

Modeling the Institutional Foundation of Parliamentary Government Formation

Matt Golder Pennsylvania State University
Sona N. Golder Pennsylvania State University
David A. Siegel Florida State University

That neither the assumptions nor the predictions of standard government formation models entirely correspond to empirical findings has led some to conclude that theoretical accounts of government formation should be reconsidered from the bottom up. We take up this challenge by presenting a zero-intelligence model of government formation. In our model, three or more parties that care about office and policy make random government proposals. The only constraints that we impose on government formation correspond to the two binding constitutional constraints that exist in all parliamentary systems: an incumbent government always exists and all governments must enjoy majority legislative support. Despite its deliberately limited structure, our model predicts distributions over portfolio allocation, government types, and bargaining delays that approach those observed in the real world. Our analysis suggests that many formation outcomes may result from the institutional foundation of parliamentary democracies, independent of the strategic behavior of party leaders.

The legislative elections in Israel in 2009 failed to produce an obvious government coalition.¹ Instead, the leaders of the two largest parties, Kadima and Likud, both claimed victory and began negotiating with potential coalition partners over cabinet positions and the policies that any new government would pursue. This initial period of negotiations provided some information about the viability of various potential government coalitions. The head of state then appointed the leader of the conservative Likud Party as *formateur* because he seemed more likely than the centrist Kadima leader to be able to successfully put together a coalition that would enjoy the support of a legislative majority. The negotiations over the new Israeli government lasted seven weeks, during which time the incumbent Kadima prime minister remained in power.

This Israeli example highlights some important features of how governments form in parliamentary democracies. First, it is often the case that there is no obvious government alternative and politicians will explore a great many options before arriving at a final government. Given that it is not unusual for a legislature to have five or more parties, many potential coalitions will typically be considered. Although some will be instantly

rejected, others will be investigated, and this process takes time. Delays often occur because party elites are considering different options, trying to obtain the most advantageous coalition. During this time, the incumbent government stays in office in a caretaker capacity. Even though caretaker governments are typically not expected to implement new policies, it is almost certainly the case that having the incumbent Kadima party remain in office was likely to please some legislative parties more than others. For example, we suspect that Kadima members, who were expecting to soon move into the opposition, had a higher tolerance for a prolonged government formation period in this instance than Likud members did.

Contrast this Israeli example (or many other government formation stories we could have told) with standard theoretical accounts of the government formation process. The typical approach is based on the canonical (closed-rule) Baron-Ferejohn model in which three legislative parties bargain over forming a new government by making alternating offers (Baron and Diermeier 2001; Baron and Ferejohn 1989). A proposer (*formateur*) is first selected using an exogenous selection mechanism. This proposer then has the sole right to make a proposal, typically involving

¹The data and all computer code necessary to replicate our results are publicly available on the authors' homepages: <https://files.nyu.edu/mrg217/public/>, <https://files.nyu.edu/sln202/public/>, and <http://myweb.fsu.edu/dsiegel/>.

some distribution of office benefits and a government policy position, to the other parties. The leaders of these other parties cannot hold their own negotiations or make alternative proposals. All parties then vote on the proposal made by the formateur. If the proposal receives majority support, then a new government forms and implements its policy immediately. If the proposal does not receive majority support, then the exogenous selection mechanism is called upon again to designate a new proposer and the process begins once more. Prior to the proposal of a new government being accepted, all parties are assumed to receive identical payoffs, often normalized to zero.

Although this approach to modeling government formation has contributed greatly to our knowledge and does particularly well in highlighting the strategic interactions that are possible between parties, it has several limitations. One is that the structure of the bargaining game, in which a single actor has the monopoly power to make proposals at any given point in time, does not accurately reflect how party leaders actually bargain over who gets into government. Though all models are necessarily simplifications of reality, the fact that the removal of the alternating bargaining assumption from these models results in their being unable to produce equilibria is obviously problematic given the failure of this assumption to hold empirically. More importantly, standard models lead to equilibrium predictions that are often at odds with the empirical evidence from parliamentary democracies that has accumulated over the last six-and-a-half decades. The standard predictions are that governments will be minimal winning, that ministerial portfolios are allocated in a non-proportional manner, and that governments form swiftly. In reality, minority and surplus majority governments form quite regularly (Laver and Schofield 1998; Strøm 1990); cabinet portfolios are nearly always allocated proportionally (Warwick and Druckman 2001, 2006); and there is significant variation in how long it takes governments to form (Diermeier and van Roozendaal 1998; Golder 2010; Martin and Vanberg 2003).²

²We recognize that some models in the Baron-Ferejohn tradition produce something other than the standard predictions that empirical scholars typically associate with this area of research. For example, Carroll and Cox (2007) and Morelli (1999) both present models in which portfolios are allocated proportionally in some equilibria. Similarly, there are models that produce governments other than minimal winning coalitions (Bandyopadhyay and Oak 2008; Baron and Diermeier 2001; Diermeier and Merlo 2000), one of which also predicts delays in government formation (Diermeier, Eraslan, and Merlo 2003). That said, we know of no existing model that simultaneously produces predictions that approximate empirical distributions of government delay, government type, and portfolio allocation.

That neither the assumptions nor the predictions of the current theoretical approach correspond closely to the empirical findings has led some to suggest that it might be fruitful to consider a different approach to modeling government formation. For example, Laver, de Marchi, and Mutlu emphasize the gap between standard models of government formation and the empirical evidence when they remark that “the profession’s canonical theory of bargaining in legislatures is contradicted by one of the profession’s strongest and most robust empirical laws [Gamson’s Law]” (2011, 288). They conclude that the topic of government formation “must be reconsidered from the bottom up.” We take up this challenge by presenting a new type of government formation model.

A key feature of any government formation model has to do with how parties negotiate. Existing models typically assume that party leaders are utility-maximizing agents who make optimal choices as they negotiate strategically using very specialized bargaining protocols. The problem is that almost nothing substantive is written on how parties actually bargain over who gets into government and what policy should be implemented. One reason for this is that analysts are seldom party to the secretive backroom dealings that underpin the typical government formation process. In our model, we circumvent this problem by placing no constraints on the types of government proposals that can be made. In this way, our model bears some similarity to those presented in Schofield and Sened (2006) and Penn (2009). Specifically, our model, like theirs, diverges from the standard approach and views government formation as a lottery over possible coalitions, policies, and distributions of perquisites (in our case, cabinet ministries).

The only constraints we impose on government formation correspond to the two constitutional constraints that exist in *all* parliamentary systems: (1) an incumbent government always exists, and (2) all governments, in addition to having the support of their member parties, must enjoy majority legislative support. Our model represents a “null” model in that any model of parliamentary government formation must impose these constraints and any model different from ours must add at least some additional structure. Our goal in this article is to abstract away from the exact details of the bargaining process and instead examine how the two constitutional constraints that are common to all parliamentary democracies shape government formation.

To do this, we draw on the economics literature dealing with “zero-intelligence agents,” which indicates that the institutional structure of markets can

itself determine outcomes (Bosch and Sunder 2000; Farmer, Patelli, and Zovko 2005; Gode and Sunder 1993, 1997). A zero-intelligence agent is one who acts randomly subject to minimal constraints.³ In their seminal article, Gode and Sunder (1993) report results from an experiment in which human traders are replaced by zero-intelligence agents that submit random offers. Gode and Sunder demonstrate that the efficiency of auctions can be raised close to 100% simply by imposing a budget constraint—not permitting traders to sell below their costs or buy above their values. They conclude that “the allocative efficiency of a double auction derives largely from its structure, independent of traders’ motivation, intelligence, or learning” (1993, 119).

We apply the same basic approach to investigate how the institutional structure of parliamentary democracies influences government formation. Specifically, we present a model in which parties that care about office and policy make random government proposals, but only those proposals that are both preferable to the status quo government for the proposer and receive legislative majority support are enacted.⁴ By allowing all parties to make proposals simultaneously and averaging across a great many formation opportunities, any patterns that emerge should be the consequence of the common institutional foundation shared by all parliamentary democracies.

Our choice of how to model “bargaining” may seem unusual for a model aiming at empirical verity. Given that our approach diverges quite markedly from that observed in much of the formal literature, we should reiterate the purpose of our model. Our model makes no pretense to accurately describe or explain the bargaining behavior of actors involved in the government formation process (Gode and Sunder 1997, 604). We hold no illusions, for example, that party leaders really gather together and simultaneously announce randomly chosen coalitions and government policy. And surely more goes into the choice of government than a comparison to the status quo. Moreover, we do not see our model as a direct challenge to existing models that rely on utility maximization and strategic behavior. Instead, the goal of our model is simply to discover whether the

two constitutional constraints that exist in all parliamentary systems are sufficient to generate empirically realistic predictions about the types of governments that form, portfolio allocation, and delays in government formation, independent of context-specific factors, norms of behavior, or bargaining protocols that limit the variance in the types of proposals that parties can make.⁵

In the next section, we outline what we know empirically about parliamentary government formation. We present our model in section three. The fourth section is devoted to discussing the model’s predictions and investigating how well they compare to real-world data. Given the bare bones nature of our setup, it is quite remarkable how well the predictions of our model match what is observed in the real world. Our model suggests that the two institutional constraints that characterize parliamentary government formation—the existence of an incumbent government and the requirement that governments enjoy majority legislative support—can generate many of our most robust empirical regularities, independent of the strategic behavior of party leaders and their bargaining protocols. In addition to showing the similarity between our results and real-world data, we also illustrate how each of our major results varies with factors such as party system diversity, party system size, and the weight that parties assign to policy relative to office. We do this not only to help explain what is necessary for our model to produce the results that it does, but also to display the wide range of theoretical and empirical questions that are opened up by our model.

