Computational models of coalition formation

Roni Lehrer

University of Mannheim and University of Frankfurt

12 March 2024

Computational Models and Game Theory: Different approaches to theory building

The goal of theory building for coalition formation is to develop simplification of reality ("models") that identify the key variables and mechanisms that lead to the formation of specific governments. Among other questions, scholars are often interested in how government parties allocate (ministerial) offices within the coalition, and which policies the government will implement.[1]

For decades, formal, game-theoretic approaches have been dominating these endeavors (see Laver 1998 for an overview). In recent years, computational models, i.e., computer simulations of coalition formation, have started to enrich the theoretical toolbox because they can overcome obstacles that game theory faces (de Marchi and Laver 2020).

The obstacles stem from the fact that rationality and utility maximization lie at the heart of game theoretic theory building. Therefore, the algebraic search for equilibria takes center stage in game theorists' work (Morrow 1995, 7–8). This task gets intractable rather quickly when parties compete in multiple arenas at the same time (e.g., parliament and government formation) or when more than a handful of players get involved (Laver and Sergenti 2011, 4–5). This, however, is the empirical reality for virtually all parliamentary democracies in which coalitions form (Bakker, Jolly, and Polk 2012). Game theoretic models of government formation are thus often severe oversimplifications of reality.

Computational methods use computer algorithms to simulate the specified coalition formation game. Hence, they do not require the assumption that players are rational utility-maximizers, and they can handle models that are far more complex (e.g., modelling thousands of voters or the integration of elections, government formation, and policymaking in one model). To derive hypotheses, researchers can vary the input fed to the computational model and investigate how it changes the model's dynamic and outcomes. Effectively, obtaining theoretical predictions from computational models requires skills in coding and applied statistics rather than algebra. All computational models of government formation relevant to this chapter are agent-based models (ABMs). ABMs are a well-developed field of computational theory building

[1] Other relevant questions include which parties are included in government, when and why oversized or minority governments form, what party controls the premiership, and many more. However, computational models mostly address office division and policy compromises, and therefore this chapter focuses on these aspects.

with applications in various social sciences including economics, sociology, epidemiology, and of course political science (De Marchi and Page 2014).

The chapter continues with a brief introduction to ABMs and how researchers can derive hypotheses from them. Then, it summarizes the computational critique on game theoretic coalition formation models. Next, it provides an overview of key ideas and ABMs of coalition formation. The chapter closes with avenues for future computational modelling research of government formation.

The workhorse of computational modeling: Agent-Based Models

ABMs have decisive building blocks.[2] The first indispensable ingredient are agents. Agents represent entities that make choices and are thus similar to players in game theoretic models. A typical government formation ABM contains at least three agents (e.g., parties that seek to form a government). Other government formation ABMs, however, model thousands of agents (e.g., voters as well as parties).

Second, each agent has certain attributes that are relevant to the ABM. For instance, voters often have ideal positions that inform their vote choices. In many applications, parties further hold a certain share of government portfolios. Overall, agent attributes may be exogenously given and fixed throughout the simulations (e.g., voters' ideal positions) or attributes may be endogenous to the ABM (e.g., a party's vote share as determined by aggregating voters' votes).

Third, ABMs have a sequence of actions. More specifically, ABMs specify at what point each agents gets to act, and how their actions affect their own as well as other agents' attributes or the model's outcomes. For instance, in many government formation ABMs, voters get to cast votes for one of several parties first. These votes are then aggregated and determine parties' seat share in parliament. Hence, voters jointly affect parties' attributes. Parties' seat shares, in turn, are the foundation for government bargaining in which parties get to act but not voters.

Finally, each agent requires a behavioral rule. Behavioral rules describe what an agent does when it is her turn. Behavioral rules often reflect a specific idea of agent abilities or agent limitations. Hence, they may be based on agents with very little information at hand or very poor skills to understand what outcomes their choices yield. However, they may also represent highly complex decision-making processes mirroring the ones assumed for utility-maximizing players of game-theoretic models. To illustrate, imagine a party that is invited to join an existing government. It gets to decide whether to join the coalition or to remain in the opposition. A very basic behavioral rule is to simply accept the invitation no matter what. But the party could also flip a coin that determines whether the party accepts or rejects the invitation. A third option is that the party computes whether joining the government yields policy outcomes that are closer to the party's

[2] This section draws heavily on Page and de Marchi (2014) who provide an excellent and more elaborate introduction to ABMs.

ideology. Finally, the party could check if any of the other government parties reneged on past coalition agreements and join the government only if every coalition party has a history of kept promises.

Designing an ABM requires that researchers combine these building blocks – agents, attributes, sequence of actions, and behavioral rules – in meaningful, plausible, and useful ways. It is researchers' responsibility to argue why certain real-world players are depicted by agents while others are left out. They also need to motivate agents' behavioral rules, for instance with reference to the results of prior research. Similarly, readers should judge ABMs with the same standards in mind that help to assess other theoretical work: Does the model provide a plausible explanation for what we observe in the real world?

