Coalition Agenda-Setting

David Fortunato
University of California, Merced
dfortunato@ucmerced.edu

Mariyana Angelova
Universität Mannheim
mariyana.angelova@uni-mannheim.de

Abstract

Building on the extant literature on coalition formation and policy-making as well as the literature on democratic responsiveness, we derive three models of coalition agenda-setting to understand how coalition cabinets divide the legislative agenda across member parties. Using expansive original data on several decades of legislative proposals and public opinion in seven parliamentary democracies we execute critical tests of these models. The data support a critical events model of agenda-setting where parties are able to leverage increases in popularity and proximity to opposition parties into greater control of the legislative agenda. These effects, however, are conditioned on institutional context.

The difficulties of multiparty governance, particularly as pertains to collective decision-making in areas where preferences are diffuse, are well documented in comparative politics. Coalition cabinets force different parties, with distinct support networks demanding diverse and sometimes incompatible policies, to compromise and cooperate in order to effectively govern. One of the most salient realities of this construction is the inability for any one party to act as a policy dictator, even within policy areas it controls as minister with portfolio. As Martin (2004) points out, each minister is endowed with complete negative agenda control, as she (with her considerable staff) determines the content of policy in her jurisdiction, but lacks complete positive agenda control, as the submission of her legislative proposals are subject to cabinet approval. That is, the cabinet, and therefore all member parties of the coalition, ultimately control the legislative agenda as a collective, thus effectively curtailing any one minister's (or party's) ability to effect policy change in a vacuum.

Because the legislative agenda is the gatekeeper between each minister and her policy goals, its construction is integral in determining the policy payoffs that each member of coalition enjoys as well as the overall shape of policy outcomes that a government will produce. Yet, while we understand generally that cabinets will prioritize policies on which they agree over policies on which they are divided (Martin 2004), we do not understand how agenda access is allocated *across individual members* of cabinet and whether or not that allocation is dynamic. More specifically, we do not understand whether governments are cooperative or competitive in their division of the agenda or whether governments react to shifts in the political landscape throughout their tenure. This is a shame as understanding how coalition partners divide access to the legislative agenda and whether that division is fluid is critical not only to understanding to how laws get made but also, and perhaps more importantly, whether governments are responsive to the electorate throughout their tenure.

In this manuscript we present answers to these open questions by deriving critical tests of three models of agenda formation that are suggested by the extant theoretical and empirical literature. The first, which we call the mandate model, holds that electoral outcomes (the confluence of expressed voter preferences and electoral institutions) are the ultimate determinant of policy outcomes

and suggests that government agenda should be more or less fixed from its initial organization (Martin 2004; McDonald, Mendes, and Budge 2004; Powell 2000). The second, the responsible government model, suggests that parties in government should be responsive to the preferences of citizens, regardless of whether or not these preferences impact the relative bargaining power of cabinet members, such that the agenda will conform to the shape of public opinion (Calvo 2007; Erickson, Wright, and McIver 1993; Page and Shapiro 1983). Finally, the critical events model holds that parties bargain over access to the legislative agenda in light of their future prospects for office and are therefore responsive to shifts in public opinion that translate into shifts in expected seat distributions (and therefore cabinet prospects) given a new election (Diermeier and Stevenson 2000; Lupia and Strøm 1995). We then present original data combining information on over 18,000 government proposals — spanning over 50 years, 40 cabinets, and 40 parties in 7 European democracies —with monthly public opinion data to execute our empirical tests. These data, the first of their kind in comparative legislative research, allow us to assess the relative fit of our three models and determine which, if any, provide a framework for understanding real world coalition policy-making.

The Legislative Agenda in Parliamentary Democracy

Over the past several decades, research on the United States Congress has been defined by a vigorous debate over who dominates the legislative agenda: log-rolling coalitions of "high demanders," the floor as a whole, or the majority party (Weingast and Marshall 1988; Gilligan and Krehbiel 1990; Cox and McCubbins 1992, respectively). The motivation for this debate is simple: when preferences are fixed, as they are nearly always assumed to be, the key to understanding policy outcomes is understanding the organization of the legislative agenda. That is, if we can predict who will get to initiate a bill, we can predict what the resulting policy outcome will be.

However, while some of these American theories have been exported to parliamentary contexts to analyze majority-minority or government-opposition voting patterns (e.g., Chandler et al. 2006; Cox et al. 2008), students of coalition politics showed comparatively little interest in the legislative agenda, preferring, instead, to study coalition formation (Laver and Shepsle 1994; Martin and

Stevenson 2001; Strøm 1990) or dissolution (Brown, Fendreis, and Gleiber 1984; Lupia and Strøm 1995; Warwick 1994). This focus on the formation and dissolution of governments was integral to the growth of our overall understanding of parliamentary democracy, but left the literature wanting for theoretical and empirical analyses of the actual business of coalition governance. Indeed, an assumption — typically implicit, but occasionally explicit — of this focus on formation and duration is that policy outcomes are determined by which parties are in cabinet and how long they remain there. But in order for this to be true, access to the legislative agenda would have to be divided equally across departments, allowing each minister to implement the same degree of policy change. We know this to be untrue.[1]

One of Martin's (2004) several contributions was to point out that one of a cabinet's most prized resources is time and that time is limited. Therefore coalitions are only able to pass so many proposals before their time in cabinet must come to an end (either naturally or prematurely). As such, perhaps the most important decision a coalition cabinet must make is how to organize its legislative agenda. It is not the composition of the cabinet alone, but the composition of the cabinet and the distribution of the legislative agenda across departments, that ultimately shapes policy outcomes.[2] Martins' conclusion is that multiparty cabinets prioritize bills on which coalition members broadly agree and bills in policy areas that are more salient to members, a conclusion suggesting the notion that agenda access is divided equally amongst departments is invalid.