What We Know Empirically About Government Formation

In what follows, we draw on the vast empirical literature addressing parliamentary government

³Gode and Sunder (1993) were the first to use the term “zero intelligence.” Their highly cited article has led to a large and growing literature on zero-intelligence agents in economics, computer science, and other fields. To our knowledge, we provide the first zero-intelligence model in political science.

⁴Although we situate our model in the zero-intelligence literature, we should note that our random proposal mechanism can be justified on other grounds (Compte and Jehiel 2004).

⁵Note that we are not claiming that context-specific factors, norms of behavior, and alternative bargaining protocols are unimportant. Even if our model were to produce results that closely matched empirical regularities, we recognize that there would likely remain a significant role for these other factors in predicting government outcomes. Indeed, Gode and Sunder’s (1993) market-constrained zero-intelligence agents produce variance in outcomes far greater than human market agents, thereby implying that additional, unmodeled factors must account for the reduction in variance. Along similar lines, it is entirely plausible that a norm toward fairness could constrain the variance around, say, the proportionality of cabinet portfolio allocation produced by our model.

formation to outline what we know about both the *process* by which governments form and the associated *outcomes* of this process.

The Government Formation Process

In parliamentary democracies, the government comprises the prime minister and her cabinet. New governments emerge either because an election occurs or because an incumbent government resigns. But what is the structure of the government formation process? Any answer to this question must start with the constitutional constraints that bind in any formation process.

Binding Constitutional Constraints. There are only two constitutional constraints that bind in any parliamentary government formation process. These constraints bind for *all* governments irrespective of whether the governments form in postelection or interelection periods. The first constraint is that all governments, in addition to having the support of their member parties, must enjoy the support of a legislative majority. In some parliamentary democracies, the government must explicitly demonstrate that it has majority legislative support in an investiture vote before it can officially take office. But even if there is no formal investiture vote, all parliamentary governments must enjoy at least the implicit support of a legislative majority due to the ability of the legislature to call votes of no confidence. If a government ever loses such a vote, because it cannot garner the support of a legislative majority, then it must resign. Ultimately, a parliamentary government may be removed from office any time a majority of legislators decides that this is what should happen. As a result, any incoming government must be able to survive a vote of no confidence and, hence, enjoy the support of a legislative majority even if it never has to explicitly demonstrate this through an actual vote.

The second and final constitutional constraint that binds any parliamentary government formation process is that there is *always* a government. Constitutionally, the incumbent government remains in place until it is formally replaced by an alternative. Whether a government resigns or is defeated in a vote of no confidence, it remains in power as a caretaker government until a new duly mandated government replaces it. Although there is variation in how much new policy a caretaker government can implement, it remains the case that the policies implemented by the incumbent government stay in place until they are replaced by the policies of a new government. As a result, the incumbent caretaker government always

represents the status quo or reversion point in any government formation process.

“Nonbinding” Modeling Constraints. Standard government formation models impose additional “nonbinding” constraints on how governments form. In general, these constraints come in the form of technical assumptions about the nature and sequence of play in whatever bargaining model is being employed. One of the most distinctive constraints imposed by many models is that there is a common knowledge formateur who is exogenously determined and who has the sole right to propose governments. In effect, these models assume that formateurs are chosen either sequentially in order of party size starting with the largest party (Austen-Smith and Banks 1988) or probabilistically where the likelihood of being selected is proportional to the formateur’s share of legislative seats (Baron and Ferejohn 1989).⁶

Empirically, though, few constitutions state how formateurs are to be selected. Although it might seem reasonable that formateurs will be selected in order of legislative size as some models do, there are only two European countries—Greece and Bulgaria—that constitutionally require this. In some countries, such as Italy, Portugal, Austria, Finland, and France, the head of state has enjoyed considerable discretion over formateur selection. In other countries, such as the Netherlands and Belgium, the head of state appoints an informateur who talks to party leaders and gathers information about viable governments. The head of state then chooses a formateur based on the informateur’s advice. In many cases, negotiations over government formation are simply described as “free-style bargaining” among the various parties. In general, the vast majority of empirical evidence indicates that formateurs emerge endogenously as part of the underlying bargaining process.⁷ As a result, we do not

⁶Rarely is evidence provided to support either of these selection mechanisms. Although Diermeier and Merlo (2004) claim that the probabilistic selection rule receives empirical support, their analysis is open to question because they use *Keesings* to identify formateurs. Despite “sustained and determined efforts” to reconstruct data on formateurs using *Keesings*, Laver, de Marchi, and Mutlu “concluded unambiguously that ... it is not ... systematically possible to observe *ex ante* exogenous *formateur* status in primary data sources” (2011, 292). They go on to conclude that formateur selection rules “are not testable using a variable for exogenous formateur status, coded from historical sources.”

⁷With the exception of Morelli (1999), Diermeier, Eraslan, and Merlo (2003), Yildirim (2007), and Bassi (2008), existing models of government formation adopt an *exogenous* selection mechanism for determining the formateur. For a far more complete discussion of the modeling literature in regard to the choice of formateur, see Penn (2009).

include an exogenous formateur selection rule in our model.

Irrespective of how the formateur is chosen, it is almost certainly unrealistic to expect that the formateur has a monopoly right on making government proposals and that other political parties do not hold their own private negotiations. Even if the head of state has designated formateur status to the leader of one party, the leaders of other parties are nearly always bargaining with each other in private over what a viable coalition would be. Indeed, our Israeli example from earlier indicates that the leaders of Kadima and Likud both claimed victory and began *simultaneously* negotiating with potential coalition partners.

It is important to recognize in all of this that government formation is the principal and most important game to be played in any parliamentary democracy. As a result, it is hard to imagine why any real-world party leader would sit by if she thought that she could put together a viable government. The recent U.K. government formation process is of particular interest here. Historically, it has been a custom that the incumbent British PM gets first shot at forming a government when there is a hung parliament. In February 2010 prior to the general election, this custom was formalized for the first time in revisions to the Cabinet Manual (Chapter 6) by Gus O'Donnell, the head of the civil service. The fact that the political parties ultimately ignored this *formal* rule suggests that even when they do exist, rules about formateur selection and proposal rights are de facto nonbinding when it comes to something as important as government formation.

There are several insights to take away from this discussion. The first is the distinct lack of structure that characterizes the typical government-formation process. The second is that negotiations over potential governments occur simultaneously rather than sequentially. And finally, much of the bargaining among party leaders occurs in private with the result that scholars have very little information about how negotiations actually proceed in practice.

The Outcome of the Government-Formation Process

What do we know about the outcomes associated with the typical parliamentary government formation process? In what follows, we focus on the various types of government that emerge, the degree of proportionality in portfolio allocation, and delays in the government-formation process.

Government Types. Historically, most government-formation models predict that minimal winning coalitions will form. This prediction, though, is only partially borne out in the real world. In fact, fewer than half the governments that formed in West European parliamentary democracies between 1945 and 1998 conform to this prediction: about 16% of the governments were single party majority governments and about 30% were minimal winning coalitions. It turns out that about 37% of West European governments have been minority governments and a further 17% have been surplus majority governments (Clark, Golder, and Golder 2009, 419). Surplus majority governments even occur 15% of the time when a single party controls a legislative majority.

Recently, formal scholars have made great strides in constructing models that predict governments other than minimal winning coalitions (Bandyopadhyay and Oak 2008; Baron and Diermeier 2001; Diermeier and Merlo 2000). However, most of these models are unable to produce the *full* range of government types that we see in the real world—single party majority, single party minority, minority coalition, minimal winning, and surplus majority. To our knowledge, Diermeier, Eraslan, and Merlo (2003) present the only model capable of producing the full range of observed government types. No existing model predicts the *distribution* of government types.⁸

Allocation of Portfolios. One of the strongest regularities in all of political science—Gamson's Law—states that cabinet portfolios are distributed among government parties in proportion to the number of seats that each party contributes to the government's legislative majority. Gamson's Law implies that if we were to regress a party's share of portfolios against its share of the government's legislative seats, then we would obtain a slope close to one and an intercept close to zero. Empirical studies have consistently found support for this (Warwick and Druckman 2001, 2006).

⁸We should note that many of these models employ a restricted number of parties, policy space, or timing of offers for reasons of tractability—it is not trivial to derive comparative statics in a model containing both policy preferences and distributive goods. For example, although Jackson and Moselle (2002) are able to prove the existence of certain classes of equilibria in an unrestricted legislative-bargaining model with distributive goods and a one-dimensional policy space, they are forced to turn to examples with three parties and specific parameter values to detail the comparative statics of these equilibria. A model like ours that simultaneously produces comparative statics on portfolio allocation, government type, and bargaining delay on an unrestricted domain is an advance in this respect.

