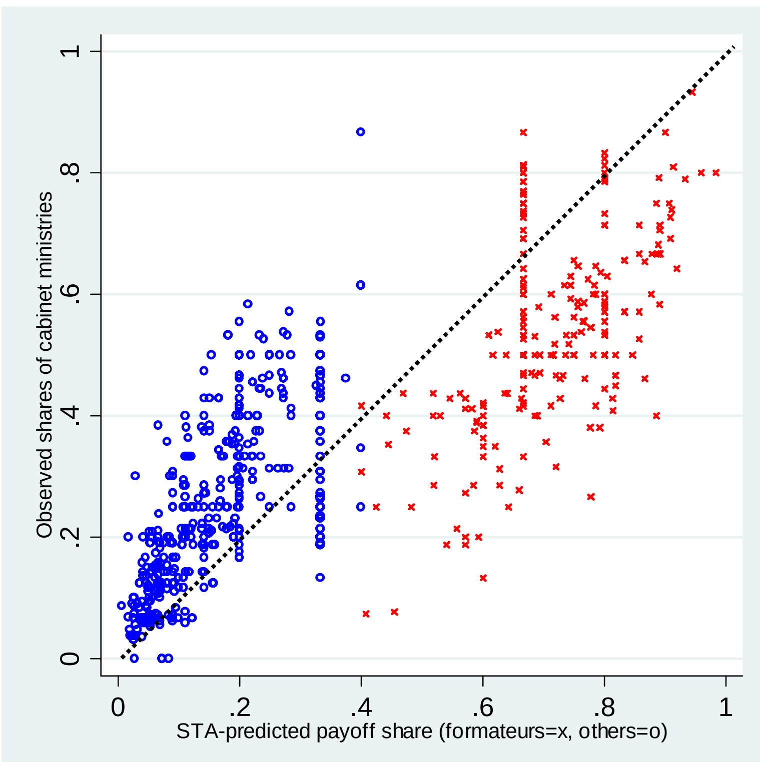

Figure 2: BF model predictions versus observed portfolio payoffs

The problem of untangling what is wrong with these predictions is complicated by the fact that STA's empirical analysis extends traditional empirical work on Gamson's Law in two directions at the same time. The first direction is to introduce a *formateur* dummy into the predictions of parties' portfolio allocations – this is the seminal BF extension. The second direction is to substitute parties' "theoretical voting weights" and specifically "minimum integer weights" (MIWs) for their raw seat shares.[6] This is a new departure and not a feature of the seminal BF model. BF use raw seat shares and, in an assumption supported empirically in work by Diermeir and Merlo, set recognition probabilities proportional to these not to theoretical voting weights (Diermeier and Merlo 2004). Very strikingly, furthermore, STA state in a footnote that a *formateur* advantage is *not* observed in field data if raw seat shares are used rather than MIWs (STA fn 23). This is easily seen from Figure 3, which extends Figure 1 to the full STA case universe and shows the classic Gamson's Law regression. Comparing parties with the same raw seat shares, *formateur* parties do not get more. STA's empirical findings depend crucially on the use of MIWs rather than raw seat shares.

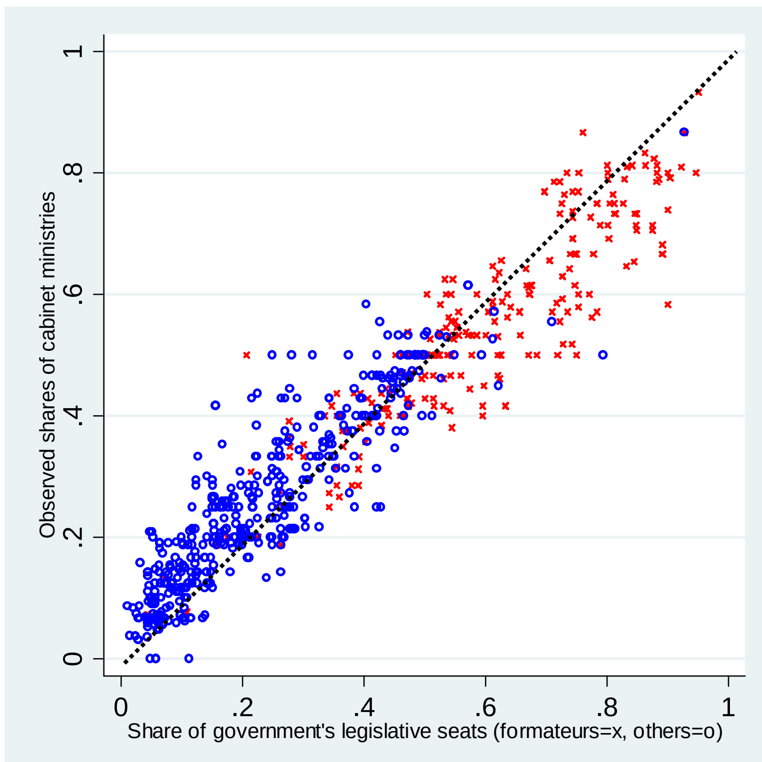

Figure 3: Gamson's Law: observed party shares of cabinet ministries, by observed share of cabinet's legislative seat total

We thus identify two distinct empirical claims by STA: (i) there is a *formateur* bonus; (ii) theoretical voting weights should be used rather than raw seat shares when predicting parties' portfolio allocations (STA: 993). Warwick and Druckman (2006) found the following in relation to these claims. First, in line with Figure 3, the *formateur* bonus predicted by BF-style models

[6] The vector $\boldsymbol{M}$ of MIWs, m_i, for each party i is defined as the smallest set of integers that generates, for a given winning quota Q_l, the same set of winning coalitions $\boldsymbol{C}$ as does the raw seat vector $\boldsymbol{L}$. Clearly, $\boldsymbol{M}$ is associated with a new winning quota, Q_m. An implication of this definition is that, for a (dummy) party d that is an essential member of no winning coalition, $m_d = 0$.

largely disappears once analyses are controlled for the fact that *formateur* parties tend strongly to be larger than non-*formateur* parties; STA did not control for this. Second, raw seat shares are *better* than MIWs in predicting portfolio allocations, when both are included in the same statistical analysis; STA left this as "a matter for future study" (Ansolabehere et al. 2005: 558). Warwick and Druckman's results are striking because they also measure the varying salience of different cabinet portfolios and find that portfolio payoffs remain proportional to legislative seat shares, taking account of the fact that some portfolios are worth more than others. As we now show, however, empirical tests of BF-style bargaining models face far a more serious problem than this, causing us to question fundamental BF-style modeling assumptions.