Aside from Martin's (2004) conclusions regarding the prioritization of agreeable and important legislation and Bräuninger and Debus's (2009) analysis of when the cabinet will choose to legislate (rather than allowing individual members to take the lead) we know very little about how the agenda is set in parliamentary democracies, save that it is dominated by the cabinet (Cox 1987) and that, in general, that cabinet gets what the cabinet wants (Heller 2001; Saiegh 2009). As Martin and others

[1]See, for example, the descriptive statistics regarding the number of policy proposals from different ministries in Martin and Vanberg (2011).

[2]Of course, we now understand the pivotal role that the legislature plays in the review phase of policy-making as well (Martin and Vanberg 2005, 2011).

have pointed out, however, in multiparty governance, the norm in the majority of parliamentary democracies, understanding *what* the cabinet wants is the challenge and this is the broad question that we address here. Before we may consider what the literature implies about how coalitions go about deciding what they want at a given point in time, however, we must first consider the role that individual ministers play in the structuring of coalition policy-outcomes.

Ministerial Autonomy and Models of Coalition Policy-Making

Until recently, the preeminent model of coalition policy-making has been the ministerial autonomy model (Austen-Smith and Banks 1990; Laver and Shepsle 1990, 1996). In this model, ministers are allowed to act as dictators within the jurisdiction over which they hold portfolio. This autonomy is presumed for two reasons: first, ministers enjoy substantial informational advantages as they control their department and the wealth of resources, particularly the expertise of the civil servant network, therein; second, ministers (and their staff) are responsible for the crafting of policy proposals, which endows them with an enviable positive agenda power. Further, this presumption generally concords with our qualitative accounts of coalition policy-making, in which ministers are overwhelmingly characterized as dominant (Laver and Shepsle 1994).

More recently, the coalition compromise model has become the "industry standard" in the literature (Martin and Vanberg 2011; Powell 2000). The typical conception of this model is that ministers are similarly autonomous in crafting their policy proposals, however, the cabinet is able to limit individual ministers' overall policy authority by triggering various monitoring institutions, such as junior ministers, parliamentary committees, intra-coalition committees, etc., such that the final policy outcomes is likely to resemble some weighted aggregation of coalition policy preferences. This scrutiny of ministerial bills, however, is typically regarded as a worthwhile tradeoff for, even though the minister may be denied her most preferred policy, she is still able to signal her preferences and reap credit-claiming or competence benefits from the construction of initial proposal.

What is most important for our purposes is that both models state that ministers, assumed to be perfectly representative of their party, enjoy complete autonomy in the preparation of their

proposals. Indeed, the empirical tests of these competing models constructed in a recent article by Martin and Vanberg (2014), state this assumption explicitly and subsequently recover evidence supporting it — evidence that is echoed by nearly every other empirical examination of coalition policy-making (Bräuninger and Debus 2009; Martin 2004; Martin and Vanberg 2005; 2011).[3] Therefore, our examination of the legislative agenda also assumes that ministers are autonomous in the construction of their legislative proposals (and that each proposal is perfectly representative of their party's ideal policy outcomes) and will focus on which minister is allowed to submit a legislative proposal and when.

Three Models of Coalition Agenda-Setting

As discussed above, one of a cabinet's most precious resources, if not its most precious resource, is time. There are only so many days in the legislative session and so many hours in a legislative day for a cabinet to affect its policy program. As such, the allocation of its time, or distribution of access to the legislative agenda, is one of the most central points a government must address. But, the question of how a coalition cabinet allocates agenda access across its individual members is an open one comparative political research.

This is not to imply, however, that the literature is without guidance on the question. A survey of the extant literatures on democratic choice and accountability, democratic responsiveness, and coalition formation and duration suggest three potential models for how coalition cabinets may divide the legislative agenda across its members.

[3]There are exceptions, of course. For example, Goodhart (2013) finds that monetary policy seems to be shaped cabinets as whole (or cabinet leadership) rather than finance ministers. This study, however, only examines outcomes and not proposals, outcomes that are the result of not only the proposal, but of its reshaping in the legislative review phase of the policy-making process. Thus, studies such as this, and their findings, do not contradict our assumptions about the positive agenda power of minsters with portfolio.

The Mandate Model

In his seminal work, Powell (2000) describes elections as the instrument linking voters to policy outcomes. In Powell's view, parties come together post election to form a government whose policy outputs will reflect the preferences of their aggregated supporters. In practice, this means that a government's policy output will approximate the seat-weighted policy preferences of the parties entering cabinet. Indeed, Powell's own empirical estimations of government policy outputs, as well as those who have followed this approach (i.e., Golder and Stramski 2010; Martin and Vanberg 2014; McDonald, Mendes, and Budge 2004), employs this formula explicitly.