When combined meaningfully, plausibly, and usefully, these building blocks enable researchers to study very diverse mechanisms and outcomes, some of which are very different from what game theorists study. This is because ABMs face far less limitations in how complex a model can be without becoming analytically intractable. With respect to government formation, for instance, ABMs allow to study individual bargaining rounds and thus bargaining delays, how the distribution of behavioral rules among agents affects government composition, or how voter behavior feeds into the government formation arena. Of course, ABMs can also model various outcomes that game-theoretic models focus on.

From ABMs to Hypotheses

Deriving hypotheses from ABMs follows an experimental logic that political scientists are familiar with from applied statistics: We can consider ABMs as experiments because researchers can vary model input (treatments) while holding everything else constant. We can also observe how changing model input affects model output (outcomes). If we change only a single input parameter between model runs, we can attribute the changes in model outcomes to the change in the input parameter. Using this experimental logic to analyze the relationship between ABM input and outcomes, we can easily derive hypotheses from ABMs. In fact, this logic resembles the logic of comparative statics in game theoretic models.

Let us clarify this approach using an example in which our guiding research question is whether we should expect more government parties in two-party systems or in three-party systems. To answer this question, we build a simple ABM in which parliament is composed of 100 seats. Parliamentary seats are randomly divided between parties. A party is randomly chosen to be the formateur. If the formateur party controls a majority in parliament, it forms a single party government. Otherwise, the formateur choses randomly one of the remaining parties and offers it to join a coalition government. The invited party always accepts the offer and if the two-party coalition controls a majority in parliament, the model outcome is a two-party government. If the two-party coalition is still a parliamentary minority (which is only possible in three-party systems), the model outcome is a three-party government.

Let us apply the experimental logic to this ABM to derive a hypothesis: We call a scenario a simulation with a given model input, e.g., a given number of parties. We run the two-party scenario, record how many government parties it produces, and save the result in a variable which we call N_2. Then, we run the ABM with three parliamentary parties, count the number of government parties, and save the result as N_3. Now, we compute $N_{diff}=N_3-N_2$ which we can interpret as the increase in government parties due to increasing the number of parties from two parties to three. Our hypothesis thus could read: "There are N_{diff} more government parties in three-party systems than in two-party systems." As is easy to see, this is a statement that we can test empirically with real-world data.

While the above strategy clarifies the general logic of deriving hypotheses from ABMs, we need to make one more important adjustment to avoid obtaining flawed hypotheses: Recall that we assigned seats between parties randomly. Due to the different potential divisions of seats, there are 101 different random compositions that the two-party scenario of this ABM gives rise to.[3] In 50 of these random compositions, the formateur party controls a majority in parliament and is thus guaranteed to form a single-party government. In the remaining 51 random compositions, the formateur party lacks the parliamentary majority to form a government by itself, and a two-party government needs to form.

Which of these random compositions' outcomes should we use to compute the difference in government parties between two-party and three-party systems? If we opt for one random composition over the other, our estimates will be unrepresentative of the other 50 % of random compositions. Instead, we can simulate *all* of the different two-party random composition and use the average number of government parties as N_2. This ensures that each random composition contributes as much to the average as the next, and we will obtain a fair representation of random composition in N_2. Technically, this method is called a grid sweep because we layout a grid of all model inputs and sweep through them one by one.

Many ABMs, however, have too many random compositions of each scenario to simulate each of them at acceptable costs (e.g., in terms of time, funds, electricity consumption, and computing power). Instead, researchers can rely on randomly simulating a (potentially high) number of random compositions for each scenario. They can use the average outcome of these random compositions as outcome value for the scenario. In our example ABM, we could, for instance, model 50 random compositions of the two-party scenario and average over them. This makes it dramatically less likely that either of the two outcomes bears overly heavy influence on the outcome value we use for the two-party scenario.

We can also expand ABM analyses beyond a comparison of two scenarios. An immediate extension is to investigate how the number of government parties changes as we allow more than three parties in parliament. A pair-wise comparison of different scenarios will then soon become rather laborious. Again,

[3] Let x be the number of seats the randomly chosen formateur party gets. Then, the other party gets (100-x) seats. Since x can only be an integer in the interval [0, 100], there are 101 random compositions that can be distinguished.

however, applied statistics come to our rescue: We can use regression techniques to learn about the relationship between the number of parliamentary parties and the number of government parties. We can simply run the ABM with different parameter input (i.e., number of parliamentary parties) and different random numbers (i.e., different distribution of seat shares among these parties) many times. After computing the scenario average, we build a dataset of scenario averages and see how they vary with the number of cabinet parties. The relationships that the regressions reveal become our hypotheses for empirical research. They can be tested just like any other hypothesis that is derived form any empirically relevant theory.