Endogenous *formateur* coding

The grounding assumption of *BF*-style bargaining models is an *exogenous* mechanism that first selects a unique *formateur* and then reveals this as common knowledge to all agents. We very rarely observe this explicitly in the real world, with the result that the *formateur* status of each agent in a real setting is typically both ambiguous and difficult to observe in primary sources. However, empirical analyses of BF-style bargaining over government formation fundamentally require coding the "*formateur* status" of each political party, as observed at the start of the bargaining process. The data on *formateur* status that form the basis of STA's empirical, as well as much other published work on this matter, were supplied by Warwick (Ansolabehere et al., 2005: 556). Consulting Warwick and Druckman (2001: 634), we find that *formateur* status was coded from *Keesing's Contemporary Archives*. Consider, however, the following entry in *Keesing's*, describing the formation of a German government in 2005. Crucially, this deals with events leading *up to but not including* the eventual formation of a government. It is thus a description of legislative bargaining, taken from the primary source in this field, but one that does not use the benefit of hindsight about the eventual outcome of the process under analysis:

> After the results were declared, Schröder controversially claimed that he was the victor because the SPD remained the largest single party, discounting the fact that the CDU and the CSU formed a single group in the Bundestag. Merkel responded that, as the leader of the largest parliamentary group, she had the right to head a new government. However, talks between her and the Greens on Sept. 23 on the formation of a "Jamaica" majority coalition – named after the black (CDU/CSU), yellow (FDP), and green colors of the Jamaican flag – quickly failed. At the same time, the FDP maintained its refusal to enter a "traffic light" coalition with the SPD and the Greens. The only viable option for a majority government, therefore, was a "grand coalition" of the CDU/CSU and the SPD, although at end-September Merkel and Schröder were both still insisting that they should be Chancellor.[7]

Who, on this basis, should be coded as the exogenously determined common knowledge *formateur*? The answer is far from clear and this problem is generic. *Keesing's* almost never contains statements of the form "… after the September election in X, the *formateur* was Y". Primary sources contain discursive accounts of contemporary events such as the one quoted above. These discursive accounts must be read by a human expert who then generates a binary variable for each party by coding its *formateur* status. Table 1, generated from the STA replication dataset using the same case universe as their published results, shows the relationship between a party's coded *formateur* status and whether or not it held the position of Prime Minister (PM) in the eventual government. The pattern is as startling as any we ever see in the

[7] *Keesing's Record of World Events*, Vol.51, 2005 (September) – Europe - Germany

social sciences, strongly suggesting that row and column variables measure precisely the same thing. This in turn raises the possibility that *formateur* status was coded, not as an *exogenous* independent variable but, *endogenously*, on the basis of whether or not the party took the PM position at the end of the government formation process. "Exogenously" determined *formateur* status and the endogenous control over the PM position, while theoretically distinct, are observationally identical in these data. While no written coding instructions survive, personal communication with Warwick confirmed that eventual control of the PM position was the default criterion for coding *formateur* status, which explains the remarkable pattern in Table 1.

Table 1: Formateur status and eventual control of PM position[8]

	Party controls eventual PM?		*Total*
	No	Yes	
Party is:			
Non-*formateur?*	1369	0	1369
Formateur?	1	249	250
Total	1370	249	1619

This has two crucial consequences. First, at a fundamental methodological level, the key "independent" variable in this dataset – and this is the main dataset that has been used to evaluate BF-style bargaining models using field data – was endogenously coded in light of the very effect it is claimed to predict. This negates the validity of any causal inference drawn from the empirical findings. Second, the same variable appears on both sides of the regression equations estimated both by STA and by Warwick and Druckman (2006). *Formateur* status is the model's key independent variable; the very same thing, in the guise of control over the PM position, is part of the dependent variable, the share of cabinet positions. The effects of doing this are exaggerated when, as in some of the STA (2005: 933) regressions, the impact of the PM position on the dependent variable is weighed three times more highly than any other cabinet post.

Table 2 replicates (in models A-C) the core results published by STA and Warwick and Druckman (2006), and then (in models D-F) corrects the endogeneity problem in models A-C by subtracting the PM position from the dependent variable. Model A perfectly retrieves STA's published result (Ansolabehere et al., 2005: 557). The significant positive coefficient on the *formateur* dummy is what is used by STA to infer that BF-style *formateur* models are effective at predicting portfolio payoffs. Models B and C retrieve Warwick and Druckman's published findings, using the STA replication dataset. Controlling for raw legislative seat share dramatically reduces the effects of both theoretical voting weight and *formateur* status on

[8] The single off-diagonal case arises from the Ciampi 1 government, forming in Italy in 1993, where the PM is described as a "technician". The number of cases is larger than that in STA's published regressions because the regressions include only parties in government, while Table 2 includes all parties in the relevant legislatures.

portfolio payoffs (Warwick and Druckman (2006: 654).[9] Model C confines the case universe to legislatures with five or fewer parties, in which MIWs and seat shares are not highly correlated. In this setting, the *formateur* effect loses statistical significance, although this is to a large extent the result of reducing the number of cases. This led Warwick and Druckman to infer that field data on portfolio allocation do not allow us to infer a significant *formateur* effect once we control for the fact that *formateur* parties tend strongly to be large.[10] Models D-F re-estimate models A-C, subtracting the PM position from the dependent variable, and thereby confining it to one side of the relevant regression equations. While the other regression coefficients are robust to this change, the coefficient for *formateur* status is now effectively zero in all models. This allows us to infer that published empirical conclusions about the *formateur* effect in field data depend crucially on the endogenous coding of *formateur* status. STA's empirical findings are entirely driven by the fact that the *formateur* is also invariably PM.