Despite being one of the strongest and most robust empirical laws in all of political science, Gamson's Law is not an implication of virtually any bargaining model of coalition formation. In fact, standard models have the potential to produce quite a high degree of disproportionality, depending on things like party discount factors (Austen-Smith and Banks 1988; Baron and Diermeier 2001; Baron and Ferejohn 1989; Snyder, Ting, and Ansolabehere 2005). We recognize that a few models do produce Gamson's Law (Bassi 2008; Carroll and Cox 2007; Morelli 1999), but even these only predict it for some equilibria, not all equilibria.

Bargaining Delays. A characteristic of most parliamentary systems is that elections do not typically determine the identity of the government but instead set the stage for a prolonged period of bargaining over the composition and policy goals of the cabinet (Golder 2010). Consider the government-formation process following the June 2010 elections in Belgium. As is usual, the king appointed an informateur to gather information and report back on a suitable formateur to manage the actual government-formation process. The nonpartisan king has been kept unusually active appointing numerous informateurs and formateurs as negotiations over possible coalitions repeatedly failed. As of September 2011—over a year later—Belgium still had no new duly mandated government.

Although the long duration of this particular Belgian government formation process is somewhat unusual, it is not uncommon for there to be a delay of several weeks or months while party leaders bargain over the choice of government. While it takes a few days on average for a cabinet to form in France, Norway, and Sweden, it takes close to three months for them to form in the Netherlands. On average, it has taken close to a month for governments to form in Western Europe (Clark, Golder, and Golder 2009). Significantly, these delays can be quite consequential (Golder 2010). For example, the political and economic uncertainty generated by bargaining delays has been shown to affect things like exchange rate markets, stock market volatility, the types of assets in which actors invest, and government stability.

Although the evidence indicates that there can be a considerable delay in the government-formation process and that these delays can have serious consequences, many models predict that governments form immediately. Furthermore, those models that allow for multiple formateurs and hence delays (Austen-Smith and Banks 1988; Baron and Diermeier 2001; Diermeier, Eraslan, and Merlo 2003) do not

predict the actual length of delay and, in fact, typically make use of the rationality of parties to limit the length of time over which negotiations take place.

A Zero-Intelligence Model of Government Formation

In this section, we present a zero-intelligence model of government formation. Our single-district model consists of a series of elections, each containing two stages: an election stage and a government-formation stage. In the election stage, m parties, where m is a parameter of the model, are assigned platforms in a two-dimensional Euclidean policy space $X \subseteq \mathbb{R}^2$.⁹ The N voters, who are evenly spaced in the policy space, are all proximity voters; that is, they vote for the closest party in the policy space, deciding ties via a fair die. The election stage thus determines the vote share that each party receives, which is translated directly into a legislative seat share via a perfectly proportional rule. Note that the election stage involves no choices by actors; it simply takes policy locations for m parties as parameter inputs and produces seat shares as outputs. The election stage is deliberately simple and nonstrategic since our primary goal is to examine the government-formation process from the point at which legislative seats have been allocated.

The election stage in our model defines a “party system” consisting of a set of parties, policy positions, and seat shares. Once the election stage is complete, the government-formation process begins. The government-formation stage takes as an input the party system as well as an endogenous status quo government and its policy position. The utility functions of the parties are linear combinations of their cabinet shares in the eventual government and a quadratic loss term that results from policy divergences between their ideal points and the eventual government policy.¹⁰ Let c_i be party i 's proportion of cabinet seats; it takes the value 0 when a party is not in government. Similarly, we assume that \mathbf{x} is the government policy in two dimensions, and \mathbf{z}_i is party i 's ideal point, which is also its platform. We assume that cabinet portfolios are infinitely

⁹Analysis using a one-dimensional policy space yields qualitatively similar results to those that we present.

¹⁰As we have no priors on the true functional form of this policy loss, we also analyzed the model using a piecewise-linear loss term. Results were qualitatively similar.

divisible. With these assumptions, party i 's utility is $u_i = cs_i - \beta(\mathbf{x} - \mathbf{z}_i)^2$. β is the utility weight, relative to a normalized value of 1 for the cabinet share, that parties place on policy concerns.

The government-formation stage consists of the following five steps, which define the transition probabilities of a Markov process, with the status quo governments as states:

1. Each party *simultaneously* chooses a potential coalition, always including itself, at random. All possible coalitions, including the one with no other members but itself, are equally likely.
2. Each party chooses a random distribution of cabinet portfolios and a policy position for this potential coalition. All distributions and policies within X are equally likely.¹¹
3. Each party then compares the utilities associated with each of the proposed coalitions to the utility associated with the status quo. With the exception of the first election, the status quo is always the previous election's government, which serves as the caretaker government just after the election. The initial status quo before any election has occurred consists of the median voter's policy with cabinet shares equal to zero for all parties. Any proposed government that strictly beats the status quo for all potential members of the governing coalition and additionally for a number of parties controlling a simple majority of legislative seats draws the attention of the head of state.
4. If only one proposed government has drawn the head of state's attention, then this is installed as the duly mandated government and becomes the new status quo for the next election. If more than one proposed government fits the bill, then the head of state chooses one randomly and installs it

¹¹One might argue that assuming a uniform distribution over proposals is a rather strong assumption. However, we believe that assuming a uniform distribution over proposals is, in a very specific sense, the weakest assumption one can make in a model where proposals must be made. Note that *any* model that requires the production of proposals is necessarily equivalent to a model that delineates some distribution, possibly degenerate, over possible proposals. Given this, the uniform distribution is the least informative distribution available to us. Indeed this is precisely why it is employed in zero-intelligence models and why a "flat" prior is so often used in Bayesian statistics. We should note, though, that we replicated our analysis with a variant model in which parties cannot make proposals that include a policy in either dimension that is more extreme than the proposed coalition's most extreme party's ideal point in that dimension. Other than a decrease in the mean duration of bargaining delay, results were qualitatively similar.

- as the new government and status quo.¹² The endogenously determined successful formateur is the proposer of the government alternative that takes office. If no proposed government has the support of its coalition and a majority of the legislature, then all proposed governments are discarded and all parties start from scratch in step 1.
5. This process continues until either a government is installed or 100 periods pass. In the latter case, the caretaker government is reinstalled. This may be interpreted as the need for a functioning government taking precedence over continued bickering. This cutoff is typically invoked in the model when a particular arrangement of parties and policies favors a narrow range of proposals, leading to the failure of parties to find anything better in subsequent elections.¹³

In our analysis, we model a repeated government-formation process for a given party system, i.e., for a fixed set of parties, policy positions, and seat shares. With the exception of the initial status quo, governments emerge endogenously and the choice of government responds dynamically to the previous status quo government. We then repeat this process for different party systems.

Results

With more than three parties, random coalition choice, and majority rule voting, no tractable analytic solution exists from which to draw insight. We therefore turn to simulation methods (Bendor et al. 2011; Laver and Sergenti 2011; Miller and Page 2007).

¹²We explored two variations to this rule. In one, we had the head of state choose the proposal with the policy closest to the median voter. In the other, we had the head of state take a straw poll of all parties and choose the plurality winning proposal. These different rules did not yield appreciably different results. We also analyzed a variant of our model in which all parties had the chance to make some number of additional proposals in an attempt to beat each new status quo before it was installed as the next government. Although we obtained qualitatively similar results for the distribution of government types and the allocation of portfolios, we obtained somewhat different results for bargaining delay. The reason for this is that in this variant of the model easily beaten status quos lead to *longer* delays, not shorter ones, as new winning proposals are introduced again and again. Given a lack of empirical knowledge as to exactly how the bargaining process unfolds, we chose to stick with the simpler modeling assumption.

¹³We also examined the situation in which the government-formation process could continue for 1,000 periods before the caretaker government was reinstalled; our results were qualitatively similar.

The Markov process defined in the previous section is ergodic when discretized in the simulation, implying that for any set of parameters, all simulated outputs (government types, Gamson's law regression parameters, and delay lengths) will be draws from unique limiting distributions.¹⁴ Our process is thus well defined. In this case, the Markov process will converge to its unique limiting distribution from all initial state vectors, so we do not have to worry that our results are dependent on the initial conditions, specifically the arbitrary initial status quo. The benefit of this is that the model can be thought to have substantial predictive power (Bendor, Diermeier, and Ting 2003, 271).¹⁵

Our simulation procedure is as follows. First, we choose a set of $2 \times m + 2$ parameters corresponding to the number of parties (m), the value parties place on policy relative to holding cabinet ministries (β), and the two-dimensional ideal points for each party (\mathbf{z}_i). For simplicity and to make interpretation easier, we assume the same value of β for all parties. Each parameter is drawn from a corresponding uniform distribution. The number of parties ranges from three to ten, covering a wide range of real-world party systems. Ideal points are drawn from $X = [-3, 3] \times$

¹⁴Though the theoretical model may have an infinite state space, instantiation on a finite-state machine (i.e., a computer) implies that the simulation is a finite-state Markov process. Define a proposal formally as $(\mathbf{cs}, \mathbf{x}) \in A$, and $A' \subseteq A$ as the top cycle of A . There is then a nonzero probability that a status quo in A' may carry over to the next government formation. Further, by definition, it is possible to construct a series of status quos that reach any other status quo in finite time. Accordingly, the process is aperiodic and all states in A' communicate. Again by definition, all proposals not in A' must be beatable by those in A' ; thus all such states are transient, and the set A' absorbing. Accordingly, the process is ergodic and produces a unique limiting distribution on A' to which it converges from all initial conditions. See Bendor et al. (2011) for a discussion of the importance of ergodicity for simulated results.