What does this model — conceiving a coalition government's policy output as the seat weighted aggregation of individual party preferences — imply for the legislative agenda? Given the assumption that ministers are autonomous when crafting their proposals and the argument that elections determine the direction government will take, this model of suggests that the distribution of the legislative agenda will conform to the relative weights of the parties in cabinet. That is, because elections function as legislative mandates, each party in government is endowed with their relative policy-making authority by the outcome of the last election (Powell 2000), therefore each party will be allowed to submit bills in proportion to the number of seats they bring to the coalition. In a manner of speaking, this approach allows the formalized will of the electorate to act as the final arbiter of agenda allocation. Because elections, and not merely estimated public opinion, endow parties with their legitimacy to govern, the allocation of agenda access across parties should remain static throughout the legislative term. Note that this is the agenda construction implied by the canonical model cabinet formation (Laver and Shepsle 1990, 1996), where ministers are assumed to be automatons within their jurisdiction, when portfolio allocations are roughly Gamsonian. We refer to this organization of the agenda as the *mandate model*, and it yields the following hypothesis:

- Mandate hypothesis: a party's share of the legislative agenda is determined by the number of seats it brings to coalition

One could consider the *coalition compromise model* a special case of the mandate model of

coalition policy-making (Martin 2004; Martin and Vanberg 2011; Thies 2001). In this view, policy outcomes should reflect the aggregated preferences of the cabinet's member parties, as in the general case, but Martin, Thies, and Vanberg are more explicit in their behavioral expectations for the government.[4] These more recent works suggest that the cabinet will use its negative agenda powers to prevent ideologically extreme (in reference to the location of the coalition compromise) ministers from flaunting the coalition agreement. Thus, the compromise model suggests agenda allocation should be determined by ideological proximity to the seat weighted aggregation of the coalition as a whole.

- Compromise hypothesis: a party's share of the legislative agenda is determined by its distance from the coalition compromise

The Responsible Government Model

The second model of coalition agenda-setting, which we refer to as the *responsible government model*, is suggested by the rich research on legislative responsiveness, typically focused on the American case (e.g., Bond et al. 2003; Erickson, Wright, and McIver 1993; Lax and Phillips 2012; Page and Shapiro 1983), but not limited to it (e.g., Calvo 2007). This long and well-known research, stretching from Downs (1957), generally argues that government decision-making should be responsive to public opinion, usually through an electoral mechanism. If parties are vote maximizers, then their behaviors should conform to the preferences of the public in order earn electoral support. The more normative approach, however, does not require parties to be electorally motivated per se, their behaviors reflect the preferences of their constituents in order to fulfill their role in the democratic process. In other words, parties *should* conform to public opinion, because it is the right thing to do. Both mechanisms of the responsible government model, unlike the mandate model, allow legislative focus to vary within legislative periods should the will of the public change.

[4]Recall that Powell's (2000) concern was not legislative behavior but the congruence of policy outcomes to the preferences of the electorate.

Much of the literature on the issue of policy responsiveness (or congruence), which is overwhelmingly empirical in its approach, does not differentiate between these mechanisms — the importance of the question of whether or not legislators or governments are responsive is self-evident and a sufficient motivation. Indeed, whether individual legislators or majority parties are responding to changes in public opinion because they are vote seekers or because they believe it is the right thing to do, has, in general, little impact on the shape of the actions they take (but see Bishin 2009). Thus, much of the empirical literature has not differentiated between normative and vote seeking motivations because it was largely irrelevant to predicting behaviors. This is not a luxury that we have, however, because the two mechanisms have different implications for how *coalition governments* may divide the legislative agenda amongst its members.

The reason is because the electoral fortunes of cabinets, though certainly correlated (Duch and Stevenson 2008), are not necessarily tied in all contexts. Therefore, a coalition that distributes agenda access across partners in response to public opinion out of normative concerns for democratic process will behave differently than a vote maximizing coalition. Consider a two party cabinet, normatively responsive to public opinion, where partner A is to the left of partner B. Should voters' preferences for policy shift to the left sometime during the legislative period, the cabinet should respond (given our assumptions about how ministers draft their proposals) by allocating a greater share of the agenda to the more left party A, therefore moving the cabinet's mean policy output to the left.

This strategy should not carry over to all cabinets composed of vote maximizing parties, however, only a small subset. The reason is because this strategy can only be an effective vote maximizer for the cabinet as a whole if the member parties have *formally tied* their electoral fortunes. In other words, for a cabinet composed of two parties who have not committed to a pre-electoral coalition for the next election, it is difficult to imagine a scenario in which one party would willingly cede a portion of its agenda privileges to its partner in hopes that it would benefit its own electoral fortunes. Though policies are made jointly, parties are held accountable individually — except when parties have agreed in advance to pool their electoral prospects. Under this circumstance, the coalition

may re-allocate agenda divisions to follow public opinion in hopes that the parties will collectively benefit in the following election. Without the pre-electoral pact, however, such a reallocation will, at best, benefit only the party or parties who receive an increases proportion of the agenda at the expense of the opposition. Alternatively, it may benefit one cabinet party at the expense of its partners — an occurrence the cabinet as a whole is likely to guard against. Thus, the responsible government model yields two distinct empirical implications, conditioned on the motivation.