Table 2: Portfolio shares, voting weights, formateur status and legislative seat shares

	A: STA Table 3	***B: All govts***	***C: ≤5 parties***	***D: STA Table 3***	***E: All govts***	***F: ≤5 parties***
Dependent variable	*Party proportion of cabinet portfolios*			*As models A-C minus PM*		
Formateur status	0.15** (0.05)	0.07** (0.02)	0.09 (0.04)	0.08 (0.05)	0.00 (0.02)	0.00 (0.05)
Share of MIW in parliament	1.12** (0.13)	0.26 (0.16)	-0.08 (0.19)	1.20** (0.10)	0.27 (0.18)	-0.16 (0.21)
Share of seats in parliament		0.94** (0.16)	0.88** (0.16)		1.02** (0.18)	1.00** (0.18)
Constant	0.07** (0.02)	0.08** (0.02)	0.18** (0.04)	0.08** (0.02)	0.08** (0.08)	0.21** (0.05)
R^2	0.72	0.81	0.71	0.64	0.76	0.60
No. of observations	680	680	197	680	680	197

*Data source: Replication dataset for STA. A case is a party-in-government. STA replication code was used to generate Model A. Models B – F were generated by substituting variables and constraints in STA replication code. Figures in parenthesis are robust standard errors, clustered by country (as in STA). ** = statistically significant at 0.01 level; * = statistically significant at 0.05 level*

[9] Following STA, regressions reported in Table 2 do not weight portfolio payoffs by salience. Warwick and Druckman find the voting weight coefficient remains significant. This is because they use the more tolerant approach of clustering standard errors by government, as opposed to the STA approach of clustering by country.

[10] *Formateur* parties in the STA case universe have a mean seat share of 0.344, non-*formateur* parties of 0.118, a difference of means statistically significant at well beyond the 0.0001 level.

Setting out to address the problem of the endogenously coded "independent" variable at the heart of empirical tests of BF-style bargaining models, we made sustained and determined efforts, using *Keesings*, to generate a new set of *formateur* codings that do *not* make use of the knowledge of the government that eventually formed, using only reports that relate to events prior to government formation. We have concluded unambiguously that this is simply not possible – that a coding of exogenous *formateur* status cannot be derived without using information about the government that eventually formed, and thus that it is not possible to observe exogenous *formateur* status in primary data sources. This in itself does not imply that the core BF assumption is wrong; party leaders just might all have the same subliminal "knowledge" of what Nature is telling them about the randomly selected *formateur* without this ever being left on the record. But it does have the scientifically crucial implication that BF-style models are not testable using a variable for exogenous *formateur* status that is coded from historical sources.

4. THEORETICAL PROBLEMS WITH THE STA PROOF STRATEGY

The severe empirical problems identified above lead to the conclusion that something may be wrong with the STA model, as applied to government formation in parliamentary democracies. Two broad possibilities arise: that the basic BF-style modeling assumptions are wrong in the setting of government formation; that these general assumptions are valid but there are specific problems with the STA proofs that attempt to show that continuation values are proportional to voting weights. We consider the STA proof strategy here and return to BF-style modeling assumptions in our conclusions.

STA's results derived for "voting weights" in general, not MIWs in particular

A striking feature of STA's core propositions is that neither the definition of voting weights and party types, nor any explicit feature of the relevant proofs, constrains voting weights to be "theoretical voting weights" in general or MIWs in particular. The crucial definition of voting weights (STA: 984) constrains these only to be positive integers – true for both raw seat totals and MIWs. No proof deployed by STA uses any specific feature of MIWs or any other theoretical voting weight; these proofs can equally be read as if "weights" are raw seat totals. STA only introduce MIWs only *after all core results have been proved*, at which point they say that "*in what follows* we will use minimum integer weights" (STA: 988, emphasis added).

Puzzlingly, this implies that STA's propositions, if they are true, are simultaneously true in any given case for a range of different types of voting weight, including both raw seat totals and MIWs. This in itself implies axiomatically that the core STA propositions, as stated, must be false. Equilibrium continuation values, or indeed any other quantity, cannot simultaneously be proportional to two different sets of weights that are themselves not proportional to each other.[11] Thus, if we take party weights as MIWs, STA appear to have proved the propositions they set out to test empirically, which they contrast with Gamson's Law. If we take party weights as raw seat totals, they appear to have proved Gamson's Law. Something is seriously wrong.

---

[11] It is of course possible that STA were *implicitly* referring to MIWs when using the notion of voting weights in their proofs. But the fact remains that no feature of their proofs uses any specific property of MIWs, as opposed to any other type of voting weight such as raw seat totals, that agents might have in mind when bargaining over government formation

Notation

Before going further, it is useful to introduce some notation (following STA, 2005). Assume a legislature N where w is the total number of votes in N, and Q_m is the number of votes needed to form a winning coalition (usually, Q_m is a simple majority). Each legislator i is of type $t \in T$, with all legislators of the same type having the same voting weight, w_t. Let $t(i)$ denote the type of i, n_t the number of legislators of type t present in N, and v_t the continuation value of a legislator of type t. For a *formateur* i of type t and a proposed coalition C, $\underline{v}_t$ is the minimum total price paid by the *formateur* to its partners in C; i.e., $\underline{v}_t = min\ v(C \backslash i)$.[12]

Difficulties generated by non-homogenous voting games

The problem that STA appear to have simultaneously proved two contradictory propositions arises because they use a very distinctive "replication" technique to prove core propositions. This technique is adopted to deal with the problem of "non-homogenous games"; these are games, typically arising in legislatures with five or more parties, in which all MWCs do not have the same aggregate weight. A recognized proposer in a non-homogenous game must choose between MWCs of different weights, while some parties may be members of *no* smallest-weight MWC, generating significant ambiguities for analyzing BF-style bargaining models.