Iterative ABMs

Iterative ABMs repeat a part of the ABM several times and agent behavior in a given iteration depends on agent attributes that were determined in the previous iteration. A party, for instance, may behave differently when it entered government last iteration than when it has been a government member for several iterations. These iterative ABMs introduce another randomness-related issue to model outcome representativeness. For an example, consider an ABM of Downsian spatial party competition in which two parties seek to maximize their vote share by shifting their policy positions. Parties' initial positions are randomly drawn. We suppose that parties take turns moving their positions, i.e., party A shifts first, then party B. This concludes a model iteration, and another iteration begins with party A shifting, then party B, and so on. We assume that parties choose the vote-maximizing policy position given its rival's policy position. Since the rival party's policy position is determined in the previous iteration, this is an iterative ABM. As game theory teaches us, agents' behavioral rule implies that parties will position both more centrist than their competitor but as close to it as possible. Eventually, both parties will converge to the median voter position and stay there forever.

One central issue with iterative games is which iterations to include into the dataset we use to derive hypotheses. In the above example, we know that it will usually take several iterations for parties to approach the median position though they will converge there eventually. The standard solution is to simulate a given number of model iterations, called burn-in, that are not included in the dataset for hypothesis derivation. It is beyond this chapter to explain in detail how researchers can determine if a model run has burnt-in. The underlying logic of most respective tests is to see if model outcomes would change substantially if more iterations were discarded as burn-in.

How many model iterations should we record after the burn-in period? The answer to this question depends on the type of ABM at hand. In the two-party spatial competition mentioned above, we know that party positions do not change once their converged. This is ideal indication to use only a single iteration per model run when deriving hypotheses. If we add a third party to the same model, however, game theory predicts that these parties constantly leapfrog each other. However, they do so with some regularity. Hence, we would seek to record as many iterations as needed to cover this cycle at least once.

With this basic knowledge of ABMs at hand, we are now ready to learn about the challenges of classical government formation theories, and how ABMs can help to overcome these issues.

Coalition Formation Theory and its Computational Critique

We start by portraying standard coalition formation theories in a way that facilitates understanding its computational critique: Forming a coalition government entails elements of cooperation and conflict. On the one hand, parties need to collaborate for any of the parties to become a government member and to obtain higher policy and office payoffs.[4] On the other hand, government parties compete for policy influence and offices against each other. This is exactly the situation that game-theoretic bargaining models describe (see Muthoo 2000 for an easy to read introduction).

Unlike most economic bargaining models in which some good is supposed to be traded and two players bargain for the price, coalition bargaining is characterized by many goods (e.g., multiple policy dimension and offices) as well as many players who jointly decide by majority rule (de Marchi and Laver 2020). As the famous "chaos theorems" show (McKelvey 1979; Schofield 1978), majority decisions are unstable in this context: Only in special circumstances, there is no majority that supports an alternative to the incumbent government. This is bad news for game theorists because this implies that virtually no equilibria exist that can help us to derive hypotheses.

Clever game theorists found a solution to the "chaos" issue in the constitutions and unwritten rules of government formation in many countries: the formateur. In real-world politics, the formateur is a well-respected politician, often the leader of a strong parliamentary party, who is tasked by the (non-partisan) head of state to form a government. If she succeeds, she usually becomes Prime Minister. While a formateur seeks to form a government, other potential Prime Ministers tend to abstain from bargaining on alternative coalitions in public (Golder, Golder, and Siegel 2012).

Game theorists adopt the idea of a formateur and make her a player who has a monopoly over proposing governments (including its policy position and office distribution) to other players. If the formateur's offer is accepted, the government forms and payoffs are realized. If it is rejected, a new formateur is chosen and the bargaining game continues. The exact rules for selecting a formateur can differ. Some models select formateurs randomly, others start with the largest party or the incumbent PM party, and yet others randomize the selection but make probabilities proportional to legislative seat shares (see Eraslan and Evdokimov 2019 for an overview). The common denominator of these type of bargaining models is that they predict that the formateur has a significant advantage in government formation. While other coalition

[4] We simplify here and do not count members of a caretaker government as government parties.

partners get enough government spoils to make them support the government, any remaining benefits are retained by the formateur (Bassi 2013, 777).

Computational theorists criticize these theories because an exogenously chosen formateur is not only at odds with what coalition bargaining looks like in reality (Bassi 2013, 778; de Marchi and Laver 2020). The assumption of an exogenously chosen formateur is also somewhat surprising in game theoretic models because it is certainly not utility-maximizing for non-formateur party leaders to refrain from making offers to other parties. Finally, the prediction that formateurs come of best in coalition governments is also empirically heavily challenged (Laver, de Marchi, and Mutlu 2011; Warwick and Druckman 2006). Hence, scholars of coalition bargaining theories require alternative solutions to the "chaos" issue. Computational models provide multiple straightforward solutions.[5]

Computational models of coalition formation differ from standard game theoretic models either by making the formateur choice endogenous to the model or by eliminating the idea of a formateur altogether. Beside these differences, computational theories of government bargaining introduce novel pathways to modeling the entire bargaining process over office and policy, and extend the joint modelling of multiple political arenas.

