CHAPTER ONE

# Modeling Multiparty Competition

We hold these truths to be self-evident:

- *Politics is dynamic*. It evolves. It never stops; It is never at, nor *en route* to, some static equilibrium. Politics evolves.
- *Politics is complex*. Political outputs today feed back as input to the political process tomorrow.
- *Politicians are diverse*. In particular, different politicians attack the same problem in different ways.
- *Politics is not random*. Systematic patterns in political outcomes invite systemic predictions, making a political "science" possible.

Politics in modern democracies is largely the politics of representation. It concerns how the needs and desires, the hopes and fears of ordinary citizens affect national decision making at the highest level, doing this via public representatives who are chosen by citizens in free and fair elections. Representative politics is to a large extent about party competition: about how a small number of organized political parties offer options to a large number of voters, who choose at election time between alternative teams of public representatives. Party competition is therefore a core concern for everyone, be they professional political scientist or ordinary decent civilian, who cares about politics in democratic societies.

We believe that party competition is a complex and evolving dynamic process that can be analyzed in a rigorous scientific manner. More precisely, we analyze the dynamics of *multi*party competition, by which we mean competition for voters' support among more than two parties, opening up the possibility that no single party wins a majority of votes cast. Figure 1.1 plots some observations of multiparty competition in the Netherlands over the period 1970–2005. The left panel shows positions of the three main Dutch parties on a left-right scale of party ideology, estimated from their party manifestos.[1] The right panel shows support for these same parties in the Dutch electorate, estimated using Eurobarometer surveys.[2] While some of the plotted "variation" in party sizes and

[1] These parties are the Liberals (VVD), the Christian Democrats (CDA), and the Labour Party (PvdA).

[2] The Comparative Manifestos Project and the Eurobarometer survey series, the sources of these data, are discussed at some length in chapter 11 below. Several smaller parties represented in the Dutch legislature have been omitted in the interest of clarity.

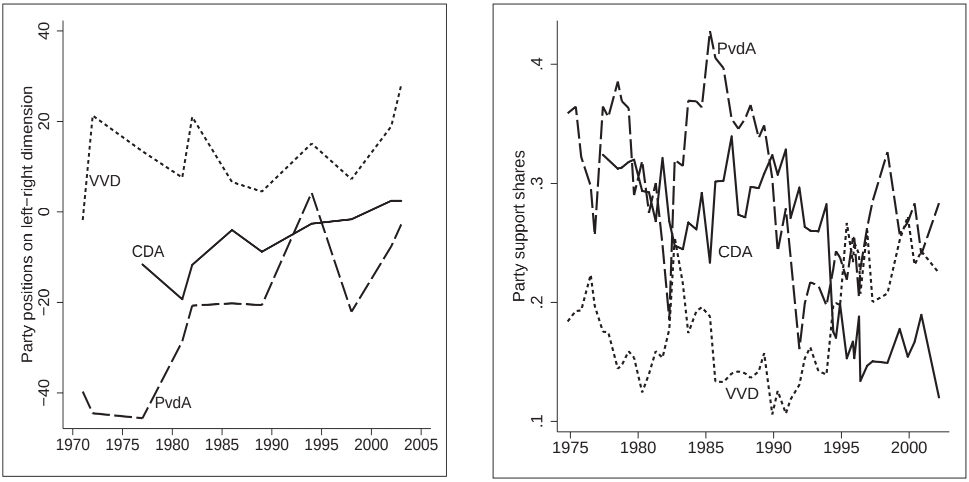


Figure 1.1. Dynamic party competition in the Netherlands, 1970–2005.

policy positions is surely the result of measurement error, by no stretch of the imagination was the Dutch party system "flatlining" in steady state during the period under observation. It was clearly a dynamic system and, as a result, there were frequent changes in the partisan composition of Dutch governments. These dynamics are clearly a central concern for all political scientists analyzing Dutch politics during this period, be they theorist or country specialist. Equivalent plots can be generated for any party system in which we might be interested.