Consider the non-homogenous majority rule game (4, 3, 3, 2, 2), for which $Q_m = 8$. If the leader of the largest party is recognized as *formateur*, s/he can choose as partners: the two small parties (and an MWC of aggregate weight 8); one small and one medium party (and an MWC of aggregate weight 9); the two medium parties (and an MWC of aggregate weight 10). Does s/he see all these MWCs as equivalent, or see MWCs with different weights as being "different"? If s/he sees them as equivalent, then continuation values of the medium and small parties must be the same, since these parties are perfect substitutes for each other as partners for the largest party. A simple extension of this argument gives all parties equal continuation values, of 1/5.[13] If s/he has read STA's papers and expects continuation values to be proportional to weights, then s/he will see coalitions with the two smallest parties as the "cheapest" alternative yielding the highest retained surplus, and will strictly prefer the two small parties as coalition partners when recognized as *formateur*.

If recognized *formateurs* do believe continuation values are proportional to weights and strictly prefer the smallest MWCs however, we show in the Appendix for this case that this in turn implies continuation values non-monotonic in weights. Thus STA's core propositions *are not equilibrium beliefs* for recognized *formateurs* in the case of this non-homogenous game, whatever one believes about which parties the *formateur* prefers. In the first case, the continuation values are equal regardless of party weights; in the second case, they are monotonically *decreasing* in party weights. Given that STA set out to show analytically that continuation values must be proportional to party weights, this raises difficulties for their proofs. STA's computational algorithm implementing their model (discussed in the Appendix) generates

[12] We retain STA's notation here, using w for voting weight and $\underline{w}$ for the quota, because it is very important to keep in mind that STA do not, in their definition of voting weights, distinguish between seat shares and minimum integer weights. This is obscured by the use only of examples expressed in MIW format. This distinction will become very significant below. Thus we adopt STA's usage of w_i when the distinction between l_i and m_i has been left undetermined.

[13] The answer of 1/5 for all parties is what one would arrive at following the approach outlined by Baron-Ferejohn and assuming uniform probabilities for recognition.

the same non-monotonic continuation values we find in this case; computational and analytical implementations of the STA model generate different results for non-homogenous games.

For homogenous games, there is no problem in assuming that *formateurs* evaluate all MWCs (which by definition have the same aggregate weight) as identical, regardless of the individual weights of the parties that comprise them. There is no ambiguity in equilibrium beliefs about the continuation values of potential coalition partners, since only parties with the same weight are perfect substitutes for each other in MWCs. As we have illustrated with the above example, however, there is considerable unresolved ambiguity in non-homogenous games about what recognized *formateurs* might believe about MWCs with different aggregate weights. On different assumptions, parties with different weights may, or may not, be seen as perfect substitutes for each other in MWCs of which the recognized *formateur* is a member.

"Replicated" weighted voted games

Maria Montero has independently proved that continuation values in BF-style bargaining games are proportional to voting weights, assuming recognition probabilities proportional to voting weights and *constraining results to strong*[14] *homogenous games* (Montero 2006). Seeking proofs that extend to non-strong and non-homogenous as well as strong homogenous games, STA resort to analyzing what they call *replicated* voting games:

> We address this problem by examining the behavior of "replicated" voting games. That is, we examine equilibrium strategies as the number of players of each type is multiplied by some positive integer, $r \in \mathbf{Z}_x$. The basic game described above has $r = 1$, and a game with r replications has rn players, a total weight of rw, and a threshold for victory of $r\underline{w}$. We show that the effect of nonhomogeneity becomes small as r increases, thus allowing us to derive some general results. (STA: 984-985)

Note that STA do not analyze *repeated* games, but what they call *replicated* games, in which the number of players of each type is multiplied by some positive integer r, yielding a new game with many more players. All core propositions are proved for a "suitably chosen" r – from a range of integer values that has no effective upper bound (STA: 999). The replication device, while introduced to deal with non-homogenous games, is then used in all proofs, which make no distinction between homogenous and non-homogenous games, or between strong and non-strong games.

There are two serious problems with a proof strategy that involves replicating the game of interest r times. The first is that the replicated game typically has a completely different bargaining structure from that of its $r = 1$ version. This is because, while voting weights may be replicated, *the set of winning coalitions is not*. Completely new types of coalition emerge in replicated games, while other types of coalition may disappear on replication. Thus even the

[14] A strong game is defined as one in which the complement of every losing coalition is winning, and thus that there are no pairs of blocking coalitions. The complications posed by non-strong games for BF style bargaining models are: the need to model what happens in the event of blocking coalitions; the fact that there may be pairs of parties that are never found in the same MWC – impossible for strong games. Consider for example the non-strong majority voting game (3, 2, 2, 1) for which $Q_m = 5$. The largest and smallest parties share membership of no MWC. (We thank Maria Montero for this point and example.) Many published bargaining models implicitly assume strong games by assuming a simple majority quota and an odd number of legislators. Non-strong games are common in real legislatures. For the record, 132 of the 329 legislatures analyzed by STA in their published results generated non-strong games.

strong homogenous (1, 1, 1) game investigated by BF and many subsequent authors becomes a radically different *non-strong* game (1, 1, 1, 1, 1, 1) when $r = 2$. It is also easy to see why the "effect of non-homogeneity becomes small as r increases" by considering the strong non-homogenous game (2, 2, 2, 1, 1, 1) with $Q_m = 5$. The $r = 2$ version of this game is (2, 2, 2, 2, 2, 2, 1, 1, 1, 1, 1, 1) with $Q_m = 10$, and is a non-strong *homogenous* game. Setting $r = 2$ in this case does not so much "reduce the effects of non-homogeneity" as it defines a completely new *homogenous* game using the same agent types.[15] More generally, it is always the case that when r is even, replicating a strong game turns it into a non-strong game with blocking coalitions.

