The Governance Cycle in Parliamentary Democracies
Cambridge Studies in Comparative Politics
de Marchi and Laver, 2023

Introduction

> The field covered in this book is very limited, and we approach it in this sense of modesty… Its first applications are necessarily to elementary problems where the result has never been in doubt and no theory is actually required. At this early stage the application serves to corroborate the theory. The next stage develops when the theory is applied to somewhat more complicated situations in which it may already lead to a certain extent beyond the obvious and the familiar. Here theory and application corroborate each other mutually. Beyond this lies the field of real success: genuine prediction by theory. It is well known that all mathematized sciences have gone through these successive phases of evolution. (Von Neumann and Morgenstern 2007: 8).

We started this three-year project motivated by a long-standing puzzle in the literature on parliamentary democracies. Governments are responsible to parliaments and must maintain the support of a majority of legislators. The social choice tradition of theoretical work on this subject implies that governments should have a hard time not only getting into office but staying there – because some legislative majority can always find a "better" alternative. We do sometimes see this type of instability – historically in Italy and recently in Israel, for example. Typically, however, governments do form and do stay in power, sometimes for quite a long time. In a

DRAFT

nutshell, despite what some theorists have called "the generic instability of majority rule" (Schofield 1983), why do we not tend to observe this instability in a real world populated by living, breathing, politicians? And why does this stability sometimes vanish, flipping the world into a state where it becomes difficult to build and maintain a government?

Work in the tradition of non-cooperative game theory attacks this puzzle by describing it in terms of a radically simplified bargaining environment and making restrictive assumptions about actions available to senior politicians.[1] We ourselves have a hard time, however, imagining that highly-motivated, successful, and seasoned senior politicians – playing for the very highest stakes in the business – are bound by the assumptions found in these models. To take just one example, the most influential approach to this problem (Rubenstein 1982, Baron and Ferejohn 1989, Baron 1991, Ansolabehere, et al. 2005, Snyder, et al. 2005) simplifies the bargaining environment by assuming that some exogenous mechanism allows one, and only one, politician at a time to make a take-it-or-leave-it offer. Given an exogenously imposed sequence of offers, the implication is that the first proposer (*formateur*) has a substantial advantage.

We do not, however, think it is in any way reasonable to assume that the most seasoned and sophisticated politicians in the business, doing the thing they most care about, sit quietly on their hands waiting for the talking stick to be passed in their direction. In contrast, we believe that – whether this happens in the cold light of day or in the dark shadows of what we think of as "backstage" – no senior politician can be banned from making any proposal to anyone at any time. Not only will several proposals typically be under consideration at the same time, no

[1] In the last few years, a number of very interesting studies have been conducted in this tradition, with a focus on providing an explanation for the empirical regularity of Gamson's law (Martin and Vanberg 2020, Cox 2021).

DRAFT

exogenous mechanism can dictate the order in which these are considered. In short, assuming that an exogenously imposed *formateur* sequence exists might *theoretically* address the generic instability of majority rule, but does this at the cost of a model that does not *realistically* apply to the world we live in. Our approach here is to privilege the level of realism needed for a plausible description of government formation, recognizing that this will result in complex theoretical models that may well be analytically intractable using traditional deductive methods. In doing this, we leverage two strands of previously published work.

The first strand implies that more realistic and complex strategic contexts are, perhaps counter-intuitively, more likely to produce structure and thereby prevent chaotic outcomes (Kollman, et al. 1992, 1998; Miller and Page 2009). In very simple settings, the intuition is that human actors will quickly see all possible courses of action and act in their own self-interest. In the context of majority voting over the making and breaking of governments, this underlies predictions of multiple equilibria and cyclical majorities. Complex and difficult environments, in contrast, typically mean that it is hard for people to identify and evaluate every conceivable possibility in real time. This forces them to search for "good" solutions which are not provably the "best" and limits their ability to find counter-proposals, thereby enhancing stability.