## We Need a New Approach to Modeling Party Competition

Formal models of party competition have been an abiding preoccupation of political scientists since the early 1960s. A vast body of existing work has added hugely to our understanding of party competition. Our own *substantive* interest, however, and we believe the substantive interest of most people who want to understand party competition in democratic societies, concerns crucial features of party competition that these models typically assume away as a price to be paid for analytical tractability. We ourselves are interested in party competition among many more than two parties. We are interested in "multidimensional" political environments in which politicians and voters care about more than one type of issue. We see politics as a continuously evolving dynamic process that never settles at some static equilibrium, to be perturbed only by random shocks. Pursuing these interests poses formidable theoretical challenges.

We show in chapter 2 that dynamic models of multiparty competition, especially when voters care about a diverse set of issues, are analytically intractable. They are not just "difficult" to solve, they *cannot* be solved using conventional analytical techniques.

The *analytical* intractability of the relevant theoretical models does not make us any less interested, *substantively*, in dynamic multiparty competition. Indeed, this very intractability gives us an important and liberating theoretical insight. If analysts cannot use tractable formal models to find optimal courses of action in this setting, then *neither can real people making real decisions about real party competition*. These people still need to make decisions about what to do. If no formally provable best-response strategy is available, real humans must employ informal decision rules or heuristics.[3] To preview a decision rule we investigate extensively in this book, a party leader might decide to move party policy toward the position currently advocated by some larger rival party, on the grounds there must be more voters who prefer this rival's policy position. We find that this decision rule (which we call Predator) is sometimes very, very good and sometimes perfectly horrid. It is certainly not a "best" response in any conceivable situation but, in the analytically intractable setting of dynamic multiparty competition, it is one of many potentially good rules that politicians may use in certain circumstances when they set party policy positions.

## Agent-Based Modeling

Analytical intractability of the decision-making environment, and the resulting need for real politicians to rely on informal decision rules, suggests strongly that we use *agent-based modeling* to study multiparty competition in an evolving dynamic party system. Agent-based models (ABMs) are "bottom-up" models that typically assume settings with a fairly large number of autonomous decision-making agents. Each agent uses some well-specified decision rule to choose actions, and there may well be considerable diversity in the decision rules used by different agents. Given the analytical intractability of the decision-making environment, the decision rules that are specified and investigated in ABMs are typically based on adaptive learning, rather than forward-looking strategic analysis, and agents are assumed to have bounded rather than perfect rationality (Gigerenzer and Selten 2001; Rubinstein 1998; Simon 1957). ABM is a modeling technology that is ideally suited to investigate outcomes that may emerge when large numbers of boundedly rational agents, us-

[3] We use these terms interchangeably in what follows.

ing adaptive decision rules selected from a diverse portfolio of possibilities, interact with each other continuously in an evolving dynamic setting (MacGregor et al. 2006).

Putting a particular ABM to work by manipulating its parameters and observing the associated outcomes typically involves *computing* the outcomes of these interactions if the underlying model is analytically intractable—as is usually the case. Such computation, does not, of its essence, involve *electronic* computers. One of the most influential early ABMs analyzed housing segregation by scattering black and white chips and then moving them around on what amounted to a large chess board (Schelling 1978). This model was computational in the sense that an abacus is a computer, implemented by moving pieces around a chessboard. As originally published, it did not rely on using an *electronic* computer.[4] Scatter a number of black and white chips at random on a chessboard; these chips represent people of different color. Assume people have some view about the color of their neighbors; say, for example, they are unhappy if fewer than a quarter of their neighbors are the same color as them. The modeled behavior is simply that unhappy agents move to a randomly chosen close-by empty square that makes them happy. A model "run" begins with chips scattered at random. With an equal number of black and white chips, the typical person will find that 50 percent of neighbors are the same color and will be happy to stay put. There will however be some people in the random scatter who find that fewer than a quarter of their neighbors are the same color; they will move to a square that makes them happy. Everyone is given a chance to move, and to move again, using this rule until there is no unhappy agent who wants to move. The results are striking and unexpected. Even if everyone merely wants at least a quarter of their neighbors to be the same color, modeled population movement typically results in a steady state in which on average about 60 percent of a typical agent's neighbors are of the same color. If we change the key model parameter and assume people to be unhappy, and to move, when they are in a local minority (fewer than 50 percent of neighbors are the same color) then people find that on average 88 percent of neighbors are the same color in the typical steady state that emerges. The deep substantive insight from Schelling's ABM is that intense spatial segregation can arise when people do not seek this at all, but simply prefer not to be in a small minority. More generally, this model shows very nicely that simple decision heuristics can interact to generate complex and unexpected "emergent" patterns of social behavior. This is the core insight of agent-based modeling.