Before we move on to detailing some of the mentioned differences, note that the critique computational modelists voice against the game-theoretic literature is not principled in nature (e.g., critique toward the idea that party leaders are rational actors). By contrast, scholars who apply computational methods to build government formation theory are massively influenced by the game-theoretic literature and cite it widely. In fact, some authors of computational models have even significantly contributed to the development of game-theoretic government formation models as well.

ABMs of government formation: From critique to contributions

Endogenizing the investiture agenda

Lehrer and Schumacher (2018) present a model that focuses on how parties shift policy positions between elections, and how this affects political representation. Since they also seek to learn how well the government represents citizens in this context, their ABM includes a full-fledged government formation model in which a non-partisan head of state aggregates parties' preferences over coalitions. Effectively, they endogenize the order in which candidate governments are presented to parliament for investiture votes.

Lehrer and Schumacher's (2018)starting point is that parties can value both policy and office. Parties want government policy to be as close as possible to the party's ideal position which may be different from the policy position the party currently advertises to win votes. Parties also wish to hold all seats at the cabinet

[5] In fact, some game-theoretic approaches also provides alternative solutions (e.g., Bassi 2013; Cox 2021).

table. There is a model parameter, α, that determines to what extent a given party's utility depends on office and policy. Parties may fully disregard office and only receive utility from the government policy position being close to its ideal position. Other parties may only care about their office spoils, and yet other parties may value a mix of office and policy.

An iteration of the Lehrer-Schumacher ABM starts with parties announcing their stated policy positions in the two-dimensional policy space for the upcoming election. Then, voters vote for the party that locates closest to them. Since the electoral system is perfectly proportional, the election result also reveals parties' seat shares in parliament to all parties. This concludes the policy-based competition between parties for votes, and the government formation stage begins.

Lehrer and Schumacher (2018) significantly constrain the choices of government policy positions and office distribution. A government's policy position must be its member parties' mean stated policy positions (weighted by party seat shares). Further, a government's office distribution is perfectly proportional to the seat shares the parties contribute to the government (Gamson's law). Since every party knows the election result and all parties' stated policy positions, they know each potential government's policy position and office distribution and hence are aware of the utility each potential government yields for them.

Parties use their knowledge of potential governments' utility to rank them. Parties may rank a government highly even if they are not a member of it, e.g., when the policy component of their utility function outweighs the office component. Next, parties communicate their ranking to the non-partisan head of state. She weights parties' rankings according to their seat share, and suggests the highest ranked government. As a consequence, the model endogenizes the order in which potential governments are presented to parliament for an investiture vote.

Parties then vote on the candidate government that the head of state suggests. Parties support a candidate government if it yields more utility than the incumbent government, and reject it otherwise. Since caretaker governments are not able to fulfil all tasks as regular governments do, the utility parties receive from caretaker governments is discounted by a factor which is also a model parameter. If a candidate government wins both the support of all of its member parties and a majority in parliament, it is installed as government. If it fails, the head of state continues with the next best candidate government. If all candidate governments fail to win sufficient support within the suggested government and in parliament, the outgoing government becomes a caretaker government.

Recall that Lehrer and Schumacher's (2018) focus does not lie on studying government formation but rather on party policy shifts and their implications for political representation (including the government's). Accordingly, they present rather few results with respect to the governments that form. The results they present (Lehrer and Schumacher 2018, Supporting Information 3), however, suggest that governments span more ideological distances between their most extreme members when the number of parties in parliament increases. This effect, however, is somewhat weaker when parties are more policy-motivated or when they take (virtually) identical positions.

Further, they report a negative correlation between the number of parties in the party system and the share of majority governments that form. Interestingly, this effect is weaker when parties are strictly policy-oriented, and non-existent if parties in the entire party system take very similar positions.

To my knowledge, there is no more elaborate analysis of the ABM data, Lehrer and Schumacher (2018) produce. Similarly, there are no empirical tests of these hypotheses either.

Little structure and random behavior: a zero intelligence model

The critique of too much structure, i.e., a formateur who has a monopoly over suggesting governments, in game theoretic coalition formation models inspires Golder, Golder and Siegel (GGS, 2012) to ask how the standard institutions of parliamentary democracy structure government formation. To answer this question, they build a government formation model that deliberately deemphasizes agents' ability to make informed choices. In particular, they rely on a "zero-intelligence" model (Gode and Sunder 1993). In these models, researchers try to impose as little structure as possible in order to highlight the influence of basic constraints. For instance, economists have studied how basic market rules structure outcomes rather than utility-maximizing behavior. The corresponding ABM contains agents that bid in markets randomly subject only to the constraints that they do not sell below costs and that they do not buy above value. They conclude that market efficiency is surprisingly high in this setup (Gode and Sunder 1993).

In the context of parliamentary government formation, GGS argue that two constraints exist in any parliamentary political system (2012, 430): First, there is always an incumbent government, even if it is a caretaker government. Second, any government requires the support of all cabinet parties as well as a majority in parliament.