The second strand flows from this and puts much more emphasis than many formal theorists on actual human behavior. In complex games, it can be difficult or impossible for real humans to figure out the perfectly optimal strategies assumed by most formal theoretic models. Canonical examples of such games – including bridge and poker – resist these purely deductive approaches. As a result, formal theorists often simplify games, thereby making it possible to assume that behavior consists of straightforward deduction. Recent advances in reinforcement learning in artificial intelligence (AI), however, allow researchers to attack complex games, in

DRAFT

many cases achieving performance better than that of expert humans (Billings, et al. 2003, Brown, et al. 2017). Rather than simplifying the game of democratic governance so radically that it can be solved by backwards induction, therefore, we set out here to explore effective – though not provably optimal – behavioral rules that scale to complex games.

Populating a more complex and realistic game with agents using behavioral rules that mimic those which might be used by expert humans, we find that in many settings our model produces stable outcomes – Condorcet winning cabinets – not generic instability. This approach had been proposed, but not previously completed, by other scholars as a potential answer to similar puzzles in the social sciences. Douglas North emphasized the central role of behavior and complexity in his later work and saw this as one answer to the question of how institutions emerge (Denzau and North 2000, Munger 2020). Similarly, Elinor Ostrom focused on the use of heuristics to solve common pool resource problems in complex environments (Ostrom 2010). Increased availability of computation and advances in reinforcement learning have now made this approach much easier to implement.

We were encouraged from our early work as part of this research program (de Marchi and Laver 2020) that not only showed that stable governments were possible but also allowed us to predict the length of negotiations – bargaining delays – during government formation. We took this as quite strong support for our approach. While there are *empirical* accounts of the duration of government formation negotiations in parliamentary systems, there is no rigorous theoretical model explaining the empirical findings.[2] In our research, we took a direct output of

[2] The main arc of empirical work on duration of negotiations starts with Diermeier and Van Roozendaal (1998) and continues with Martin and Vanberg (2003) to Golder (2010). On the general difficulty in linking theory and empirics in this area, see Diermeier et al. (2003).

our theoretical model – the likelihood of finding a Condorcet winner in each empirical setting – and used this as a measure of bargaining complexity to predict the length of negotiations. A purely analytical output of our theoretical model thus predicts a real-world dependent variable, providing evidence for our causal mechanisms.

The motivation for this book was to extend our model and use it to predict other important real-world outcomes that have resisted prior efforts. In particular, we were interested in predicting, on a case-by-case basis, which coalition is likely to form after elections in parliamentary systems. This is hard, given the huge number of different possible proposals for government. While making predictions is not the current fashion in political science, it is our belief that this is the best way to test the main implications of any theory. Predict an outcome of substantive interest and, as new data arrives, continually expose the theory to new tests. The composition of governments that form in parliamentary systems is a nearly perfect dependent variable in this sense. It is of enormous substantive interest and easily observable. There are a relatively large number of cases and new data are generated by every new election.

Even better for our purposes, government formation in parliamentary democracies is a complex strategic game played by highly motivated politicians. Our expectation is that if we can successfully analyze this, we most likely will be able to apply our model to other contexts in which bargaining occurs. Examples include alliance formation in the international system and intraparty bargaining in the US Congress.

Complex strategic models of the sort we investigate in this book, along with the behavioral algorithms we use to describe the decisions made by senior politicians, involve a large number of modeling choices. The choices could, in principle, limit the generalizability of our model and produce "curve-fitted" results. By focusing on predicting a real-world outcome

DRAFT

such as government formation we can, however, see if our model is general and examine the impact of different parameter values. *Theoretical* generalizability is increased by committing to an *empirical* target. Simply put, we are trying to fulfill Von Neuman and Morgenstern's original goal: "genuine prediction by theory". Of the universe of formal models we could write down about government formation, we are interested here in those that have a direct mapping to real-world political phenomena.

In this book, we describe both the methodology we used and the substantive results we derived. Methodologically, there are several good treatments of how to build computational social science models in the area of *empirical* work (Alvarez 2016). At present, however, there are few treatments of computational social science applied to *theory*.[3] We hope this book, and the availability of our code and data, will help to fill that gap.