[4] A version of this model implemented in NetLogo for electronic computers can however be found in the NetLogo models library.

All good things come at a price. The price paid for using computational as opposed to formal analytical models, and thus for using agent-based modeling, is that computation involves calculating model outputs for particular parameter settings. An analytical result, if it is general, is a beautiful thing that is good for all valid parameter settings. Strictly speaking, computational results are good only for those parameter settings that have actually been investigated. Inferences about parameter settings that have not been investigated—and thus more general theoretical inferences we might want to draw from the model—are, in effect, interpolations. This is one reason why we never use computational methods when analytical results are available for the substantive problem that interests us.

The distinction between analytical and computational methods should not be overdrawn, however. A longstanding set of observations that compare models of computation with systems of formal logic, collectively known as the "Curry-Howard isomorphism," shows us that computer programs and formal proofs are in essence the same thing (De Groote 1995). Both take a set of explicit premises and manipulate these, using some system of formal logic, to prove theorems based on these premises. Consider, for example, the area, $A$, of a circle with radius $r$. It is well known that we can prove analytically the proposition: $A = \pi \cdot r^2$ for *any* positive real $r$. We can also prove $A \approx \pi \cdot r^2$ for *any given* positive real $r$ by various computational methods. With *infinite* computing power at our disposal, we could prove $A \approx \pi \cdot r^2$ for *any* positive real $r$.[5] This would not be an "elegant" proof according to most standards of elegance, but now we are talking about aesthetics. With less-than-infinite computing power, we can sample a huge number of positive real values of $r$, compute $A$, and show in every single case that $A \approx \pi \cdot r^2$. We can draw the *statistical* inference, at a specified level of confidence, that $A \approx \pi \cdot r^2$ for any positive real $r$. If for some reason it happened that we could not prove analytically that $A = \pi \cdot r^2$, then this computational/statistical inference would be immensely valuable to us. If we wanted to increase our confidence in this inference, we could simply do more computing and sample more values of $r$. Of course, we could never be *perfectly* confident in this conclusion. We can show that $A \approx \pi \cdot r^2$ when $r = 2.0000001$ and $2.0000002$; you could claim it is possible $A \neq \pi \cdot r^2$ when $r$ is set between these values, at

[5] The approximation arises because $\pi$ is a transcendental number that cannot be stored to any arbitrary level of precision in a digital computer. However, a number very close to $\pi$ can be stored as a high-precision floating point number. For the same reason, the area of any circle calculated using the classical formula $A = \pi \cdot r^2$ can be *written down* as a real number only using an approximation that deploys some arbitrary level of precision specifying the number of decimal places we are prepared to use.

2.00000015. Strictly speaking, this would be true.[6] We could however show statistically, with access to enough computing power, that the *probability* of this exception is extraordinarily small. Furthermore, we could drive down this probability to as low a level as makes you feel happy—simply by doing more computing.

This is an issue we take very seriously indeed in this book since we do want our computational results to have effectively the same scope and precision as those derived from analogous analytical work. We address this by specifying careful procedures for systematically varying parameter settings, and rigorous methods for estimating model outputs of interest associated with these settings. If we carefully design and execute our computational work in this way, then the scope and precision of our results depend only on the volume of computation we are willing and/or able to deploy. Since we want our own results to have the same scope and precision as typical results from formal models in this field, we are both willing and able to deploy a huge amount of computing power, taking advantage of the Harvard-MIT Data Center's high-performance cluster in order to do this. An important consequence of this is that we are confident that the computational results we present in this book can be "taken to the bank," in the formal statistical sense that, if we were to do very much more computing, or if many other people were to repeat our procedures, essentially identical results would arise. Thus, while this is a book above all about the substantively fascinating topic of multiparty competition, it is also an exercise in how to use computational methods in general, and ABMs in particular, in a way that allows us to draw confident general conclusions.