How do GGS translate these constraints into a zero-intelligence ABM? At the beginning of each model iteration, each party makes a government proposal composed of a random set of government parties (itself included), a random distribution of portfolios among these parties, and a random government policy position in a two-dimensional policy space. Parties care about their share of government portfolios and proximity between the government's policy position and their ideal position. In light of these considerations, each party compares the proposed governments to the incumbent government.[6] If no proposed government wins all of its cabinet parties' support as well as majority support in parliament, parties return to making random government proposals.[7] If a single proposed government meets these criteria, it is installed as the next government and the ABM returns to parties making random government proposals, with the new

[6] In the first round, a dummy government is used as incumbent government.
[7] After 100 rounds without a successful government formation, the game ends.

government serving as incumbent government against which proposals are evaluated. If multiple candidate governments win majority support and cabinet support, the head of state chooses among them randomly.

GGS find that their model performs well when compared with real-word data. They report that the proportions of government types (e.g., minimal winning, minority, oversized governments) corresponds well to what we observe across real-world democracies. Further, the model predicts a correlation between parties' seat share and the distribution of portfolios that is rather similar to empirical observations.

In a critique of GGS, Martin and Vanberg (2014) argue that a government formation model's goal is to correctly predict which government will form given specific circumstances. Hence, the observation that the model's proportion of different government types corresponds to empirical reality is weak evidence for its empirical accuracy at best. The model could predict the correct proportion of government types, yet, predict for each specific government formation opportunity a highly implausible government, e.g., an ideologically very diverse coalition when a single-party majority government can be formed.

Little structure and almost random behavior: a little intelligence model

Warwick (2019) adjusts the GGS zero-intelligence model in order to study why bilateral governments, i.e., governments that bring parties from the two major political blocs together, form rather infrequently. To answer this question, he augments the GGS model with the assumption that governments need to divide portfolios in proportion to parties' parliamentary seat shares, and that government policy positions need to be the seat-weighted average of government parties' policy positions. Both additional assumptions water down the idea of zero-intelligence agents who are neither able to bargain over a government policy position nor distribute portfolios in any systematic way. Instead, GGS expected them to randomize both decisions. Warwick's argument for this adjustment is that empirical evidence suggests that coalition governments tend to gravitate toward these results across countries (703-704). Warwick finds that these additions (as well additional assumptions concerning voters' motivation) have very little effect on whether governments include parties from the two major political blocks.

Analyzing the ABM output, Warwick (2019) finds that government policy positions are more responsive to mean vote choices than in the GGS model. Further, the government seat shares or portfolio shares–the two are assumed to be proportional–are rather realistically distributed between left and right-wing parties.

Modeling the government cycle and log rolling

De Marchi and Laver's (2023) book advances the ABM approach in several ways significantly. We discuss them in turns.

Government formation: log rolling and tabled issues

With respect to government formation, De Marchi and Laver set up a model in which in all party leaders can suggest a candidate government each at the same time. Hence, each party leader has to find a potential government to suggest. This implies that party leaders need to decide on which parties should form the government, what the government's policy position should be, and how portfolios should be divided between parties.

Which government parties will party leaders suggest? Government party leaders' behavioral rule dictates that they suggest an incumbent government if they are a member whenever it controls a majority in parliament. If it does not, party leaders compare their utility from two potential governments (called proto-governments) and suggest the one they prefer. The first proto-government included in this comparison is composed of those parties that exceed the majority threshold most narrowly. Hence, it does not matter how ideologically diverse the government is or whether it includes many more parties than other majority governments. It only needs to exceed the majority threshold more narrowly than any alternative combination of parties. The second proto-government includes those parties that are closest to the suggesting party ideologically, and together control a majority in parliament. Obviously, these behavioral rules for proto-government selection reflect ideas from government formation with office-motivated or policy-motivated parties respectively.

Which government program do party leaders suggest? With respect to finding a government's program, De Marchi and Laver's argue that previous theories of government bargaining severely limit the number of policy issues that parties bargain over. Typically, researchers assume that individual issues can be grouped into ideological dimensions such as the left-right dimension or a liberal-conservative dimension. Irrespective of these dimensions' substantial interpretations, they allow researchers reduce the complexity of their bargaining models severely.

De Marchi and Laver (2020), by contrast, maintain that real government bargaining is less about ideological dimensions and more about working through a long list of individual issues that appear in parties' election manifestos. Importantly, either government parties can find a joint position on an issue, or they can agree to disagree and table the issue. Tabling an issue may be particularly interesting when a specific issue has little urgency to be dealt with. External shocks, such as natural disasters, scandals, or significant changes in public opinion, may change an issue's urgency rather quickly.