There is no result in the STA paper that demonstrates that different legislatures formed by different values of r are in an equivalence class. Nor is such a result possible. The winning coalitions generated and thus the likelihood of different coalitions forming are quite different, and these are not the most severe problems. Recognition probabilities are not constant across replicated legislatures, nor are continuation values. Moreover, the expected benefit for the *formateur* changes as the game is replicated. In the simple (1, 1, 1) game, the *formateur* retains 2/3 of the dollar in the case where $r = 1$, 1/2 when r = 2, and 8/15 when $r = 3$.[16] For different r values, it is clear that monotonicity does not even hold as r increases. Valid logical inferences about the $r = 1$ game which is the actual subject of interest cannot be drawn from a hypothetical new game in which $r > 1$ (introduced to facilitate the analysis), since this is a completely different game.

The second problem arising from STA's use of replicated games shows us why it is possible to "prove" propositions that are simultaneously consistent with different things – continuation values simultaneously proportional to ***L*** and ***M***, for example. This follows from Lemma 1 (STA 996-7), used in two subsequent lemmas (STA 998-9), with the set of three lemmas then used repeatedly in proofs of the main propositions. This lemma deals with a crucial quantity for BF-style bargaining models, the "cheapest" offer a recognized *formateur* can make to partners in winning coalitions. STA label this quantity for a type t agent as $\underline{v}_t$ and their proofs depend on setting an upper bound on this quantity. The lemma states that, in a stationary equilibrium and for any real $\varepsilon > 0$, there exists a finite r_ε such that for any $t \in T$ and $r \geq r_\varepsilon$, it follows that $\underline{v}_t \leq (r Q_m - w_t)/(rw) + \varepsilon$.[17] Setting ε arbitrarily small and rearranging, this gives us the result that it is possible to find a large enough r such that $\underline{v}_t \leq Q_m /w - w_t/rw$.

Clearly, as $r \to \infty$, $\underline{v}_t$ becomes arbitrarily close to a constant, Q_m /w, for any w_t and approaches ½ for simple majority games. Equally clearly, "suitable" values of r can be chosen to set very different upper bounds on $\underline{v}_t$. This shows us why the core STA propositions can be simultaneously true for different sets of weights. The upper bound on $\underline{v}_t$ is determined solely by the expression w_t/rw in the limit. The same upper bounds can be derived for different values of w_t by choosing "suitably different" values of r.

---

[15] The list of examples could easily be extended. Consider a classic homogenous majority rule "apex" game such as (2, 1, 1, 1). The $r = 2$ version of this games is (2, 2, 1, 1, 1, 1, 1, 1) and is clearly not an apex game. More generally it is easy to see that and that any apex game, by definition, ceases to be an apex game on replication.

[16] Again, this assumes MWCs are used. Using STA's computational algorithm to calculate continuation values in non-homogenous game complicates matters, since the assumption of smallest-weight winning coalitions can result in continuation values that are non-monotonic in r.

[17] STA use the notation $\underline{w}$ in place of $Q_{m.}$

5. WHAT IS TO BE DONE?

We have reached an *impasse*. The canonical non-cooperative model of bargaining in legislatures has no empirical support in field data on portfolio allocation, correcting for problems with the endogenous coding of its key independent variable. The recent *n*-party extension of this model by STA relies on a proof strategy involving replicated games, which allows “rival” propositions to be proved simultaneously. We believe these problems to be so serious that they imply a need to reconsider bargaining in *n*-party legislatures from a fundamental perspective, rather than to tinker under the hood of current models. This is a crucial challenge for all scholars who are interested, not in bargaining models *per se*, but either in comparative political analysis based on rigorous theoretical foundations or, more generally, in the scientific testing of theoretical bargaining models in real world settings. Confining our attention to bargaining in *n*-party legislatures *over government formation in parliamentary democracies*, we now map a way forward within this intellectual territory. Seeing a large part of the problem as arising from the deep empirical implausibility of existing modeling assumptions, move forward by stating some general empirical “truths” about government formation that we take to be self-evident and uncontroversial, in an attempt develop a set of realistic premises for models of bargaining over government formation.

1. Gamson’s ***Law*** *is true.* Many empirical analyses over the past thirty years or so, from Browne and Franklin (1973) to Warwick and Druckman (2006), have found GL to be true in field data on government formation. This empirical relationship is extraordinarily robust to details of measurement and estimation – indeed it is one of the most robust relationships uncovered and repeatedly investigated by an entire generation of political scientists. Two points are salient in this context. First, GL is so pervasive and easily observed in operation that it may be usefully treated both as common knowledge and as an equilibrium belief about how cabinet seats are likely to be divided. Indeed, for this reason and somewhat unusually in the social sciences, we can be confident in asserting that, independent of any particular bargaining model, GL does *empirically* characterize the equilibrium in the real world portfolio allocation game. Second, GL holds equally in types of coalition that look very different from each other, viewed from different theoretical perspectives. Thus, as we have seen, most non-cooperative bargaining models (incorrectly) predict the formation only of MWCs. However, results available from the authors and easy to generate from the STA replication dataset show that the classic GL regressions hold for: all coalition governments; MWC governments only; surplus majority coalitions only; minority coalitions only. GL thus applies in classes of commonly-occurring coalition governments that cannot be explained by conventional bargaining models.

2. Gamson’s ***Model*** *is false.* Gamson’s informal model of coalition bargaining does not *theoretically* characterize government formation. Since payoff shares are proportional to legislative seat shares, Gamson’s model predicts that the government will comprise the majority coalition with the smallest aggregate seat share. This proposition has been known to be empirically false for over thirty years, since the first empirical investigations of formal models of government formation (Taylor and Laver 1973). Recently, Fréchette *et al*. (2005a,b), analyzing experimental data from the laboratory and nesting predictions of government formation within predictions of portfolio payoffs, show convincingly that Gamson’s model, seen as a joint model of government formation and portfolio allocation, predicts the wrong governments. While Gamson's Law is empirically true, it is not true for the reasons put forward by Gamson.