- Normative responsible government hypothesis: a party's share of the legislative agenda is always increased by its popularity
- Vote-maximizing responsible government hypothesis: a party's share of the legislative agenda is increased by its popularity if the coalition has entered into a pre-electoral pact to reform

The Critical Events Model

Our third model, the *critical events model*, is motivated by the rich research on coalition bargaining and termination.[5] Importantly, we take an "event," as Lupia and Strøm (1995) and Diermeier and Stevenson (2000) did, to mean a shift in public opinion that alters the bargaining weights of coalition participants, rather than the stochastic occurrence of a crisis that may end cabinet (i.e., a war, economic collapse, scandal, etc.) as was the common understanding of earlier work (Brown et al. 1984). More specifically, "[e]vents are interpreted as common knowledge information about what would happen if parliament were dissolved and an election were held immediately" (Diermeier and Stevenson 2000: 628). This model views coalition governance as repeated negotiations between current and potential cabinet members where the status quo (current coalition) is evaluated against cabinet replacement and parliamentary dissolution. If we allow, however, that the status quo government may re-allocate its resources within the present configuration of member parties, then this theoretical framework, and its empirical implications, may be extended to understand the legislative agenda.

[5]See Laver (2003) for an excellent review of this substantial literature.

Diermeier and Stevenson's (2000) extension of the Lupia-Strøm (1995) model of cabinet termination, which assumes that the incumbent cabinet has the power to call elections, isolates two conditions for determining if and how the present coalition will terminate: (A) "there is some decisive coalition the prefers an election over the current government" and (B) all members of the standing "governing coalition prefer an election over the best offer they can receive in bargaining" (Diermeier and Stevenson 2000: 629). When A is satisfied, but B is not, the cabinet will terminate in replacement; when both conditions are satisfied, the cabinet will terminate in dissolution. Notice, however, that termination relies only on condition A. If we consider the status quo to also represent the present allocation of agenda access parties within the cabinet, then the model bears nuanced predictions for agenda-setting.

The central questions are: What makes an alternative coalition more appealing than the status quo for a particular cabinet party? And, on what grounds can a member of the incumbent cabinet make a credible treat to dissolve the cabinet? There are several answers. First, a party's popularity may increase. Consider a two-party majority cabinet composed of some parties A and B, where party C stands in opposition. Now consider an event, a shock to public opinion that increases the popularity of party A at the expense of party B such that if elections were called, B would expect to lose a portion of its seat share to party A thus increasing the relative strength of parties A and C, vis-á-vis, party B. Given this new information, A may impose a demand on B: readjust the agenda to reflect this shift in A's expected seat share or A will dissolve the cabinet. If party B accepts, then it remains in cabinet, albeit with a weakened share of the agenda. Should B refuse, A and C may dissolve the cabinet and move for new elections, formally increasing their strength relative to party B, or attempt to form a new cabinet directly. This yields the first critical events hypothesis:

- Critical events popularity hypothesis: A cabinet party's agenda access will increase with its popularity

The Diermeier and Stevenson model also implies that this relationship will be moderated by time, particularly the time remaining in the constitutional interelection period (CIEP). As constitutionally mandated elections approach, the relative value of maintaining the current cabinet decreases for all

parties. Thus, one may expect party B's willingness to acquiesce to A's demands over the agenda should also decline and the ability of parties to profit from their popularity should decrease in kind.

- Critical events popularity CIEP hypothesis: A cabinet party's agenda gains from popularity will decrease as the CIEP comes to a close

Further, the effects of popularity should be tempered by the larger electoral context, particularly whether or not the current cabinet have formally allied for the coming election. If the cabinet parties have agreed to pool their electoral fortunes and reform on the condition that they win enough seats, then threats to dissolve the present cabinet to form an alternative lose credibility. In a sense the parties have tied their hands by imposing external costs on themselves. If a pre-electoral coalition is announced and one of the parties proceeds to dissolve the cabinet because its partners refused some agenda demand, it is likely that voters would punish that party for not honoring its agreement. At the very least, voters would look askance at future promises made by that party.

- Critical events popularity PEC hypothesis: A cabinet party's agenda gains from popularity will decrease in the presence of a pre-electoral coalition agreement

Alternative coalitions are made attractive by more than seatshare. Indeed, there is an extensive literature on cabinet formation that we may look to for empirical guidance on what makes some coalitions more attractive than others. What is perhaps the most salient and familiar finding in this well-known literature is the critical role of ideological compactness on coalition formation. Coalitions form to govern, after all, and the extent to which parties agree on important issues will determine the extent to which the cabinet can be effective at making policy and, more importantly for present purposes, the extent to which parties would be willing to coalesce. Thus, a particular cabinet member can make a more credible threat to dissolve the current government and reform an alternative, if that alternative is ideologically compatible. Returning to our stylized three party example above, suppose that the parties are ordered on the ideological continuum such that $A < B < C$. In this case, where A and B compose the current cabinet, B may make a more credible

threat to dissolve the current cabinet than A because B is better suited to coalesce with opposition party C. Thus, even when popularity is accounted for, we would expect cabinet partners more ideologically proximate to the opposition to be able to extract greater agenda access than their counterparts.