With this in mind, they devise a government bargaining model, which is a key element of their government formation model. It starts with a high number of policy issues over which parties have binary preferences. That is, they either agree or disagree with a suggested policy change. Further, specific issues can be more or less salient to a party. Using De Marchi and Laver's notation, let A(**1**, 1) and B(0, **0**) denote two parties' positions on two issues and their respective saliencies. Party A supports both issues whereas party B opposes both. While party A attaches high salience to the first issue, party B cares more about the second issue. It is easy to see that if these parties bargain on a common policy position in order to form a coalition government, they will likely agree to the joint policy position G(1, 0). Effectively, parties give in on issues they care about less to win concessions on issue their care about more.

We can now extend this logic to government bargaining with multiple parties. To begin with, all issues on which all parties in a potential coalition agree are immediately included in their government program. Then, all issues on which more than one party would have to flip its position in order to obtain unanimity are tabled. Put differently, the model assumes that only issues with a single outlier can be resolved by bargaining. The ABM chooses one of the remaining single-outlier issues at random, and the corresponding outlier party checks if another single-outlier issue exists, that allows for a mutually beneficial bilateral trade. In particular, each of the trading parties have to trade a position on an issue that carries less salience than the issue on which their position prevails. This bargaining continues until no beneficial trades between two parties are left. Any remaining issues are tabled.

Governance Cycle

De Marchi and Laver's second major critique of previous models of government formation concerns that they do not jointly model what they call the governance cycle: elections, government formation, and the struggle of governments to maintain parliamentary support. While models of the individual elements are parsimonious and insightful, joint models of multiple stages allow analyzing what consequences actions in one political arena have for some other arena. For instance, an election may result in a government formation context that requires parties from opposed blocks to cooperate to form a government (e.g., because radical pariah parties deprive classical blocks of their options to form majority coalitions). Such a unity government is arguably less stable than governments that form within well-known blocks.

This critique of game-theoretic models is not new (Lupia and Strøm 1995), and two arenas have been modelled jointly (Austen-Smith and Banks 1988). However, even more arenas is beyond what game-theoretic models have been able to deliver so far.

De Marchi and Laver (2023), thus, present an ABM in which they model elections, government formation, government survival. The most interesting implementation of such linkages occurs in the context of government formation and government survival. Their ABM allows for what they call agenda shocks, i.e., sudden changes of the issues that are on the government's agenda. De Marchi and Laver cite the start of the COVID pandemic as an example. In early 2020, issues were suddenly on the agenda that many governments did not expect or prepare for.

The model captures the effects of agenda shocks by forcing governments to take a position on an issue that it tabled. Recall that undisputed issues are immediately accepted as part of the government program, and hence issues will only be tabled if there is some disagreement between government parties over the issue. As a consequence, forcing the government to implement one of the two possible policies with respect to the previously tabled issue makes at least one government party worse off. The more government parties are made worse off, and the more salient the "untabled" issue to them, the greater the probability that a government member prefers leaving the government to maintaining it. Since there is an almost deterministic relationship between the number of random agenda shocks and the probability of government termination (the more agenda shocks to a government, the more likely government termination becomes) de Marchi and Laver (2023) do not explicitly model these steps.

They also investigate how preference shocks affect government survival. Preference shocks include random changes to parties' position on specific issues, the salience they put to specific issues, and their value of policy over office. To learn how strongly these shocks affect government termination, they run several simulations.

Agent-Based artificial intelligence

De Marchi and Laver's final contribution takes a step toward more utility maximizing behavior in ABMs. They adopt a Monte Carlo Counterfactual Regret (MCCFR) approach that was developed by scholars of artificial intelligence to solve complex games. Effectively, MCCFR plays a game-theoretic game over and over and learns how to make good choices for all players. At a given node in a game, its output is a probability distribution over actions that maximizes the utility of the player whose turn it is to act. MCCFR can provide such probability distributions for many nodes and many players even in games that are too complex to solve algebraically. It obtains these probability distributions by iterating through the following steps: At a given node, it plays the current probability distribution over actions (starting with a uniform distribution in the first iteration), evaluates which alternative choices would have decreased (or improved) the current player's utility, and adjusts the probability distribution accordingly by making beneficial actions more likely and harmful actions less likely. In the next iteration of the MCCFR procedure, the updated probability distribution is used to sample an action and so on. The Monte Carlo element of MCCFR refers to not having players play every potential game history in every iteration but instead make them sample a single history.

Note that the results the MCCFR procedure returns are similar to results which rational actors who "look down the game tree" produces. Of course, the MCCFR technique does not return best responses according to their formal definition, however, it approaches them reasonably well and is applicable in an ABM context.

Unfortunately, MCCFR procedures become very computing intensive for complex games such as de Marchi and Laver's (2023) election cycle model. To cope, de Marchi and Laver (2023) massively simplify the game in multiple ways. For instance, they limit the number of policy issues to fourteen, they do not allow for tabled issues in coalition agreements, they limit the number of cabinet seats to eleven, and they assume that all parties value office and policy equally when evaluating potential coalitions. Most importantly, however, they use the MCCFR procedure to obtain a probability distribution for two players only.[8] All other players use heuristics.