¹⁵We can now see why there is no tractable analytic solution to this problem. An analytic solution would entail, at a minimum, determining the limiting distribution of the Markov chain, which would require specifying the transition matrix between (very high-dimensional) state vectors. While we can specify the probability that any particular proposal is made during any government formation attempt, the use of majority rule to determine which proposals might be enacted introduces complications. For nearly all initial conditions, for any proposal there exist other proposals that a majority would prefer, given a two-dimensional policy space and an $m - 1$ simplex of possible cabinet share distributions. To specify the transition matrix, one must determine the set of proposals preferred by a majority to every status quo proposal; this set need not be convex. While this is theoretically possible for a finite proposal space, it is not tractable. Thus, much as we turn to (Monte Carlo) simulations to understand the properties of estimators in quantitative analyses, we do so here to understand the dynamics of our model.

$[-3, 3]$, and β is drawn from the set $[0, 1]$ to represent parties' policy interests. As described in the previous section, the initial status quo sets all cabinet shares to zero, and the government policy to that of the median voter, which here is located at $(0, 0)$.

Second, we execute an election followed by a series of 2,500 government formations; we call this a "run." Each run represents a simulated data set from a single "party system" about which we collect various quantities of interest relating to formation outcomes. Although the initial parameters describing the party system are fixed during a given run, the proposals made by the parties vary from one formation process to the next. Governments emerge endogenously in accordance with the model outlined previously. As a result, government choice in each formation process responds dynamically to the previous status quo government.

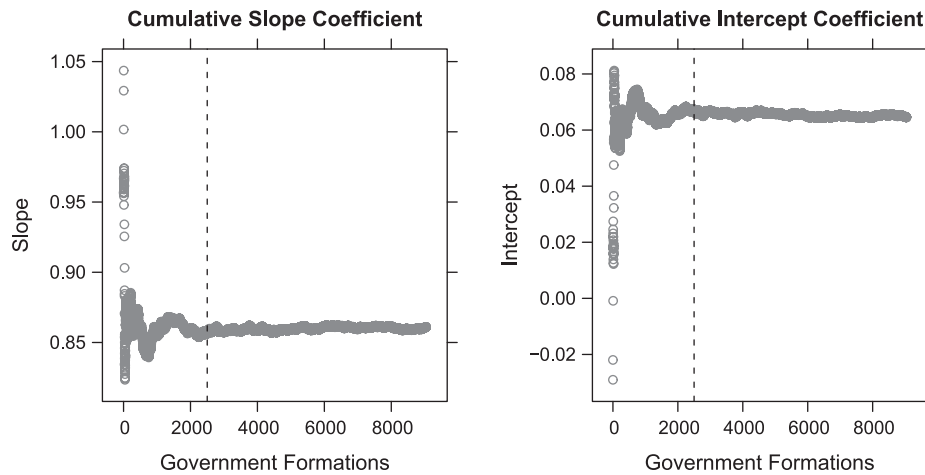
As an illustration, Figure 1 plots the cumulative estimated slope and intercept coefficients from models in which we regress the share of cabinet portfolios held by governmental parties against the share of the government's legislative seats that these parties control for one set of parameter values (party system) over the course of many government formations.¹⁶ The vertical dashed line in each panel indicates the end of our normal run of 2,500 government formations and the estimated coefficients that we obtain at this point. As one can see, both the slope and the intercept coefficients have settled down to their steady-state values of approximately 0.86 and 0.06, respectively, by this point, and do not subsequently vary much over time. The settling down process illustrated in Figure 1 obtains for all sets of parameter values. Further, for a particular set of parameter values, these results reproduce across runs.¹⁷

During the government formations, and at their end, we record a variety of outcome measures. The value of each of these outcome measures taken at the 2,500th government formation produces an estimate of the mean of the limiting distribution for

¹⁶The run used five randomly chosen parties located at $(-0.0058, 2.3)$, $(2.5, -2.4)$, $(-1.1, 1.5)$, $(0.25, 0.66)$, and $(-0.47, -1.0)$, with $\beta = 0.25$. By cumulative, we mean that the slope and intercept coefficients are reestimated after new data points are added following each new government formation.

¹⁷While the steady-state value for each run is only a draw from the well-defined limiting distribution corresponding to that set of parameter values, we have found that these steady-state values vary comparatively little across runs. For example, the standard deviations of the means of the output measures across only 10 runs was between just 5% and 10% of their mean values. Accordingly, to conserve computational time we use only one run per set of parameter values in the results below; this allows us to explore more of the parameter space.

FIGURE 1 Cumulative Estimated Slope and Intercept Coefficients Across Government Formations for One Set of Parameter Values



Note: Figure 1 illustrates the cumulative estimated slope and intercept coefficients from models in which we regress the share of cabinet portfolios held by governmental parties against the share of the government’s legislative seats that these parties control for one set of parameter values over the course of many government formations. The vertical dashed line in each panel indicates the end of our normal run of 2,500 government formations and the estimated coefficients that we obtain at this point.

each measure, and thus a meaningful result on the nature of government formation for a particular set of parameter values (party system). Focusing on the outcome measures from the final government formation minimizes any effects in early government formations that might be driven by the exogenously determined initial status quo (Laver and Sergenti 2011). In other words, we utilize a “burn-in period” sufficient to carry us into the limiting distribution (Laver 2005, 269).¹⁸

Because empirical regularities occur *across* party systems, and so across parameter values, we simulate 50,000 party systems, each with *randomly* chosen parameters. We use random parameters rather than parameterizing our model in such a way as to best fit real-world data in order to avoid biasing our results. While we could, of course, guarantee empirical congruence of our model with the real world by carefully choosing parameters, any subsequent comparisons of our model’s outcomes to real-world data would not be a fair test of our model.

As one can imagine, though, our simulation procedure yields a large space of possible party systems, some of which may not occur with great

frequency in the real world. Because of this, we report two different sets of simulation outcomes in the results below. The first set of outcomes is obtained by averaging across all 50,000 sets of runs, treating each as having an *equal probability of occurring*. The second set is designed to better reflect the types of party systems we actually observe in the real world. In this second set of outcomes, we *weight* each run by the relative frequency with which the effective number of legislative parties in that run appears in real-world data (1945–98) from 17 West European countries (Strøm, Müller, and Bergman 2003). Obviously, we would like to do more to better approximate the frequency with which different party systems occur in the real world. However, this would require knowing things like the weight that parties place on policy relative to office and how this is distributed across different party systems. Unfortunately, this type of information does not exist.

We now turn to our results regarding (1) the proportions of each type of government observed; (2) the proportionality of portfolio allocation; and (3) bargaining delays in the formation process.

Government Types

In addition to producing the full variety of observed government types—single party majority, single party

¹⁸We also performed analyses in which we included *all* simulated data and not just those from the final period. Results were qualitatively similar.

TABLE 1 Types of Parliamentary Governments

Government Type	Proportion Occurring:		
	Real-World Data Western Europe	Simulated Data	
		Unweighted	Weighted
Single Party Majority	.130	.086	.157
Minimal Winning Coalition	.307	.196	.257
Minority Coalition	.118	.461	.254
Single Party Minority	.229	.150	.230
Surplus Majority	.216	.107	.102

Note: The real-world data in the first column are for 17 West European countries—Austria, Belgium, Denmark, Finland, France (5th Republic), Germany, Greece, Iceland, Ireland, Italy, Luxembourg, Netherlands, Norway, Portugal, Spain, Sweden, and the United Kingdom—and cover the period from 1945 to 1998 (Ström, Müller, and Bergman 2003). The unweighted simulated data in the second column represent the proportion of each government type in the last election of 50,000 simulated party systems. The weighted simulated data in the third column come from the same simulated party systems, but now each system is weighted by the frequency with which its effective number of parties occurs in the observational data from Western Europe.

minority, minimal winning, minority coalitions, and surplus majority—our model also produces an empirically realistic distribution of these government types. In Table 1, we provide information on the proportion of governments that fall into each government type category. In the first column, we show the proportion of governments comprised by each government type using real-world data from 17 West European countries (1945–98). In the next two columns, we report the average proportion of governments comprised by these different types from our simulated elections, showing both unweighted and weighted results.