To summarize, the substantively important real-world problem that interests us is the dynamics of multiparty competition. Theoretical models are no more than intellectual tools designed to help us understand substantively important real-world problems. The technology of classical formal modeling is not a good tool to help us understand the dynamics of multiparty competition, since the resulting models are analytically intractable, with consequences for analysts and more importantly for real humans making decisions in these settings. In contrast, the empowering new technology of agent-based modeling is well suited to investigating problems that are of great substantive interest to us. Impatient for results and problem focused as we are, this book is about how agent-based modeling helps us think systematically about the dynamics of multiparty competition. We start simple and build an increasingly complex model of party competition that deals with a range of substantive matters we have

[6] Although our advice to you in this case would be that you should get out more.

wanted to think about for a long time but had not really been able to think about in a systematic way before the emergence of ABM.

## Plan of Campaign

Chapter 2 sets up the core problem in which we are interested. To demonstrate that this problem is analytically intractable, we use compelling results from a subfield of geometry that deals with "Voronoi tessellations" (or tilings) and has powerful applications in many disciplines. Largely unnoticed by political scientists, this work addresses a problem of "competitive spatial location" that is directly analogous to the problem of dynamic competition between a set of political parties competing with each other by offering rival policy programs. One result from this field is that the problem of competitive spatial location is intractable if the space concerned has more than one dimension (we return below to discuss the meaning of a "dimension" in models of party competition), implying that there are no formally provable best-response strategies for this. This is an important and widely recognized justification for deploying computational methods, and the study of Voronoi tessellations is a major subfield in *computational* geometry.

Chapter 3 specifies our "baseline" ABM of dynamic multiparty competition, which derives from an article published by one of us (Laver 2005). This assumes that each voter has in mind some personal ideal "package" of policy positions and supports the political party that offers the policy package closest to this. The dynamic system at the heart of our model is as follows: voters support their "closest" party in this sense; party leaders adapt the policy packages they offer in light of the revealed pattern of voter support; voters reconsider which party they support in light of the revealed pattern of party policy packages; and this process continues forever. This recursive model describes policy-based party competition as a complex system, and our baseline model specifies three decision rules that party leaders may deploy when they choose party policy positions in such a setting. These rules are Sticker (always keep the same position), Aggregator (move policy to the centroid of the ideal policy positions of your current supporters), and Hunter (if your last policy move increased your support, make another move in the same direction; or else change heading and move in a different direction). These rules model, in a simple way, an "ideologically intransigent" party leader who *never* changes party policy, no matter how unpopular this might be; a "democratic" party leader who always adapts the party position to the preferences of *current* supports; and a "vote-seeking" party leader who is always looking for *new* supporters and does not care what policies must be chosen in order to do this. These decision rules were specified in Laver (2005); the innovation in this chapter concerns our assumptions about the preferences of voters. Rather than assuming a single coherent voting population with a perfectly symmetrical multivariate normal distribution of ideal policy positions, we now assume that electorates comprise a number of distinct subgroups. Combining subgroups into an aggregate voting population, we produce an aggregate distribution of ideal points that is no longer perfectly symmetric. This more generic assumption about voter preferences makes a big difference to what our model predicts.

Chapter 4 develops our methods for designing, executing, and analyzing large suites of computer simulations that generate stable and replicable results. We start with a discussion of the different methods of experimental design, such as grid sweeping and Monte Carlo parameterization. Next, we demonstrate how to calculate mean estimates of output variables of interest. In order to do so, we must first discuss, among other things, stochastic processes, Markov Chain representations, and model burn-in. As we see below, we are especially interested in three stochastic process representations: nonergodic deterministic processes that converge on a single state, nondeterministic stochastic processes for which a time average provides a representative estimate of the output variables, and nondeterministic stochastic processes for which a time average does not provide a representative estimate of the output variables. The estimation strategy we employ depends on which stochastic process the simulation follows. Last, we present a set of diagnostic checks, used to establish an appropriate sample size for the estimation of the means. More observations obviously lead to more precise estimates. However, given a fixed computational budget, in terms of computer processing time and storage space, as well as the opportunity costs of not executing other simulations, we want to gather enough observations to allow precise estimates, but no more than is needed.