Simulating and analyzing their model, de Marchi and Laver (2023) predict that a higher dimensionality of the policy issue space decreases the probability of cycling majorities in parliament. A similar effect arises when politicians are more likely to table issues. Further, single party cabinets become more likely when politicians value office spoils more and are more risk-averse with respect to tabling issues. Finally, governments tend to be less stable when politicians value office spoils highly and are optimistic about tabling issues.

What is next?

Computational models of government formation are certainly powerful tools to understand which governments form and why. However, their application remains rare in the literature on government formation. This is a pity since there are many research questions that lend themselves to being addresses with ABMs.

To begin with, the existing models presented above, and their implication have not been fully studied yet. For instance, the available data from the Lehrer and Schumacher (2018) model can be analyzed in more detail to understand the dynamics of government formation in light of an endogenous investiture agenda. Further, they can be compared to game theoretic predictions of endogenous formateurs (Bassi 2013), and to real-world data. Similarly, the GGS model's (2012) data have not been fully exploited yet. It remains, for instance, unstudied to what extent centrist governments or one-sided governments are easily defeated by alternative governments. Finally, the de Marchi and Laver (2020) bargaining model can be used to study government formation delays (S. N. Golder 2010).

Further, questions of substantial portfolio distribution have not been addressed yet with computational models. Up until now, government formation theories mostly assume that all government portfolios are equally valued by parties. Empirical evidence, however, indicates that parties value portfolios differently (Bäck, Debus, and Dumont 2011). With their ability to model many policy dimensions at the same time, ABM can easily explore how party preferences over specific government portfolios enhance or limit parties ability to successfully compromise in government formation bargaining.

Despite the major steps that de Marchi and Laver (2023) take toward jointly modelling voters, government formation, and government termination, they stopped short of modelling the entire government cycle. Since voters respond to parties' performances in government (Fortunato 2021), and parties are likely to respond to poor government performance with terminating the government (Lupia and Strøm 1995), a full model

[8] In fact, they use the MCCFR technique for one player only. However, by alternating between the two largest parties in parliament, they obtain results for two parties.

of the government cycle needs to model how voters respond to government compromises – potentially in the context of policy shocks.

Finally, many computational models' predictions have only been partially tested against empirical data at best. This, however, is an indispensable element of improving scientific theories. And since political scientists have collected ample data on government formation processes, several critical steps to testing computational theories have been taken already.

Appendix: Software for ABMs

Working with ABMs requires two very different computational tasks. On the one hand, ABMs need to be simulated. On the other hand, they need to be analyzed with applied statistics to derive hypotheses. In terms of software suitability, these are very different needs. I will, therefore, discuss different software packages' pros and cons for ABMs with respect to these different needs.

R (www.r-project.org) is a free programming language that specializes in statistical computing. In recent years, R has become more and more popular in political science programs across the globe because it is available for free, flexible, and allows to implement more state-of-the-art techniques than for-profit software.

Python (www.python.org) is a very broad programming language that has a broad user base that develops and maintains a wide array of addon packages. It is widely used in several fields of scientific computing including simulations and data science.

NetLogo (https://ccl.northwestern.edu/netlogo) is a software environment that was specially developed for ABMs. It is well known to be relatively simple to use. However, other NetLogo skills are unlikely to help researchers work on non-ABM projects.

How well are these software packages suited to simulate ABMs? While one can implement ABM simulation tasks in R, it becomes more time-efficient rather quickly to learn a new coding language than simulating complex ABMs in R. This is because R is optimized to work with matrices which makes perfect sense for a statistical tool. ABMs, however, are usually iterative and sequential in the sense that one agent acts, and then another agent responds to what the first agent did. In terms of coding, these iterations and sequential processes are usually translated into loops which are particularly slow in R. Python is certainly more powerful in simulating ABMs than R. This is mostly because it is much faster than R in computing in the context of loops. Python also benefits from libraries like Mesa and AgentPy that provide basic tools that make it easier to build ABMs efficiently. NetLogo is arguably the easiest software to simulate ABMs. However, it is also the least flexible of the three alternatives. Yet, for researchers who shy away from heavy programming, NetLogo is a viable option.

And how well can we analyze ABM data in these packages? When it comes to analyzing ABM model output, R is in a very strong position. R performs strongly because it offers a number of regression and machine learning techniques that can help to find relevant patterns. Python is also well equipped to analyze ABM data. Its pandas and numpy libraries have been used heavily by data scientists of all backgrounds. Similarly, Python can draw on a number of machine learning libraries. Finally, NetLogo is no realistic alternative to analyze the ABM output because it lacks the required strong implementation of regression techniques.

There are several other software packages that researchers can use to simulate ABMs and to analyze their output. However, they have played less of a role in political science application. These include Julia, C, C++, and many more.