On the whole, our model does well at approaching the real-world distribution of government types. Of course, this, in and of itself, is not a sufficient test of our model's predictive power. Even though the distributions we obtain are similar to those in the real world, it is possible that substantial differences may exist between the distribution of governments we obtain for a particular configuration of party locations and policy weights and that which would obtain given the identical scenario in an actual government formation opportunity. So to better understand our model's behavior and to make a start at discounting such possibilities, we proceed as in any formal model by deriving comparative statics on our aggregate results as a function of our parameters. If the aggregate results vary with the parameters in a reasonable manner—matching real-world results whenever available—then we can start to have more confidence that our model, despite its simplicity, captures an element of what happens in real government-formation processes.

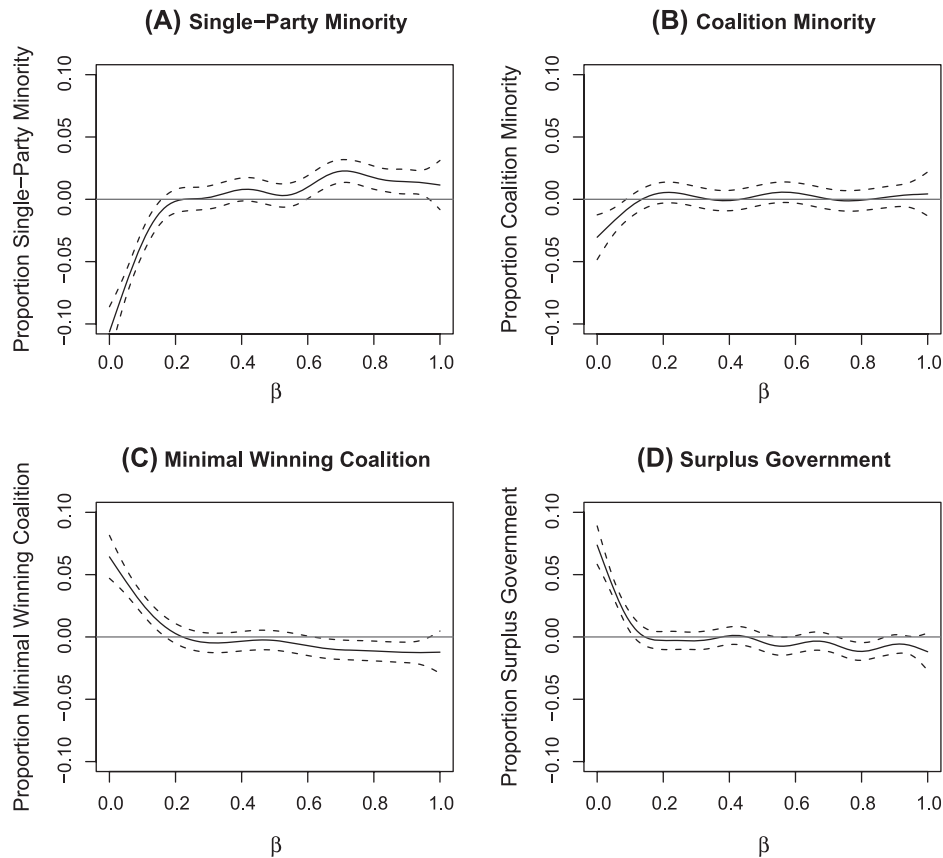
To see how this works, note that the only consistent deviations in the weighted simulation results are a bias toward producing minority governments and a

bias against producing surplus majority governments relative to what we observe in the real world. We know the existence of minority governments in all models, including ours, requires parties to value policy sufficiently highly compared to office, otherwise minority governments would not obtain the necessary majority legislative support to come to power.¹⁹ Further, we expect surplus majority governments to become less common as policy becomes more important and it becomes more difficult to find common ground among larger coalitions. It stands to reason, therefore, that at least some of the bias we find may be arising due to the wide range of policy weights that we consider in our simulations (β varies from 0 to 1). Although we do not have any empirical support for choosing more realistic values of β , we can examine how the distribution of government types in our weighted simulated data varies with β . Our expectation is that minority governments should become more likely, and majority governments less likely, as the relative weight placed on policy, β , increases.

To examine this expectation, we employ a series of generalized additive models using the weighted simulated data to examine the (possibly nonlinear) relationship between the proportion of each government type and β (Beck and Jackman 1998). In Figure 2, we plot the predicted relationship between government type and β for four of our five government types along with 95% pointwise confidence intervals. Specifically, we plot the predicted change in the proportion of each government type against β . Thus, positive values indicate a higher proportion of

¹⁹A separate analysis with $\beta = 0$, so that parties have no interest in policy, yielded no minority governments.

FIGURE 2 Relative Importance of Policy (β) and the Distribution of Government Types



Note: The solid black lines indicate the predicted change in the proportion of each government type at different values of β and are fitted smoothed functions from a generalized additive model in which we regress the proportion of each government type against the relative weight that parties place on policy (β) using our weighted simulated data; the dashed lines represent 95% pointwise confidence intervals.

a particular government type for a given β and negative values indicate a lower proportion. We do not provide a plot for single party majorities because the proportion of this government type was unaffected by the weight that parties place on policy.

Two things are apparent in Figure 2. The first is that minority governments of both types (single party minority and minority coalition) become more common as the weight on policy increases, whereas majority governments of both types (minimal winning coalitions and surplus majority) become less common. This is exactly as expected. The second is that nearly all of the action occurs for low values of β . Above about $\beta = 0.2$, the effect of any additional weight on policy is marginal. Placing enough weight on policy so that it competes with cabinet shares in the utility of parties opens up the bargaining space to allow for minority governments, a fact to which we will return later; further increasing the policy weight has little effect.

In effect, our zero-intelligence model not only produces an empirically plausible distribution of all government types, but it also produces a distribution of government types that varies with parameters such as the weight that parties place on policy in a theoretically intuitive manner.

Allocation of Portfolios

Numerous studies have sought to determine whether portfolios are allocated among governmental parties in proportion to the percentage of legislative seats that these parties contribute to the government. On the whole, these studies evaluate the proportionality of portfolio allocation by regressing a governmental party's share of portfolios against its share of the government's legislative seats. According to Gamson's Law, such regressions should produce a slope coefficient close to one and an intercept coefficient close to

TABLE 2 Portfolio Allocation

Independent Variables	Dependent Variable: <i>Cabinet Portfolio Share</i>		
	Real-World Data: Western Europe	Simulated Data	
		Unweighted	Weighted
Seat Share	0.891*** (0.009)	0.744*** (0.002)	0.771*** (0.002)
Intercept	0.045*** (0.003)	0.102*** (0.001)	0.117*** (0.001)
N	917	125,996	125,996

Note: Cells show coefficients with standard errors clustered on the government in parentheses. The real-world data in the first column are for 14 West European countries—Austria, Belgium, Denmark, Finland, France (5th Republic), Germany, Iceland, Ireland, Italy, Luxembourg, Netherlands, Norway, Portugal, and Sweden—from 1945 to 2000 (Strøm, Müller, and Bergman 2003; Warwick and Druckman 2006). The unweighted simulated data in the second column represent coefficients and standard errors using data from the last elections of 50,000 simulated party systems. The weighted simulated data in the third column come from the same simulated data, but now each party system is weighted by the frequency with which its effective number of parties occurs in the observational data from Western Europe. * $p < 0.10$; ** $p < 0.05$; *** $p < 0.01$ (two-tailed).

zero. Almost every empirical study has reported regression results that approximate these predictions. Although the slope coefficient is always slightly less than one in these studies, indicating that smaller parties are overcompensated and larger parties undercompensated, the overall picture is that portfolios are allocated in a strongly proportional manner. The next test of our model is to see how we fare in this important regard.

In Table 2, we provide information about the slope and intercept coefficients from analyses employing both real-world and simulated data. The first column reports the estimated slope and intercept coefficients from an analysis using observational data from 14 West European countries. These data consist of portfolio allocations to 917 parties that participated in 378 West European governments. In the second column, we report the slope and intercept coefficients using unweighted data from the last elections of 50,000 simulated party systems. In the third column, we report the same results, but each party system is weighted according to the frequency with which its effective number of parties occurs in our observational data. In each of the analyses, we include all government types in recognition of the fact that the determination of government type and the allocation of portfolios occur simultaneously both in our model and in the real world.

Although neither the simulated data nor the real-world data generate an intercept of exactly zero or a slope of exactly one, they come reasonably close and are in line with results from previous empirical analyses (Warwick and Druckman 2001, 2006). As before, appropriately weighting the frequency of our simulated data to approximate the distribution of

party systems in the observational data brings our simulated results even closer to the ones based on real West European governments. Although not often reported, there is some heterogeneity across countries in the slope (and intercept) coefficients that are obtained from regression models similar to the one employed here (Buttard and Golder 2011; Indriðason 2010). It turns out that the slope coefficient based on our weighted simulation data is very close to the slope coefficients obtained when using real-world data from Germany, Italy, and Denmark. Overall, we believe that the degree of concurrence between our simulation and real-world data is quite remarkable, particularly considering that party leaders in our model can propose any allocation of portfolios.

We now explore things further to see what features of our model influence the proportionality of portfolio allocation. As all subsequent results are qualitatively similar using both unweighted and weighted data, we report only analyses using weighted data in what follows.