3. BF-style bargaining models are false (in the context of government formation). We have shown above that there are both empirical and theoretical problems with *n*-party bargaining models in the BF tradition, applied to government formation.

4. Models predicting only MWC governments are false. MWC governments are very common in the real world, whether these are minority governments controlling less than a majority of legislators or "surplus" majority governments with more than the minimally necessary number of parties. The source of MWC predictions is often the assumption of a constant sum bargaining game with a reversion point normalized to zero for all agents.

5. There is always a status quo (SQ) incumbent government. This is specified in all written constitutions of which we are aware, which typically take great care to prevent a situation is which there is no constitutionally recognized government. The incumbent government remains in place, even if as a *gouvernement démissioné* that has been defeated in the legislature or has resigned, until replaced by an alternative. This is axiomatically true rather than a "mere" modeling assumption, with the consequence that some agents are almost certain to value SQ more than others. The government formation game, even when dealing only with the distribution of portfolios and before thinking at all about policy, is not constant sum.

6. Different assumed bargaining protocols produce very different theoretical predictions. BF-style models, with their implied *formateur* bonus, are fundamentally characterized by an assumed protocol with an exogenous random recognition rule. In our view this is a substantively unrealistic modeling assumption that plays a large part in the failure of BF-style models to explain field data. There is an increasing profusion of assumed bargaining protocols in the literature, and an effectively unlimited number of possible protocols that could be assumed, each generating different equilibrium predictions about the distribution of a fixed bundle of perquisites. Morelli's demand bargaining protocol implies payoffs proportional to voting weights and no *formateur* bonus (Morelli, 1999). A model of three-party legislatures designed by Austen-Smith and Banks assumes *formateurs* to be deterministically selected in strict size order, rather than chosen by an exogenous random mechanism, and predicts coalitions between the largest and smallest party (Austen-Smith and Banks 1988). Work by Carroll and Cox extends the bargaining timeline to include the possibility of pre-electoral coalitions and predicts GL payoffs when pre-electoral coalitions form (Carroll and Cox 2006). Recent economic modeling within the tradition of the "Nash Project", albeit not in the substantive context of government formation, describes plausible bargaining protocols for non-cooperative bargaining games that predict payoffs proportional to either the Shapley value or the Nash bargaining solution (Gul 1989; Hart and Mas-Colell 1996). The core intellectual issue is not whether or not we can choose a legislative bargaining protocol, from the infinite number of possibilities, that accounts for a robust empirical regularity we know to be true, in this case GL; we can be almost certain this is possible. In effect we can probably fit a model to almost anything we feel like by choosing the right bargaining protocol. Rather, what we now face is a "forest of protocols" problem and the key issue has become the *empirical plausibility of the bargaining protocol* we should assume in a given substantive setting.

In the context of government formation, we re-emphasize the universal constitutional requirement that there is always an incumbent government. Given this, we consider particularly plausible a bargaining protocol under which, unless this is explicitly contradicted by the constitution, the *de facto* first *formateur* is the incumbent Prime Minister. In the spirit of the increasingly diverse family of non-cooperative bargaining models, this implies that, if the incumbent PM can find an acceptable offer to make to a set of parties that between them command the legislative majority, then this offer will be made and accepted and the incumbent government will remain in office. Changes in government arise, on this assumption, when no such equilibrium offer is available to the incumbent Prime Minister.

*7. Legislative parties are not unitary actors***.** They are in effect "political clubs" that control membership *but not exit*, and allocate club resources between members according to some internal decision party rule. Party legislators always have an outside option and, if the outside option offers higher expectations, they are free to leave. The right of elected legislators to act independently of their parties, furthermore, is typically a binding constitutional constraint. This is particularly important for models predicting the allocation of government portfolios, for the following two reasons.

8. Party "weights" in legislative voting games are, precisely, the raw numbers of legislators affiliating to each party. Party weights are not *proportional* to this number or *derived* from it – they ARE this number. It is constitutionally unambiguous that each legislative vote cast is cast by an individual legislator. There is a *de facto* weighted voting game in a legislature only to the extent a party leader can ensure the disciplined behavior of all party legislators. Since there is always a constitutionally protected exit option for any legislator, *party weights in legislative voting games are endogenous*, not exogenous primitives, and must themselves be in equilibrium.

9. The "portfolio payoffs" from legislative bargaining comprise an integer number of senior party politicians who will sit at the cabinet table. Unlike many possible payoffs of the political game, seats at the cabinet table are: indivisible; very scarce and intensely valued by politicians; easily observed and counted. Portfolio payoffs are absolutely not, as assumed by most models, a real number consumed instantly and entirely by the party leader. It is simply inconceivable that a party leader would return from government formation negotiations and tell party colleagues "I've won eight cabinet portfolios and I'm keeping them all for myself". Thus portfolio payoffs to party leaders are in effect a small integer number of seats/votes at the cabinet table, available for distribution by the leader between party members on the basis of some internal party game. We consider this to be a crucial and hitherto ignored feature of the portfolio allocation problem, particularly when considered in conjunction with the problem of party discipline.

It seems uncontroversial to assume that a significant element in the utility function of almost any senior politician concerns expectations of holding high office and, in particular in parliamentary systems, a seat at the cabinet table. And it also seems uncontroversial to assume that an important motivation for any party leader who wants to remain party leader is to maintain the support of senior party colleagues by satisfying their expectations. Portfolio allocation thereby becomes a major feature of the party discipline game; party leaders who win a less than proportional share of cabinet portfolios, in a Gamsonian sense, are more likely to disappoint senior party members.