If ABMs are strongly iterative, e.g., because burn-in periods are rather long, researchers may find it helpful to optimize their code for speed. While coders are often encouraged to write clean code, i.e., code that it easily readable, it may become necessary to trade clean code for speed. When researchers need to simulate many different model scenarios or many random components, parallelization, i.e., having different computing cores simulate different model run can be very helpful. Many universities also provide high-performance computers that may have several hundreds of computing cores – far more than the single digit count standard computers have.

References

Austen-Smith, David, and Jeffrey Banks. 1988. "Elections, Coalitions, and Legislative Outcomes." *American Political Science Review* 82(2): 405–22.

Bäck, Hanna, Marc Debus, and Patrick Dumont. 2011. "Who Gets What in Coalition Governments? Predictors of Portfolio Allocation in Parliamentary Democracies: Who Gets What in Coalition Governments?" *European Journal of Political Research* 50(4): 441–78.

Bakker, Ryan, Seth Jolly, and Jonathan Polk. 2012. "Complexity in the European Party Space: Exploring Dimensionality with Experts." *European Union Politics* 13(2): 219–45.

Bassi, Anna. 2013. "A Model of Endogenous Government Formation." *American Journal of Political Science* 57(4): 777–93.

Cox, Gary W. 2021. "Nonunitary Parties, Government Formation, and Gamson's Law." *American Political Science Review* 115(3): 917–30.

De Marchi, Scott, and Michael Laver. 2023. *The Governance Cycle in Parliamentary Democracies: A Computational Social Science Approach*. 1st ed. Cambridge University Press. https://www.cambridge.org/core/product/identifier/9781009315449/type/book (July 27, 2023).

De Marchi, Scott, and Scott E. Page. 2014. "Agent-Based Models." *Annual Review of Political Science* 17(1): 1–20.

Eraslan, Hülya, and Kirill S. Evdokimov. 2019. "Legislative and Multilateral Bargaining." *Annual Review of Economics* 11(1): 443–72.

Fortunato, David. 2021. *The Cycle of Coalition: How Parties and Voters Interact under Coalition Governance*. 1st ed. Cambridge University Press. https://www.cambridge.org/core/product/identifier/9781108877053/type/book (March 17, 2023).

Galnoor, Itzhak, and Dana Blander. 2018. *The Handbook of Israel's Political System*. 1st ed. Cambridge University Press. https://www.cambridge.org/core/product/identifier/9781316160978/type/book (August 3, 2023).

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(1): 119–37.

Golder, Matt, Sona N. Golder, and David A. Siegel. 2012. "Modeling the Institutional Foundation of Parliamentary Government Formation." *The Journal of Politics* 74(2): 427–45.

Golder, Sona N. 2010. "Bargaining Delays in the Government Formation Process." *Comparative Political Studies* 43(1): 3–32.

Laver, Michael. 1998. "MODELS OF GOVERNMENT FORMATION." *Annual Review of Political Science* 1(1): 1–25.

Laver, Michael, Scott de Marchi, and Hande Mutlu. 2011. "Negotiation in Legislatures over Government Formation." *Public Choice* 147(3–4): 285–304.

Laver, Michael, and Ernest Sergenti. 2011. *Party Competition: An Agent-Based Model: An Agent-Based Model*. Princeton University Press. https://www.degruyter.com/document/doi/10.1515/9781400840328/html (August 3, 2023).

Lehrer, Roni, and Gijs Schumacher. 2018. "Governator vs. Hunter and Aggregator: A Simulation of Party Competition with Vote-Seeking and Office-Seeking Rules" ed. Asim Zia. *PLOS ONE* 13(2): e0191649.

Lupia, Arthur, and Kaare Strøm. 1995. "Coalition Termination and the Strategic Timing of Parliamentary Elections." *American Political Science Review* 89(3): 648–65.

de Marchi, Scott, and Michael Laver. 2020. "Government Formation as Logrolling in High-Dimensional Issue Spaces." *The Journal of Politics* 82(2): 543–58.

Martin, Lanny W., and Georg Vanberg. 2014. "A Step in the Wrong Direction: An Appraisal of the Zero-Intelligence Model of Government Formation." *The Journal of Politics* 76(4): 873–79.

McKelvey, Richard D. 1979. "General Conditions for Global Intransitivities in Formal Voting Models." *Econometrica* 47(5): 1085.

Morrow, James D. 1995. *Game Theory for Political Scientists*. Princeton University Press. https://www.degruyter.com/document/doi/10.1515/9780691213200/html (August 16, 2023).

Muthoo, Abhinay. 2000. "A Non-Technical Introduction to Bargaining Theory." 1(2): 22.

Schofield, N. 1978. "Instability of Simple Dynamic Games." *The Review of Economic Studies* 45(3): 575–94.

Warwick, Paul V. 2019. "Engendering Bilateralism: An Exploration of Possible Sources of Non-Centrism in Government Formation." *Political Science Research and Methods* 7(04): 699–715.

Warwick, Paul V., and James N. Druckman. 2006. "The Portfolio Allocation Paradox: An Investigation into the Nature of a Very Strong but Puzzling Relationship." *European Journal of Political Research* 45(4): 635–65.