What Affects the Proportionality of Portfolio Allocation?

We can use our simulated results both to better understand our model and to generate new hypotheses. As a first example, consider the thorny issue of whether there is a formateur bonus in portfolio allocation. While most models do not make predictions on the degree of proportionality of portfolio allocation across all input parameters, all but a few (Bassi 2008; Morelli 1999; Penn 2009) predict that the formateur will receive more than her fair share of the cabinet posts. The empirical evidence for this prediction is mixed at best, though. While Ansolabehere et al. (2005) find evidence for a formateur advantage when it comes to portfolio

allocation, Warwick and Druckman (2006) find evidence of a formateur *dis*advantage.

Given the empirical uncertainty that surrounds this issue and our prior success in matching real-world behavior, we now examine whether our model produces a formateur advantage by adding a dichotomous variable for formateur status to the model specification shown in Table 2. The results from this new regression are shown in the first column of Table 3. The formateur is endogenously defined as the maker of the proposal that the head of state chooses to enact as the new government. As one can see, our model produces a small formateur bonus—formateurs receive a larger share of portfolios than one would predict from simply looking at their contribution to the government’s legislative majority. This formateur bonus occurs due to a lack of foresight on the part of our parties. While actors in Penn’s (2009) model, for example, reason forward to assess the probability that future proposals might be made that would overturn a status quo government that overly benefited a proposer, leading to a smoother distribution of benefits to each party, no such smoothing mechanism exists for our myopic parties. Lacking foresight, the parties in our model make proposals purely at random, and this proposal mechanism produces a formateur advantage.²⁰

In addition to considering extant questions such as whether a formateur bonus exists, we can also use the model to develop new hypotheses. As an example, we now explore the impact of varying our model’s input parameters on the slope coefficient indicating the relationship between a party’s share of the government’s legislative seats and its share of the portfolios.²¹ Recall that we have $2 \times m + 2$ parameters: β , which indicates how much policy is weighted relative to office; the number of parties, m ; and m two-dimensional policy positions for each party. The first two— β and m —are easily varied and have clear substantive interpretations. As a result, we include them directly in our analysis. However, each of the $2 \times m$ party locations have no obvious

²⁰As we demonstrate shortly, the proportionality of portfolio allocation declines with the number of cabinet parties due to the random proposal mechanism in our model. It is the fact that formateur parties, by virtue of being in all governments, are relatively more likely to be in smaller cabinets than nonformateur parties that produces the formateur bonus. Note that the formateur bonus is not caused by the two constitutional constraints in our model—a model with no constitutional constraints produces a similar formateur bonus.

²¹The intercept tends to decrease as the slope increases, and so we do not examine it separately.

TABLE 3 Portfolio Allocation, Formateurs, and Parameter Settings

Dependent Variable: Cabinet Portfolio Share		
Independent Variables	Formateur Advantage	Effect of Parameters
Seat Share	0.714*** (0.002)	1.099*** (0.011)
Intercept	0.107*** (0.001)	0.084*** (0.006)
Formateur Status	0.076*** (0.002)	
<i>Parameters:</i>		
Seat Share \times Number of Parties (m)		-0.098*** (0.002)
Seat Share \times Effective Number of Parties		0.046*** (0.005)
Seat Share \times Ideological Diversity		-0.130*** (0.009)
Seat Share \times Policy Weight (β)		0.040*** (0.006)
Number of Parties (m)		0.027*** (0.001)
Effective Number of Parties		-0.030*** (0.002)
Ideological Diversity		0.058*** (0.005)
Policy Weight (β)		-0.010** (0.003)
N	125,996	125,996
R ²	0.56	0.57

Note: Cells show coefficients with standard errors clustered on the government in parentheses. All models use the weighted simulation data. * $p < 0.10$; ** $p < 0.05$; *** $p < 0.01$ (two-tailed).

independent meaning; rather, combined they dictate the size and ideological diversity of the party system, both of which are of substantive import. We measure these two substantively important concepts via two deterministic functions of our parameters.

The first is the effective number of legislative parties, which measures the size of our simulated party systems. The second captures the ideological diversity of the party systems and is the inverse of Alvarez and Nagler’s (2004) weighted party system compactness measure. It measures the weighted distance of all parties from the ideological center of the party system relative to the dispersion of voter preferences, and is calculated as $\frac{\sum s_j |(z_j - \bar{z})|}{\sigma_c}$, where \bar{z} is the seat share weighted mean of all the parties’ positions, and σ_c is the standard deviation of citizen preferences.

In the second column of Table 3, we report the results from a model in which we add variables capturing the raw number of parties, the effective number of parties, the ideological diversity in the party system, and the relative weight that parties place on policy to those variables already included in the regression presented in the last column of Table 2. In addition, we also include interaction terms between each of these new covariates and the share of legislative seats contributed by each party. The coefficients of interest here are those on the interaction terms, as these coefficients indicate how the size and shape of the party system, as well as the weight that parties place on policy, affect the proportionality of portfolio allocation. In effect, the interaction term coefficients indicate if and how the new covariates affect the slope of the relationship between the share of seats a party contributes to the government's legislative majority and its share of portfolios. For example, a positive coefficient on *Seat Share* \times *Policy Weight* indicates that the slope of this relationship increases, and hence that portfolios are allocated more proportionally, as parties place more weight on policy relative to office.

Given our large sample size, it is not surprising that the coefficients on the interaction terms are all statistically significant. More interesting, though, is that these coefficients are substantively significant, and robust to weighting the simulation data and altering the parties' utility functions. This means that we can make clear statements about how our model's input parameters influence the slope of the relationship between the share of seats a party contributes to the government's majority and its share of cabinet portfolios.

First, consider the variables that relate to party system complexity and size. As the negative coefficient on *Seat Share* \times *Number of Parties* indicates, the marginal effect of *Seat Share* is decreasing in the raw number of parties, m , implying that more complex party systems exhibit less support for Gamson's Law. This occurs due to the nature of proposal making in our model. The more parties there are, the less likely it is that a random draw of cabinet shares will mirror the contributed seat shares.²² In contrast, the positive coefficient on *Seat Share* \times *Effective Number of Parties* indicates that the marginal effect of *Seat Share* is increasing in the effective number of parties

(controlling for ideological diversity and the raw number of parties). This indicates that once one controls for party system complexity, more equal party sizes, as indicated by a larger effective number of parties, imply a more proportional allocation of portfolios.²³ This too occurs due to our random proposal mechanism. The more equal the seat shares, the more likely that the average random proposal for the distribution of portfolios will match the seat shares contributed by each governmental party. Interestingly, the random proposal mechanism produces outcomes that mimic what would be expected intuitively: the fairness of portfolio distributions is increasing in the need for proposers to be fair (more evenly sized parties), but decreasing in party system complexity because complexity opens up additional less-fair options of which to partake. These two effects—party system complexity and size—are by far the strongest in our analysis and account for most of the explained variance in the slope of the relationship between a party's share of the government's legislative majority and its share of portfolios.

Now consider the variables that relate to ideological diversity and policy significance. As the negative coefficient on *Seat Share* \times *Ideological Diversity* indicates, the marginal effect of *Seat Share* decreases in the degree to which the parties are dispersed in the policy space. This is because increased diversity facilitates trade-offs between office and policy, thereby allowing parties to trade off a better government policy in return for a less proportional allocation of portfolios. The positive coefficient on *Seat Share* \times *Policy Weight* indicates that Gamson's Law receives increasing support the more that parties value policy. It turns out, though, that the weight that parties place on policy actually has two opposing effects on the slope of the relationship between a party's share of the government's seats and its share of portfolios. The first is negative and arises from the fact that parties are more likely to trade off office for policy when they place greater weight on policy. The second is positive and arises from the fact that the size of governments becomes smaller as the weight that parties place on policy increases. Recall from Figure 2 that high β produces more minority governments, which have fewer parties in them. As discussed previously, our random proposal mechanism produces a more proportional allocation of portfolios the

²²One can see this simply by drawing two random, normalized vectors of length m and then computing their correlation; it is decreasing in m .

²³Note that by controlling for the raw number of parties, the effective number of parties essentially captures the extent to which legislative parties are of roughly equal size.

fewer governmental parties there are. As the results in Table 3 indicate, this positive effect appears to outweigh the negative effect, i.e., the coefficient on $Seat\ Share \times Policy\ Weight$ is positive.

In sum, support for Gamson's Law increases in the weight that parties place on policy and in the extent to which parties are of roughly equal size; it decreases in the number of parties and in the ideological diversity of the party system. To our knowledge, the existing literature says nothing about the likely relationship between Gamson's Law and party system size, ideological diversity, or the weight that parties place on policy. Thus, our model generates new hypotheses that scholars might test with real-world data.

Bargaining Delays

As we noted earlier, most existing models that rely on Baron-Ferejohn-style bargaining predict little delay in the government formation process, and none predict the distribution of delay observed in the real world. In contrast, our model, like the real world, exhibits considerable variation in the length of time that it takes governments to form. We measure bargaining delays in terms of the number of periods in which parties get to propose potential governments before a new government is installed by the head of state. Though a period in our model is an arbitrary length of time, potentially encompassing a great deal of back and forth between parties, more periods do imply longer delay, and so the number of periods should behave similarly to real-world bargaining delays. In our weighted simulation data, the mean bargaining delay is 29.84 (29.13 for the unweighted data), and delays range anywhere from 1 to the maximum of 100.