- Critical events ideology hypothesis: A cabinet party's agenda access will increase with its proximity to the opposition

This hypothesis may seem to contradict Martin's (2004) findings regarding opposition divisiveness and the government agenda. Recall, however, that Martin's model, which considered the cabinet more akin to a unified actor than the Diermeier and Stevenson model, found that bills drawn by ministers who were ideologically distant from the opposition would receive timing priority: the greater the distance between the proposing minister and the opposition, the sooner the bill would be initiated. Our concern here, however, is not with timing, but with the number of bills submitted. Thus, while the cabinet as a collective may prioritize bills that draw distinctions between their policy goals and the preferences of the opposition, individual parties should be better able to make agenda demands when they can more credibly commit to dissolving the present cabinet to pursue an alternative coalition with parties currently in opposition. This is similar to the relationship observed between individual party members and the party organization as a whole in majoritarian systems with single member districts, where members closer to the opposition suffer deprioritization of their policy goals but are also able to extract rents from their party organizations (Carroll and Kim 2010; Jenkins and Monroe 2012).

There is an important caveat to bear in mind when evaluating the critical events hypotheses, however. The model implications, particularly our popularity hypotheses, absolutely require that the cabinet be able to call new elections. If calling early elections is not a power possessed by the cabinet, as in, say, Norway, then popularity events are irrelevant as there is no device to realize the new potential seat distributions.[6] Thus, the model would only predict support for the popularity hypothesis and its CIEP and PEC modifiers in the presence of dissolution power. However, as we

[6]The French case is another where the cabinet lacks the power to call elections; that power

may observe replacements without the threat of dissolution, the ideology hypothesis should be far less sensitive to the presence of dissolution power.

Data and Measurement

To test the hypotheses discussed above we have assembled original data on over 32,000 government proposals across seven parliamentary democracies: Belgium (1988-2007), Denmark (1985-2014), France (1986-2002), Germany (1959-2009), the Netherlands (1994-2013), Norway (1996-2013), and Sweden (1987-2011).[7] Of the original 32,000 bills, many had to be purged for one or more of several reasons. First, only bills submitted by coalition governments will be considered in our empirical analysis, thus all bills submitted by single party cabinets (such as Stoltenberg I in Norway or Carlson's cabinets in Sweden) are discarded. Following Martin (2004) and Martin and Vanberg (2011, 2014), we exclude all bills subject to special initiation and procedure rules, particularly budgets and constitutional changes. We also eliminate transpositions of European Union directives, as these bills are compulsory and a product of the EU's agenda, rather than that of the domestic cabinet. Finally, several bills are unavoidably lost due to missing data on the independent variables.

This culling and attrition leaves us with over 18,000 ministerial proposals. For each proposal we note the date of submission, the proposing minister and her party and department. We then

resides with the president. In the remainder of our countries, the cabinet does possess the power to call new elections, albeit in different forms. In Denmark for instance, that power resides with the Prime Minister, in Belgium the cabinet as whole must consent. These variations, however, do not change the model implications as all parliaments allow for replacement in the incidence of lost confidence and the new cabinet may immediately call new elections. Thus, members of the current sitting cabinet can credibly commit to new elections. See Strøm and Swindle (2002) for more on these powers.

[7]We use Brüninger and Debus' (2009) data on Belgium and France and Fortunato et al.'s (2013) extension of their German data. The remainder were collected from the countries' respective legislative databases.

transform this data into ministry-month observations, where the unit of analysis is the number of bills proposed by a minister in a given month. This variable ranges from 0 to 32, though majority of the counts are zero and the modal non-zero count is one.[8]

We choose months, rather than cabinets or years, as the time interval to allow for variations in party popularity as well as variation in the CIEP. Ministries are chosen as the unit of analysis because ministries are likely to vary systematically in the number of proposals they submit. These variations could be a function of departmental prestige, giving a Minister of Justice an advantage over the Minister of Culture, for example. Alternatively, these variations could correspond to the breadth of the department's reach. They could also be a function of the times.[9] In any event, choosing the ministry-month as the unit of analysis allows us to explicitly model these dependencies hierarchically, giving us better statistical leverage on the explanatory variables we are substantively interested in. We feel this approach is superior to the natural alternative, party-month units, where we would be forced to make some qualitative assessment of departmental prestige or reach and aggregate these assessments in some way for each party in each cabinet.

Because the allocation of ministries is almost perfectly proportional in all of our cabinets, analyzing ministry-months allows to evaluate the mandate model hypothesis as the baseline. That is, because the allocation of ministries is Gamsonian, as we show in Figure 1, we would expect that each should initiate the same number of bills, ceteris paribus. Therefore, if all variables pertaining to popularity, pre-electoral coalitions, etc. lack explanatory power, the data are suggesting that all ministries receive equal share of the agenda, and, by extension, the agenda is allocated according to seat share.

Once the data is transformed, we incorporate data on party popularity as well as several other

[8]The 32 count was from Sweden's Ministry of Justice in 1993 during the Bildt cabinet. This was a particularly busy cabinet, likely because it was composed of center-right parties and followed a period of 9 uninterrupted years of Social Democratic single-party government.

[9]For instance, Germany's Minister of Finance was extraordinarily busy ironing out the ramifications of East German debt in the years following reunification.