10. Policy is important in government formation. Thus far we have confined our attention to the allocation of cabinet portfolios, seen as a fixed bundle of perquisites. However, as a matter of fact, government formation between legislative parties almost invariably involves negotiating and publishing a joint policy program, agreed by all parties in the government. Furthermore, cabinet portfolios almost invariably define "policy jurisdictions" which at the very least give the incumbent minister some agenda power in predetermined policy areas (Laver and Shepsle 1996). The empirical sequence of events in almost all government formation negotiations of which we are aware is that the joint policy program is agreed first, and portfolio allocation between parties follows. Theoretically, this seems likely to happen because cabinet portfolios have policy jurisdictions with substantial consequences for downstream policy implementation, so that the allocation of cabinet portfolios logically follows rather than precedes policy bargaining. For this reason, it does not seem productive to model bargaining over cabinet portfolios as if this was independent of bargaining over the joint policy program. While it would be theoretically satisfying to endogenize this bargaining sequence, a useful interim assumption, one that is also behaviorally realistic, is that government formation negotiations first involve bargaining over a joint policy program, and then involve bargaining over the distribution of cabinet portfolios.

Clearly, any theoretically rigorous new model of bargaining in legislatures will be parsimonious and could not possibly assimilate all of the ten general statements we have just made. So we conclude by prioritizing these.

We regard statements 1-4 above as settled facts about portfolio allocation, to the extent that it is ever possible for political scientists to establish settled facts. Gamson's Law is true; Gamson's model is false; there is no observable *formateur* bonus; raw seat totals better predict parties' portfolio shares than theoretical bargaining weights; models predicting only MWCs are flawed empirically. In our view, future theorists who ignore these findings are on a fool's errand if their intention it to develop empirically realistic, and thus scientifically testable, models.

We regard the strong constitutional regularity noted in 5 above – there is always an incumbent government – as a very promising starting point for new theoretical work. Such strong constitutional regularities are rare; to ignore this regularity, which extends to most government formation settings, seems perverse. This in turn suggests an assumed bargaining protocol that gives first *formateur* status to the incumbent PM. Following this, our strong empirical intuition is that second *formateur* status passes, in practice, to the leader of the largest opposition party. This is however no more than an informal intuition and to be more systematic we will run smack into the wall, noted above, that the systematic empirical coding of *formateurs* generates dire endogeneity problems.

Taking one further step along to road to substantive realism, we find it impossible to ignore the plain fact that cabinet portfolios are distributed between senior party politicians rather than entirely consumed by the party leader – whether as ways to ensure the loyalty of senior politicians or, in the spirit of Laver and Shepsle (1996), as ways to underwrite particular policy deals. While the assumption that parties are unitary actors has served the profession well in certain contexts, it seems unrealistic to the point of obtuseness to assume this when the payoff under investigation is, by definition, a sack of goodies that is distributed *inside* political parties.

APPENDIX: EXAMPLE DEMONSTRATING POSSIBILITY OF CONTINUATION VALUES NON-MONOTONIC IN VOTING WEIGHTS FOR A NON-HOMOGENOUS GAME

Consider a majority rule legislature with $\boldsymbol{M} = (4, 3, 3, 2, 2)$ and $Q_m = 8$. There are three party types: $T = \{1, 2, 3\}$; $w_1 = 4$; $w_2 = 3$; $w_3 = 2$. Let μ be the set of MWCs and $\mu_{\mathbf{min}}$ **be** the set of MWCs with smallest aggregate weight. There are four coalition types in μ: $\tau_1 = \{t_1, t_3, t_3\}$; $\tau_2 = \{t_2, t_2, t_3\}$; $\tau_3 = \{t_1, t_2, t_3\}$; $\tau_4 = \{t_1, t_2, t_2\}$. However, only τ_1 and τ_2 are in $\mu_{\mathbf{min}}$. The smallest party is a member of all coalition types in $\mu_{\mathbf{min}}$. STA's Propositions 2 and 3 are derived by solving the non-cooperative bargaining game for equilibrium continuation values. This game can be solved as a system of constraints, the solution approach adopted in the widely circulated computer program designed by STA to calculate MIWs and BF continuation values. This is the only program available to calculate MIWs in difficult situations. The first (budget) constraint is that the continuation values, v_t, of all five parties sum to unity:

$$v_1 + 2v_2 + 2v_3 = 1$$

The second (coalition) constraint is that the probabilities of any given *formateur* proposing one of the set of possible coalitions sum to unity. Noting both that we consider other possibilities in the text and that there are many alternative assumptions, assume here that recognized *formateurs* have read STA's core propositions, thus expect continuation values to be proportional to voting weights, and thus strictly prefer to form what they believe to be the "cheapest" coalitions, with parties in $\mu_{\mathbf{min}}$, since the surplus s/he retains from other MWCs is in this event less. Thus the only coalitions that *formateurs* will propose, if recognized, will be τ_1 or τ_2. Within $\mu_{\mathbf{min}}$, the t_1 party belongs only to τ_1 and the t_2 parties belong only to τ_2. The t_3 parties are the only ones belonging to both τ_1 and τ_2, so the coalition constraints are quite simple. If a t_3 party is *formateur*, then the probability, c_1, that it proposes τ_1 and the probability, c_2, that it proposes τ_2 must sum to unity:

$$c_1 + c_2 = 1$$

The third (substitutability) constraint arises from the fact that, in equilibrium, if one potential coalition partner has a higher "price", v_i, than another, then the *formateur* strictly prefers the cheaper partner. In equilibrium, *formateurs* will be indifferent between sets of coalition partners with the same price. Only t_3 parties can chose between possible sets of partners, in coalitions in $\mu_{\mathbf{min}}$, yielding them the same retained surplus[18]:

$$v_1 + v_3 = 2v_2$$

The final constraint is that continuation values are subgame perfect. We must take into account the likelihood an agent will be chosen either as *formateur* or as a partner in another *formateur's*

[18] In words, a 2-vote *formateur* can replace two 3-vote parties with a 4-vote and a 2-vote party.