More interesting than these simple descriptive statistics, though, is the shape of the distribution of bargaining delays. In Figure 3, we display two kernel density plots of delay distributions. The left one corresponds to real-world data taken from 17 West European democracies. The right one displays our weighted simulation data (the plot for unweighted data is similar), scaled by 2.08 so that the largest delay in our model corresponds to the largest delay (208 days) in our observational data. Other than the hump at 208, which is entirely caused by our decision to artificially censor formation attempts after 100 periods and which generally corresponds to cases in which the arrangement of parties is such that there exists a set of proposals sufficiently attractive so as to be rarely beatable should one become the status quo,

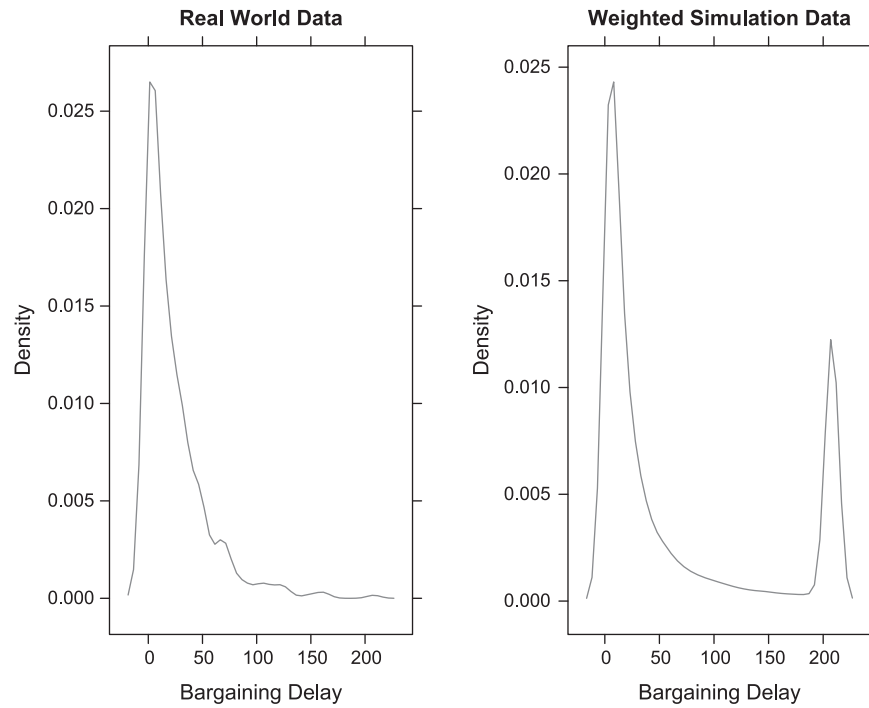
the two distributions look very similar.²⁴ Thus, the model does a good job not only of predicting varying levels of delay, but also in replicating the *distribution* of delays observed in the real world.

What causes delays in the government-formation process? In general, there are few empirical studies of bargaining delays in the real world (Diermeier and van Roozendaal 1998; Golder 2010; Martin and Vanberg 2003). Most of the studies that exist, though, focus on explaining delays in terms of either uncertainty or bargaining complexity. "Uncertainty" refers to not knowing the preferences over policy and office of all the actors whose agreement might be necessary to form a government. "Bargaining complexity" refers to things such as the number of different potential government proposals, both in terms of portfolio allocation and in terms of future governmental policy, as well as the number of politically relevant actors who have to agree on the proposed government. To some extent, scholars have tried to determine whether it is uncertainty or bargaining complexity that is the cause of delays in the formation process. For example, Diermeier and van Roozendaal (1998) argue that delays are caused by uncertainty, whereas Martin and Vanberg (2003) side primarily with bargaining complexity. More recently, Golder (2010) finds that both uncertainty and bargaining complexity matter but in a conditional way.

Our model can provide some light on this debate. Although we do not vary the level of uncertainty, we do vary the level of bargaining complexity in the sense that existing scholars measure bargaining complexity in terms of the number of parties and the ideological diversity of the party system. In Table 4, we examine how policy weight β , the raw number of parties m , the effective number of parties, and ideological diversity affect bargaining delays. The first column presents results from a model in which we use ordinary least squares (OLS) to regress bargaining delays on our covariates using the weighted simulation data. OLS may be problematic in this context, though, because the distribution of delays is not normally distributed (see Figure 3,) delays have to be positive, and some delays are right-censored. As a result, the next column in Table 4 provides results from a Cox proportional hazards model that is

²⁴In the real world, we would expect situations in which an attractive proposal is the status quo to resolve quickly, and not result in a series of failed attempts to find something better when nothing better is forthcoming.

FIGURE 3 The Distribution of Bargaining Delays



Note: Figure 3 presents two kernel density plots of delay distributions. The left one corresponds to real-world data measured in days from 17 West European democracies from 1945 to 1998 (Strøm, Müller, & Bergman 2003). The right one is based on our weighted simulated data where our bargaining delays have been scaled by 2.08 so that the largest delay in our model corresponds to the largest delay (208 days) in our observational data. The hump at 208 in the right plot is entirely caused by our decision to artificially censor government formation attempts after 100 periods; apparent positive densities after this point are just a consequence of the smoothing in the kernel density function.

specifically designed to deal with these issues. One thing to note is that a positive coefficient in the OLS model is equivalent to a negative coefficient in the Cox model and vice versa. This is because we are modeling bargaining delay in the OLS setup but the hazard rate of completing the bargaining in the Cox setup.

The results from the OLS and Cox models effectively match each other. In other words, both estimation strategies indicate that an increase in the number of parties or the weight that parties place on policy lead to longer bargaining delays, whereas an increase in the effective number of parties or party system diversity lead to shorter bargaining delays. Each of these results can be explained quite simply. More parties leads to more complexity in the party system and more wasted proposals due to the fact that obtaining a legislative majority becomes harder. This result matches the claim in the existing literature that bargaining complexity leads to longer bargaining

TABLE 4 Bargaining Delays and Parameter Settings

Independent Variables	OLS	Cox Model
Policy Weight (β)	19.5*** (0.52)	-0.82*** (0.05)
Number of Parties (m)	8.82*** (0.18)	-0.27*** (0.02)
Effective Number of Parties	-20.3*** (0.27)	0.53*** (0.02)
Ideological Diversity	-40.7*** (0.86)	1.42*** (0.07)
Constant	84.8*** (0.79)	
R ²	0.21	
N	50,000	50,000

Note: Cells represent coefficients; standard errors are shown in parentheses. The first column presents results from an ordinary least squares regression, whereas the second column presents results from a Cox proportional hazards model. The Efron method was used to handle ties in the Cox model. * $p < 0.10$; ** $p < 0.05$; *** $p < 0.01$ (two-tailed).

delays. More weight on policy implies that randomly chosen policy proposals are likely to significantly decrease parties' utilities, leading to wasted proposals and hence bargaining delay. In contrast, greater parity in party size and more ideological diversity in the party system both imply a greater likelihood that a random proposal will beat the status quo for more parties, leading to fewer wasted proposals and therefore shorter bargaining delays. To our knowledge, several of these findings are new and have not previously been evaluated in the empirical literature.

Conclusion

The government formation literature is one of the largest in all of political science. Despite significant advances in recent years, there remains a gap between the predictions produced by standard models of government formation that are built on Baron-Ferejohn-style bargaining and what scholars repeatedly observe in the real world. In response to this, there have been calls to rethink our theoretical approach to modeling government formation from the ground up. This is precisely what we do in this article. Specifically, we present a zero-intelligence model in which three or more parties that care about office and policy make random government proposals. The parties in our model are simple—they do not maximize, adapt, or learn; they do not even exhibit bounded rationality. The only constraints that we impose on government formation correspond to the two constitutional constraints that exist in all parliamentary democracies, namely that an incumbent government always exists and that all governments must enjoy majority legislative support.

Despite its deliberately limited structure, our model predicts distributions over portfolio allocation, government types, and bargaining delays that approach those observed in the real world. While existing game-theoretic models of government formation generally seek to explain the consequences of the (assumed) strategic behavior of actors involved in the government formation process, most also seek to predict formation outcomes. Indeed, extant models are often evaluated in terms of whether they can produce particular formation outcomes such as minority governments or proportional portfolio allocation. Although our model makes no pretense to explain the actual bargaining behavior of party leaders as they try to form a government, it is noteworthy that it does better at simultaneously

predicting a variety of observed formation outcomes than existing models do. Significantly, we were also able to demonstrate that the distributions of portfolio allocation, government types, and bargaining delay produced by our model vary in theoretically intuitive ways as model parameters change. Overall, the results from our model indicate that the two constitutional constraints that bind in any parliamentary government formation process are sufficient to approximate many of our most robust empirical regularities, independent of the strategic behavior of party leaders and their bargaining protocols. This suggests that structure, not behavior, may be the most important thing when it comes to explaining government formation outcomes.