Substantively, we are happy to report that our initial results are encouraging; we feel we have achieved our main goals. The models we present here are able to predict: which governments form after an election; how long negotiations take during this process; and how long these governments are likely to last. One issue that deserves attention is that we, unfortunately, do not have other, existing research with which to compare our findings. Prior theoretical work has not typically aimed at these targets. It is our hope, however, that by providing a first step in prediction, other researchers will build new models, improving on what we have done here, and measure their progress against our baseline. One way in which science progresses is when a community of researchers uses common benchmarks to compare the performance of models, linking theory to empirics.[4]

[3] The exception is Laver (2020a, 2020b) which is focused on agent-based models.

[4] For an overview, see *https://www.nsf.gov/sbe/ses/polisci/reports/pdf/eitmreport.pdf*.

DRAFT

In addition to success at prediction, we also have a set of substantive insights that are distinct from earlier research. Loosely, earlier work typically argues that government formation:

i. is led by a proposer with considerable advantages in bargaining;
ii. involves short negotiations where the first proposal is accepted;
iii. tends to result in minimum winning coalitions (not single-party, minority, or surplus governments);
iv. at most considers government policy on one or two latent "dimensions".

Our results on the birth and death of governments in parliamentary democracies show that:

i. senior politicians make and are aware of multiple proposals;
ii. negotiations may be lengthy, depending on the modeled complexity of the strategic environment;
iii. negotiations may result in minority, minimum winning, or surplus governments;
iv. policy matters a great deal and parties negotiate over high-dimensional joint policy programs. Logrolling is a central part of this and parties with negatively correlated saliences over policy have an advantage in forming coalitions.

Crucially, we link the endless "governance cycle" of elections, government formation, and government survival into an integrated modeling architecture and show that even when governments form, they are not all equally successful in negotiating a joint policy program. Salient and contentious issues may be left unresolved (tabled) and may, at some point in the future, be forced onto the agenda by an exogenous shock. Since such shocks cause legislators to think again about whether they prefer some alternative to the incumbent, the survival of governments refers directly to the (implied) process of government formation. This allows us again to connect a key output of our formal model directly to an empirical model of government

DRAFT

durations. This expands on an approach pioneered by Lupia and Strom (1995) in two ways: it generalizes from three to any number of parties and multiple issue dimensions; it models an explicit mechanism that explains the role of unanticipated shocks in bringing down governments.

Finally, we show that the behavior of senior politicians is not only quite complex, but poorly described by backwards induction (the algorithm employed by non-cooperative game theorists). We explore the strategic capabilities of senior politicians by building two quite different types of models, grounded in different behavioral assumptions. The first is an agent-based model, which uses relatively sophisticated rules-of-thumb – heuristics – to model behavior. The second is an AI approach based on reinforcement learning, which allows agents to teach themselves how to play the game and find "near-Nash" strategies. We find that the performance of both models is similar, which indicates that senior politicians – even though they are experts in their domains, have access to professional staffs, and are playing for high stakes – may not be fully strategic and instead use sophisticated heuristics to make decisions.

We therefore model senior politicians as people who, if they wish to be successful while playing a complex game, are what we call "functionally" rational. They are clearly very good at finding *effective* strategies. Given the enormous complexity of their strategic environment, however, being successful also means recognizing that it makes no sense – indeed is functionally *irrational* – to obsess on finding *perfectly optimal* strategies.

We have left many questions on the table for future work. For example, we are able to "back out" from our model unobservable features of decision-making by party leaders, including their relative taste for the perks of office or the policy outputs of government. We also hope to explore which empirical settings have more policy-motivated party leaders given that this is a crucial link in a representative democracy between elected officials and the constituents they

DRAFT

represent. In this book, we avoid detailed case-specific work and focus on general results, but differences surely exist between settings. An alternative use of our model is to calibrate it closely to some particular setting of interest and use it to power a theoretically structured case study. This could, for example, help us understand why some countries – Greece during the debt crisis or Israel today – enter into periods of instability.

DRAFT