Figure 1: Allocation of Ministries according to seat share

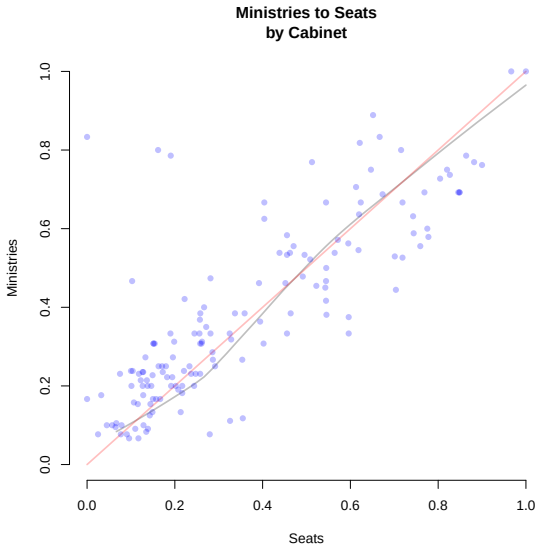

attributes of the proposing minister and her relationship to the rest of the cabinet and the opposition. The popularity data were taken from a variety of different sources that are discussed in the appendix, but all data, except our French polls, are derived from a survey asking the standard Sunday question: "If the election was this Sunday, which party would you vote for?" The French polls are in the form of approve/disapprove responses, which we transform into party popularity in a variety of ways. All transformations are discussed in the appendix and all yield nearly identical empirical results. With this data, we construct the independent variable of interest to several hypotheses, the popularity of the particular minister's party, the popularity of the other cabinet parties, and the popularity of the opposition. In some specifications, these terms will be interacted with the length of, in months, remaining in the CIEP. We have also collected data on pre-electoral coalitions, extending Golder's (2006) data throughout our sample times for our seven countries in order to differentiate between mechanisms of the responsible government model. In the models the variable is an indicator that all parties in the current coalition have consented to reform should they win enough seats and this indicator is interacted with the popularity measures.

Our final independent variables of interest require data on the ideological positions of the parties in the legislature. For this, we turn to the Comparative Manifestos Project, specifically the construction of the left-right positions and standard errors suggested by Lowe et al. (2011).[10] We use these estimates to construct our measurement of a given ministerial party's proximity to the opposition in order to gauge the attractiveness of forming an alternative coalition with opposition parties. For each minister-month, we calculate the ideological distance between the proposing minister and each opposition in the legislature using the preferences estimated from the parties' most recently published electoral manifesto. Using these distances, we calculate the average seat-weighted ideological distance between the minister and the opposition, following Martin and Vanberg (2005, 2011). That is, each individual dyadic distance is multiplied by the number of the seats controlled by the opposing party; these weighted distances are then summed and divided by the total number of seats controlled by all opposing parties combined. We call this variable opposition divisiveness. To test our coalition compromise hypothesis, we create the same measure for the minister's distance to her partners in cabinet, calling the variable cabinet divisiveness.

With all measured variables constructed, it is time to turn our attention to potential variation induced by unmeasured factors. Here, it is important to consider the structure of the data, which presents a complex hierarchy. First, there are five higher-levels in our data where we may expect error from unmeasured factors: ministries, months, parties, cabinets, and countries. These five higher levels quickly multiply when we consider that many are crossed *with* or nested *within* other

[10]These data present the problem of error in the estimates. Because we have so many observations in our data, we can model this error by simply replacing every CMP point estimate with a random draw with a normal distribution with mean equivalent to the CMP point estimate and standard deviation equivalent to the estimate's standard error. Thus, error in the CMP scaling will manifest in the parameter estimates of our ideology variables. If the size and repetitiveness our data was significantly reduced, we would likely bootstrap this process. However, that is unnecessary given that nearly every CMP derived value (e.g., the ideological divisiveness for a given cabinet) enters the data several hundred times.

levels. For example, cabinets and parties are clearly nested within countries, but they are also crossed with each other. Ministries can be considered nested within country, but crossed with cabinets, parties, and months, *or* generalized across context and therefore *crossed* with countries, cabinets, parties, and months. All told, there are over twenty groupings within the data hierarchy on which we may observe correlations across rows of the data as a function of unmeasured variables.

Clearly, estimating a model in which each of these levels is explicitly accounted for is computationally infeasible. How are we to account for these potential perturbations, then? Fortuanto and Stevenson (2013, 2014) recommend an error-clustering exercise in which the basic model is estimated in several iterations where robust standard errors are clustered at each level, thus identifying the levels which at which standard errors are most perturbed, a process similar in concept to the diagnostic recommended by King and Roberts (2014). We can simplify this process further by theoretically selecting the most efficient levels — i.e., groupings which subsume multiple alternative groupings. For instance, the cabinet-month grouping subsumes the country, cabinet, and month levels. Though estimating a hierarchical model with random intercepts at the cabinet-month level does not allow us to parse error into its constituent levels with high degrees of certainty, this concern is largely irrelevant as we are not substantively concerned with those error estimates, we are merely concerned with recovering unbiased and efficient estimates of out independent variables of interest.

The results of these diagnostic tests (detailed in the appendix) show that the most salient levels to account for that we can simultaneously estimate are the cabinet-month, the ministry (pooled across cabinets and countries), and the party-month, where effects of the ministry-month generalize to the residual. With the specification of error components complete, we move on to the specification of covariates for testing each of our hypotheses.