We conclude by raising a potentially provocative point about the empirical value of formal models, particularly those used to provide insight into processes generating aggregate data. Using institutionally specific rules that are not supported empirically, such as an alternating offers bargaining protocol, to say something substantive across different settings may not be the best way to conduct empirically relevant formal work. An alternative approach is to encompass as much institutional reality as possible into a model, and when this is not possible, make no additional assumptions. It is not that one should believe that the government formation process is “zero intelligence”; rather, it is that the randomization in our model enables us to feel more confident that the outcomes we observe are truly a function of the institutional details that we know to exist.

Acknowledgments

We thank Siddhartha Bandyopadhyay, Jonathan Bendor, Ken Benoit, Royce Carroll, Jamie Druckman, Sean Ehrlich, Jens Grosser, Michael Laver, William Minozzi, John Patty, Maggie Penn, Bing Powell, Chris Reenock, Norman Schofield, John Scholz, Michael Ting, Paul Warwick, and three anonymous reviewers for their helpful comments on this project. We also thank audiences at the 2010 annual meeting of the Southern Political Science Association, the Political Economy Seminar at Washington University in St. Louis, the Complexity Seminar in the Department of Mathematics at Florida State University, the 2009 annual meeting of the Midwest Political Science Association, Chicago, and the 2008 annual meeting of the American Political Science Association, Boston. In addition, we are grateful to High Performance

Computing at Florida State University for allowing us to use their supercomputer to run our theoretical model.

References

- Alvarez, R. Michael, and Jonathan Nagler. 2004. "Party System Compactness: Measurement and Consequences." *Political Analysis* 12: 46–62.
- Ansolabehere, Stephen, James M. Snyder, Aaron B. Strauss, and Michael M. Ting. 2005. "Voting Weights and Formateur Advantages in the Formation of Coalition Governments." *American Journal of Political Science* 49: 550–63.
- Austen-Smith, David, and Jeffrey Banks. 1988. "Elections, Coalitions, and Legislative Outcomes." *American Political Science Review* 82: 405–22.
- Bandyopadhyay, Siddhartha, and Mandar Oak. 2008. "Coalition Governments in a Model of Parliamentary Democracy." *European Journal of Political Economy* 24: 554–61.
- Baron, David, and John Ferejohn. 1989. "Bargaining in Legislatures." *American Political Science Review* 83: 1181–1206.
- Baron, David P., and Daniel Diermeier. 2001. "Elections, Governments, and Parliaments in Proportional Representation Systems." *Quarterly Journal of Economics* 116: 934–67.
- Bassi, Anna. 2008. "An Endogenous Model of Government Formation." Presented at the Summer Meeting of the Society for Political Methodology, University of Michigan.
- Beck, Nathaniel, and Simon Jackman. 1998. "Beyond Linearity by Default: Generalized Additive Models." *American Journal of Political Science* 42: 596–627.
- Bendor, Jonathan, Daniel Diermeier, David A. Siegel, and Michael Ting. 2011. *A Behavioral Theory of Elections*. Princeton, NJ: Princeton University Press.
- Bendor, Jonathan, Daniel Diermeier, and Michael Ting. 2003. "Recovering Behavioralism: Adaptively Rational Strategic Behavior with Endogenous Aspirations." In *Computational Models in Political Economy*, ed. Ken Kollman, John H. Miller, and Scott E. Page. Cambridge, MA: MIT Press, 213–74.
- Bosch, Antoni, and Shyam Sunder. 2000. "Tracking the Invisible Hand: Convergence of Double Auctions to Competitive Equilibrium." *Computational Economics* 16: 257–84.
- Buttard, Jacquelyn T., and Sona N. Golder. 2011. "Portfolio Allocation and the Vote of No Confidence." Pennsylvania State University. Working Paper.
- Carroll, Royce, and Gary W. Cox. 2007. "The Logic of Gamson's Law: Pre-Electoral Coalitions and Portfolio Allocation." *American Journal of Political Science* 51: 300–13.
- Clark, William Roberts, Matt Golder, and Sona N. Golder. 2009. *Principles of Comparative Politics*. Washington DC: CQ Press.
- Compte, Olivier, and Philippe Jehiel. 2004. "Bargaining Over Randomly Generated Offers: A New Perspective on Multi-Party Bargaining." C.N.R.S. Unpublished manuscript.
- Diermeier, Daniel, and Antonio Merlo. 2000. "Government Turnover in Parliamentary Democracies." *Journal of Economic Theory* 94: 46–79.
- Diermeier, Daniel, and Antonio Merlo. 2004. "An Empirical Investigation of Coalitional Bargaining Procedures." *Journal of Public Economics* 88: 783–97.
- Diermeier, Daniel, Hulya Eraslan, and Antonio Merlo. 2003. "A Structural Model of Government Formation." *Econometrica* 71: 27–70.
- Diermeier, Daniel, and Peter van Roozendaal. 1998. "The Duration of Cabinet Formation Processes in Western Multi-Party Democracies." *British Journal of Political Science* 28: 609–26.
- Farmer, J. Doyne, Paolo Patelli, and Ilija I. Zovko. 2005. "The Predictive Power of Zero Intelligence in Financial Markets." *Proceedings of the National Academy of Sciences* 102: 2254–59.
- Gode, Dhananjay K., and Shyam Sunder. 1993. "Allocative Efficiency of Markets with Zero Intelligence Traders: Market as a Partial Substitute for Individual Rationality." *Journal of Political Economy* 101: 119–37.
- Gode, Dhananjay K., and Shyam Sunder. 1997. "What Makes Markets Allocationally Efficient." *Quarterly Journal of Economics* 112: 603–30.
- Golder, Sona N. 2010. "Bargaining Delays in the Government Formation Process." *Comparative Political Studies* 43: 3–32.
- Indriðason, Indriði. 2010. "Live for Today, Hope for Tomorrow? Rethinking Gamson's Law." Department of Political Science, UC-Riverside. Working Paper.
- Jackson, Matthew O., and Boaz Moselle. 2002. "Coalition and Party Formation in a Legislative Voting Game." *Journal of Economic Theory* 103: 49–87.
- Laver, Michael. 2005. "Policy and the Dynamics of Political Competition." *American Political Science Review* 99: 263–81.
- Laver, Michael, and Ernest Sergenti. 2011. *Party Competition: An Agent-Based Model*. Princeton, NJ: Princeton University Press.
- Laver, Michael, and Norman Schofield. 1998. *Multiparty Government: The Politics of Coalition in Europe*. Ann Arbor: University of Michigan Press.
- Laver, Michael, Scott de Marchi, and Mande Mutlu. 2011. "Negotiation in Legislatures over Government Formation." *Public Choice* 147 (3–4): (285–304).
- Martin, Lanny W., and Georg Vanberg. 2003. "Wasting Time? The Impact of Ideology and Size on Delay in Coalition Formation." *British Journal of Political Science* 33: 323–44.
- Miller, John H., and Scott E. Page. 2007. *Complex Adaptive Systems: An Introduction to Computational Models of Social Life*. Princeton, NJ: Princeton University Press.
- Morelli, Massimo. 1999. "Demand Competition and Policy Compromise in Legislative Bargaining." *American Political Science Review* 93: 809–20.
- Penn, Elizabeth Maggie. 2009. "A Model of Farsighted Voting." *American Journal of Political Science* 53 (1): 36–54.
- Schofield, Norman, and Itai Sened. 2006. *Multiparty Democracy: Elections and Legislative Elections*. New York: Cambridge University Press.
- Snyder, James M., Michael M. Ting, and Stephen Ansolabehere. 2005. "Legislative Bargaining Under Weighted Voting." *American Economic Review* 95: 981–1004.
- Strøm, Kaare. 1990. *Minority Governments and Majority Rule*. Cambridge: Cambridge University Press.
- Strøm, Kaare, Wolfgang C. Müller, and Torbjörn Bergman, eds. 2003. *Delegation and Accountability in Parliamentary Democracies*. Oxford: Oxford University Press.

- Warwick, Paul V., and James N. Druckman. 2001. "Portfolio Salience and the Proportionality of Payoffs on Coalition Governments." *British Journal of Political Science* 31: 627–49.
- Warwick, Paul V., and James N. Druckman. 2006. "The Portfolio Allocation Paradox: An Investigation into the Nature of a Very Strong but Puzzling Relationship." *European Journal of Political Science* 45: 635–65.
- Yildirim, Huseyin. 2007. "Proposal Power and Majority Rule in Multilateral Bargaining with Costly Recognition." *Journal of Economic Theory* 136: 167–96.

Matt Golder is Associate Professor at The Pennsylvania State University, University Park, PA 16802.

Sona N. Golder is Associate Professor at The Pennsylvania State University, University Park, PA 16802.

David A. Siegel is Assistant Professor at Florida State University, Tallahassee, FL 32306.

Copyright of Journal of Politics is the property of Cambridge University Press and its content may not be copied or emailed to multiple sites or posted to a listserv without the copyright holder's express written permission. However, users may print, download, or email articles for individual use.