We report our benchmark results in chapter 5. Perhaps the most striking of these concerns the "representativeness" of any given configuration of party policy positions and uses a second result that comes from the Voronoi geometry of competitive spatial location. A set of *n* points arranged so as to generate a "centroidal Voronoi tessellation" (CVT) is an "optimal representation" of the space in which these points are located. By this we mean that the *aggregate* distance between all points in the space and their closest generating point can never be less than when the *n* generating points are arranged in a CVT (Du et al. 1999).[7] If we think that voters are

[7] The analogous problem in digital imaging is to find the most representative set of *n* points (party positions) to represent a much more detailed image comprising *m* points (voters). More generally, a CVT can be seen as a "best" simple representation of any spatially structured dataset.

more satisfied at election time the closer their own ideal policy is to the policy position of their closest party, then this implies that the electorate as a whole is most satisfied when party policy positions are arranged in a CVT. Since the "representativeness" of any party system is an important matter, both normatively and in terms of practical politics, the notion of an optimal representation gives us an important benchmark for assessing evolved configurations of party policy positions. A robust conjecture in computational geometry, concerning what is known as Lloyd's Algorithm (Lloyd 1982), is very relevant in this context. If all party leaders use the Aggregator rule for setting party policy positions, continuously adapting party policy to the centroid of *current* supporters' ideal points, then Lloyd's Algorithm tells us that the set of party policy positions will converge on a steady state that is a CVT. *Party positions in all-Aggregator party systems thus evolve to configurations that are optimal representations of the space.*[8] Other configurations of party policy positions will generically imply suboptimal representation, in this precise sense.

Thus far we have treated the set of competing political parties as exogenously given to us by God or Nature. We move beyond this in chapter 6 and define a model of endogenous party "birth" and "death" (Laver et al. 2011; Laver and Schilperoord 2007) that has the implication that *the set of surviving political parties is endogenous* to the system of party competition. We now also model competition between party leaders using different decision rules, extending work on this using computer "tournaments" (Fowler and Laver 2008). All of this requires us to extend our model to define a de facto survival threshold for political parties; an updating regime that specifies how voters feel about the party system today, given what happened today and how they felt about the system yesterday; and a distinction between "campaign ticks" of the model, during which party leaders make choices that do not have long-term consequences for their survival, and "election ticks" that do have a bearing on party survival. The resulting more realistic model of party competition with endogenous parties is *evolutionary*, describing a *survival-of-the-fittest* environment in which more successful parties survive and less successful parties do not.

Up to this stage in the argument, we have extended, improved, and generalized previously published work based on three simple decision rules: Sticker, Aggregator, and Hunter. We break completely new ground in chapter 7, defining new "species" of vote-seeking decision rule (Predator and Explorer) and specifying both these and existing rule species in terms of a set of parameterized rule "features," including satisficing and speed of adaptation. Predator rules, specified in a flawed form in Laver (2005)

[8] In this context it is very important to note that there is typically no unique optimal representation.

and redefined by us here, in essence attack the closest more successful party by moving their policy position toward it. Explorer rules are generalizations of "hill climbing" algorithms. Explorers randomly poll positions in some local policy neighborhood during campaign ticks—moving on an election tick to the best position they found during the campaign. The net result of these extensions is that we now consider competition between party leaders who may choose from one of 111 different decision rules—or, strictly speaking, parameterizations of rule-agent pairings. This dramatically expands the state space of our model and forces a major modification in the method we use to estimate characteristic model outputs. We find that which particular vote-seeking rule is most effective depends critically on parameters of the competitive environment. Chapter 7 reports another result we feel is particularly important, concerning what happens when *satiable* and *insatiable* vote-seeking party leaders compete with each other. We find well-defined circumstances in which satiable leaders, who do nothing until their party vote share falls below some "comfort threshold," systematically win higher vote shares than insatiable leaders, who always seek more votes no matter how many they currently have. This is a good example of the classic "exploitation-exploration trade-off" in reinforcement learning (Sutton and Barto 1998). Insatiable party leaders always explore the space in search of more votes, whereas satiable leaders exploit their good fortune whenever vote share is above their comfort threshold. This is the type of insight that can be derived only from a *dynamic* model of party competition.

In chapter 8, we extend our survival-of-the-fittest evolutionary environment to take account of the possibility that new political parties, when they first come into existence, do not pick decision rules at random but instead choose rules that have a track record of past success. We do this by adding *replicator-mutator dynamics* to our model, according to which the probability that each rule is selected by a new party is an evolving but noisy function of that rule's past performance. Estimating characteristic outputs when this type of positive feedback enters our dynamic model creates new methodological challenges. Having addressed these challenges, the simulation results we report in chapter 8 show that it is very rare for one decision rule to drive out all others over the long run. While the diversity of decision rules used by party leaders is drastically reduced with such positive feedback in the party system, and while some particular decision rule is typically prominent over a certain period of time, party systems in which party leaders use different decision rules are sustained over substantial periods. More generally, we continue to find party leaders choosing from a diverse rule set in this evolutionary setting. We find no evidence whatsoever of evolution toward the dominance of a single decision rule for setting party policy positions.