Results

Our first model is simple, including only the covariates necessary to coarsely evaluate the fit of the responsible party government model against the basic critical events expectations, where the mandate model serves as the baseline expectation. Recall that the responsible party government model

has two possible mechanisms: normative, where popularity increases a minister's agenda access and vote-seeking, where this is effect is observed only in the presence of pre-electoral coalition. If, on the other hand, popularity increases a minister's agenda access, but pre-electoral coalitions mitigate these effects, the data support the critical events model. Finally, if popularity has not influence, then we may conclude that the mandate model provides better fit for the data. Thus the dependent variable is the number of bills proposed in a ministry-month and the independent variables are the popularity shares interacted with a dummy indicating that the cabinet has formed a pre-electoral coalition for the coming election. As the popularity vectors sum to 1 across the proposing party, its cabinet partners, and the opposition, we must omit one share to prevent perfect multicollinearity. Thus, we include the popularity of the proposing party and the opposition, omitting the remaining cabinet parties' share of popularity. This specification allows us to model changes in the proposing parties' popularity at the expense of its cabinet partners, thus capturing its strength relative to its partners in government. The model is an error components poisson model where the random intercepts are allowed at the level of the cabinet-month, the ministry, and the party-month, as discussed above.

We include only the most simple control variable: the total number of bills proposed in a given month as an offset term, as the number of theoretically observable bills from any given ministry is bounded by the total number of bills a cabinet is able to introduce during that month.[11] Note that, because the unit of analysis is the ministry and the distribution of ministries is (nearly perfectly) Gamsonian, we need not control for the proportion of seats contributed by the proposing minister's party, this is, in effect, built in to the model with the expectation that each ministry is allotted equal agenda space, as we discussed above. The results of the model are below.

As the Table shows, the effects of popularity on agenda access are positive, statistically robust, and substantively large. Further, the interaction between the pre-electoral coalition indicator and ministerial popularity is negative. These two relationships imply that the critical events model better explains the data than the mandate model, which predicts no popularity effects, and the

[11]This means that the parameter estimate on this term is constrained to 1.

Table 1: Basic Model

Variable	Estimates	
Intercept	-3.459	(0.182)
Pre-electoral Coalition	-0.362	(0.142)
Popularity	1.068	(0.094)
Opposition Popularity	-0.251	(0.099)
Popularity × PEC	-0.772	(0.175)
Opposition Popularity × PEC	1.078	(0.204)
Random Effects		
Party-month	0.021	(0.145)
Cabinet-month	0.000	(0.000)
Ministry	0.710	(0.842)
Fit		
N	19872	
AIC	41079	
ln(Likelihood)	-20530	

responsible government model, which predicts either a strictly positive effect or a positive effect given pre-electoral coalition. In Figure 2 below, we present the effects of pre-electoral coalition over a range of party popularity, holding opposition popularity constant 0.45 and monthly output offset constant at its mean. As the figure shows, unpopular parties benefit from pre-electoral coalition by tying the hands of their partners. However, as party popularity increases the benefits of PEC turn to costs as popular parties are unable to capitalize on their potential seat shares.

As the basic specification supports the critical events model over the mandate and responsible government models, we move on to a more nuanced testing of the critical events and coalition compromise hypotheses. Our more richly specified model includes several additional covariates which we discussed in the measurement section above. First, to test the ideological hypotheses, we include our measure of seat weighted ideological divisiveness between the proposing minister and its cabinet partners and the opposition borrowed from Martin and Vanberg (2014). Recall that these variables get larger as the minister's distance from the mean, seat weighted position of its partners or the opposition increases. Recall also that the coalition compromise model predicts a negative

Figure 2: Effect of Pre-electoral Coalitions

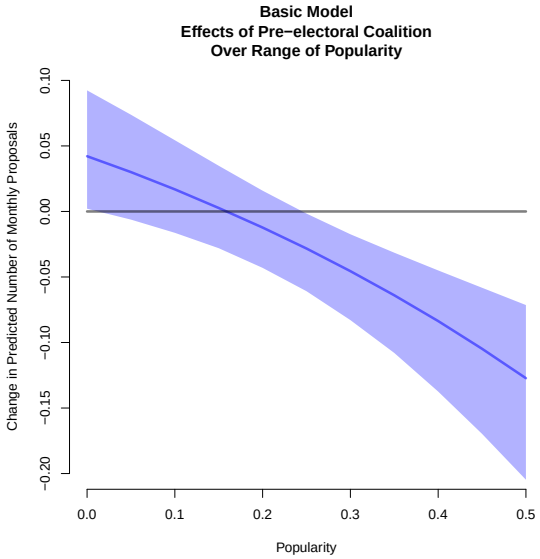

estimate on the cabinet divisiveness variable and no strong prediction on opposition divisiveness, while the critical events model predicts a negative parameter on the opposition divisiveness model and no prediction on cabinet divisiveness

We also include the (logged) number of months remaining in the CIEP interacted with the popularity terms to give us purchase on whether and how proximity to the next constitutionally mandated election effects the demands a party is able make and whether or not their partners are willing to meet them. The final change is that we estimate this model for the whole sample and then disaggregate according to whether or not the cabinet can call for new elections. Again, we the include pre-electoral coalition indicator and interactions and the same offset, as well as the same vector of error components. For the sake of comparison, however, we present the results of several models. One, the full specification estimated on the all the data. Next we estimate the full specification on two subsamples: all observations where the cabinet possesses the power the call new elections and all observations where the cabinet does not. By comparing across these two samples, we can be more certain that the relationships uncovered by the statistical models are a product

of the type of internal bargaining dynamics described by the critical events models, which require dissolution power, and not an artifact of some alternate process. Finally, we estimate an alternative model, based on the Martin 2004 analysis, for each subsample and compare it to the critical events specification. The results of these models are given in Table 2.