coalition. Assuming equal recognition probabilities for all agents, there are three continuation values[19]:

$$v_1 = \frac{1}{5}(1-2v_3) + \frac{2v_1}{5}\cdot(c_1)$$

$$v_2 = \frac{1}{5}(1-v_2-v_3) + \frac{2v_2}{5}\cdot(c_2) + \frac{v_2}{5}$$

$$v_3 = \frac{1}{5}(1-2v_2) + \frac{2v_3}{5} + \frac{v_3}{5}\cdot(c_1)$$

Solving this set of constraints produces the result:

$$v_1 = \frac{1}{8},\ v_2 = \frac{3}{16},\ v_3 = \frac{1}{4},\ c_1 = \frac{1}{2},\ c_2 = \frac{1}{2}$$

Contradiction of Propositions 2 and 3 in this case, furthermore, does not depend upon the inference from the STA model that only coalitions in $\mu_{\mathbf{min}}$ will be proposed. If the set of coalitions considered is extended to μ, then each party type has three possible coalition types to choose between if selected as *formateur*. Not surprisingly, predicted bargaining outcomes are now completely different. Solving the constraints in this setting shows that each agent now has a continuation value of 1/5, again contradicting STA's propositions, as each partner is preferred equally *regardless* of their voting weights. In both cases, "continuation values" are non-monotonic in voting weights. Indeed parties' solved "continuation values" are either monotonically *decreasing* in their MIWs or equal regardless of weights, not directly proportional to them. The STA calculator noted above generates the same non-monotonic BF "continuation values" as our first case above. What this shows, of course, is *not* that we actually predict continuation values to be monotonically decreasing in their MIWs, but that STA's core propositions – continuation values monotonic in weights – are *not* equilibrium beliefs for recognized *formateurs* in this non-homogenous game. This is not an isolated case; we have identified considerable set of cases in which direct computations contradict STA's core propositions, resulting in equilibrium continuation values that are non-monotonic in MIWs.[20] Given that STA-Strauss calculator also produces precisely the same contradictions, one must ask why STA's proof a computational algorithm are so systematically at odds

[19] In words, the 4-vote party has a 1/5 probability of being formateur, proposing the {4,2,2} coalition, and retaining the surplus of 10/14, and a 2/5 probability that a 2-vote party will be chosen and, with probability c_1, offer it 4/14. There is a 2/5 probability that a 3-vote party will be *formateur*, in which case the 4-vote party has zero probability of receiving an offer since 3- and 4-vote parties only share non-SWCs. Constraints on v_2 and v_3 have analogous interpretations.

[20] For example, in addition to {4,3,3,2,2}, the majority rule games: {5,4,3,2,2}, {3,3,2,2,2}, {7,6,5,2,2,2}, {9,6,5,4,2,2}, {6,5,4,3,2,2}, {9,5,5,3,2,2}, {7,6,5,4,2,2,2}, {8,7,6,5,4,2,2,2}

REFERENCES

Ansolabehere, Stephen, James M. Snyder, Aaron B. Strauss, and Michael M. Ting. 2005. Voting weights and formateur advantages in the formation of coalition goverments. *American Journal of Political Science* 49 (3):550-563.

Austen-Smith, David, and Jeffrey Banks. 1988. Elections, coalitions, and legislative outcomes. *American Political Science Review* 82:405-422.

Banzhaf, John. 1965. Weighted voting doesn't work: a mathematical analysis. *Rutgers Law Review* 19:317-341.

Baron, David, and John Ferejohn. 1989. Bargaining in legislatures. *American Political Science Review* 83 (4):1182-1202.

Browne, Eric, and Mark Franklin. 1973. Aspects of Coalition Payoffs in European Parliamentary Democracies. *American Political Science Review* 67:453-469.

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* 24:753-768.

Carroll, Royce, and Gary W. Cox. 2006. The logic of Gamson's Law: pre-election coalitions and portfolio allocations: Department of Political Science, University of California, San Diego.

Diermeier, Daniel, and Antonio Merlo. 2004. An empirical investigation of coalitional bargaining procedures. *Journal of Public Economics* 8:783-797.

Felsenthal, Dan S., and Moshe Machover. 2001. Myths and Meanings of Voting Power: Comments on a Symposium. *Journal of Theoretical Politics* 13 (1):81-97.

Gamson, William A. 1961. A Theory of Coalition Formation. *American Sociological Review* 26:373-382.

Gul, Faruk. 1989. Bargaining Foundations of the Shapley Value. *Econometrica* 57 (1):81-95.

Hart, Sergiu, and Andreu Mas-Colell. 1996. Bargaining and Value. *Econometrica* 64 (2):357-380.

Laver, Michael, and Norman Schofield. 1998. *Multiparty government: the politics of coalition in Europe*. Ann Arbor paperback ed. Ann Arbor, Mich.: University of Michigan Press.

Laver, Michael, and Kenneth A. Shepsle. 1996. *Making and breaking governments: cabinets and legislatures in parliamentary democracies*, *Political economy of institutions and decisions*. New York: Cambridge University Press.

Montero, Maria. 2006. Noncooperative foundations of the nucleolus in majority games. *Games and Economic Behavior* 54:380–397.

Montero, Maria, and Juan J. Vidal-Puga. 2005. Demand committment in legislative bargaining: School of Economics, University of Nottingham.

Morelli, Massimo. 1999. Demand competition and policy compromise in legsilative bargaining. *American Political Science Review* 93 (4):809-820.

Shapley, Lloyd, and Martin Shubik. 1954. A method for evaluating the distribution of power in a committee system. *American Political Science Review* 48 (3):787–792.

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

Taylor, Michael, and Michael Laver. 1973. Government coalitions in Western Europe. *European Journal of Political Research* 1:205-248.

Warwick, Paul V., and James N. Druckman. 2001. Portfolio salience and the proportionality of payoffs in coalition governments. *British Journal of Political Science* 31:627-649.

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* 45:635-665.