Moving beyond the assumption that voters care about only the party policy positions on offer, chapter 9 models the possibility that they also care about perceived "nonpolicy" attributes of political candidates: competence, charisma, honesty, and many other things besides. These characterize what have become known as "valence" models of party competition. Voters balance utility derived from each candidate's nonpolicy valence against utility derived from the candidate's policy position. The contribution of valence models has been to explain why all parties do not all converge on regions of the policy space with the highest densities of voter ideal points. Higher valence parties tend to go to regions of the policy space with higher voter densities, while lower valence parties are forced to steer well clear of these parties and pick policy positions in regions with lower voter densities. We replicate and extend the findings of traditional static valence models, with one important twist. Over the long run, lower valence parties tend to die and higher valence parties tend to survive, a finding that suggests a reappraisal of valence models as currently specified. These essentially static models show a snapshot of the party system at a given time; but the tendency of low-valence parties to disappear in an evolutionary setting suggests that these snapshots are not dynamic equilibriums that can be sustained over time.

Moving beyond voters who care about more than policy, we look in chapter 10 at party leaders who care about their own private policy preferences as well as about winning votes. In the spirit of our existing model of endogenous party birth, we take the preferred policy position of a party leader as the founding policy position of his or her party. In an intriguing echo of our findings on satisficing in an evolutionary setting, we find that party leaders who care somewhat about their own policy position may do somewhat better *at winning votes* in competition with party leaders who care exclusively about vote share. This may arise from the fact that, in an evolutionary setting, *the ideal points of surviving party leaders are endogenous*. Each surviving party leader was once a new entrant into the system at a policy position for which there was demonstrable voter "demand." Leaders who stay close to this founding position continue to satisfy the demand that originally caused the party birth. They thereby also, effectively though not intentionally, forestall new party births in this region of the policy space.

Having specified theoretical models of multiparty competition in the first ten chapters of the book, we investigate empirical implications of these models in chapter 11, comparing model predictions with changes in observed party policy positions and vote shares in ten real European party systems. Confronting theoretical models with empirical data is a central part of the definition of political science as a "science," highlighted by the influential Empirical Implication of Theoretical Models project. This is easy to say but hard to do well, and it is even harder for dynamic models that have many parameters whose values are not directly observable in the real world. We face serious problems of model *calibration*, of finding parameter settings for our theoretical models that plausibly correspond to those in the real party systems for which we can observe empirical observations. Calibration problems are compounded in this case by an acute shortage of high-quality time series data on party system outputs of interest. Given all of these problems, the best we can hope for is to find "plausible" model calibrations generating predictions that are close to reality. This is of course not a full-dress scientific test of our model, which is not feasible given the lack of good time series data on party policy positions, combined with the lack of reliable independent data for model calibration. We prefer, however, to be honest about the calibration and data problems that arise with any dynamic model of multiparty competition rather than, dishonestly we feel, making "assumptions" about model calibration that will give us lovely empirical results but that are, in effect, assumptions chosen to give those lovely results. What we call the "autocalibration" of our model, searching for model calibrations consistent with accurate predictions, does allow us to conclude that the model *can* be calibrated to generate accurate predictions and that the calibration values associated with good predictions do have good face validity.

Putting all of this together, our fundamental interest in this book is in multiparty competition, seen as an evolving dynamic system. Our fundamental intellectual objective is to explore some of the puzzles about this that can be addressed using techniques of agent-based modeling. Substantively, while readers are the ultimate judges of this, we do feel that agent-based modeling empowers us to tackle interesting and important questions that cannot be addressed so fruitfully using the techniques of classical analytical modeling. Methodologically, we do feel that carefully designed and executed computational work can generate results that have a scope and precision equivalent to those generated by more traditional techniques.

Our sincere hope is that we open an intellectual window for at least some readers, who will take the ideas and suggestions in what follows and improve them beyond all recognition.