We begin discussion by reaffirming what we learned with the basic specification above. Again, the effects of popularity in the pooled sample and the sample constrained to cabinets that enjoy dissolution power are positive, statistically robust, and substantively large. Further, both samples demonstrate the predicted relationship of pre-electoral coalition: pre-electoral coalitions inhibit a minister from capitalizing on her popularity by sapping credibility from threats to leave the current cabinet to pursue an alternative government. These relationship between popularity and agenda access, however, *do not* hold for cabinets that do not enjoy dissolution powers. This is in keeping with the critical events model implications. In order for popularity events to affect internal cabinet bargaining, parties must have some institutional mechanism allowing them to realize the potential seat distributions public opinion implies; without dissolution power, these events can have no impact.

The effects of CIEP (the logged number of months until the next constitutionally mandated election) on popularity, however, are substantively small and statistically insignificant in each model. Indeed, the parameter on the popularity-CIEP interaction is in the wrong direction in the model of dissolution power cabinets. It is possible, however, that the effects of time we observe regarding dissolutions do not effect agenda allocation in the same way. It is possible that cabinet parties, though less invested in maintaining a government about to expire, are nonetheless more likely to acquiesce to the demands of popular partners in order to ingratiate themselves in hopes of being invited into coalition again after the next election. More research is needed to understand these relationships.

Moving on to the ideological parameters, as predicted by the critical events model, proximity to opposition is beneficial in every model. The data imply that in every case, a ministerial party is privileged by its ideological compatibility to the opposition. As replacement does not hinge on the ability of a cabinet to call new elections, this relationship holds across institutional contexts.

Table 2: Full Model for Critical Events with Alternate Model for Case Comparison

Variable	All		With Dissolution Power				Without Dissolution Power			
	Full		CE Model		Alt Model		CE Model		Alt Model	
Intercept	-3.490	(0.272)	-3.473	(0.312)	-3.302	(0.209)	-2.761	(0.600)	-2.751	(0.150)
Cabinet Divisiveness	-0.031	(0.012)	-0.031	(0.014)	-0.062	(0.012)	-0.040	(0.030)	-0.042	(0.023)
Opposition Divisiveness	-0.040	(0.009)	-0.036	(0.010)	-0.054	(0.009)	-0.041	(0.023)	-0.045	(0.021)
CIEP (in logged months)	0.074	(0.068)	0.073	(0.079)			0.017	(0.169)		
Pre-electoral Coalition	-0.381	(0.147)	-0.517	(0.183)			-0.368	(0.338)		
Popularity	0.781	(0.324)	1.220	(0.391)			-0.327	(0.644)		
Opposition Popularity	0.278	(0.324)	0.012	(0.392)			0.138	(0.775)		
Popularity × CIEP	0.027	(0.101)	-0.121	(0.122)			0.148	(0.200)		
Opposition Popularity × CIEP	-0.183	(0.104)	-0.136	(0.124)			-0.076	(0.242)		
Popularity × PEC	-0.651	(0.185)	-0.631	(0.221)			0.139	(0.449)		
Opposition Popularity × PEC	1.051	(0.213)	1.346	(0.268)			0.512	(0.473)		
			Random Effects							
Cabinet-month	0.000	(0.000)	0.000	(0.000)	0.000	(0.000)	0.000	(0.000)	0.000	(0.000)
Party-month	0.018	(0.134)	0.022	(0.149)	0.035	(0.187)	0.000	(0.000)	0.000	(0.000)
Ministry	0.713	(0.844)	0.817	(0.904)	0.998	(0.999)	0.337	(0.580)	0.338	(0.581)
			Fit							
N	18628		14387		15743		4241		4241	
AIC	38753		31029		33916		7344		7332	
ln(Likelihood)	-19363		-15500		-16952		-3658		-3660	
			CE better $p < 0.000$				CE better $p < 0.917$			
			CE better $\chi^2 = 2903$				CE better $\chi^2 = 3$			

Further, as the coalition compromise model would predict, parties that ideologically extreme relative to their partners in cabinet are less able to get their proposals heard by the plenary. Thus, the same monitoring effects that we observe in reference to a cabinet's prioritization of legislation, e.g., Martin (2004), holds in a more dynamic setting as well.

Finally, we compare fits across models. For both subsamples we estimated the critical events model as well as an alternate model building on Martin (2004), capturing the internal division between the proposing minister and her cabinet partners as well as the level of ideological division between the minister and the opposition. Looking at the cabinets that have dissolution power, the explanatory variables included in the critical events model specification provide substantially better for the data over the alternative model. For the cabinets without dissolution power, however, the alternate specification provides better fit. In keeping with the critical events model, without dissolution powers, there is no mechanism for parties to realize the potential seat shares implied by public opinion.