Computer Simulations of Elections, With Applications to Understanding Electoral System Reform

by

Samuel N. Baltz

A dissertation submitted in partial fulfillment
of the requirements for the degree of
Doctor of Philosophy
(Political Science and Scientific Computing)
in the University of Michigan
2022

Doctoral Committee:

Professor Walter Mebane, Chair
Professor Allen Hicken
Associate Professor Iain Osgood
Professor Scott Page

Samuel N. Baltz

sbaltz@umich.edu

ORCID iD: 0000-0002-4669-7435

ACKNOWLEDGMENTS

It would not be possible to cover all the ways that Walter Mebane shaped this dissertation in his role as the chair of my dissertation committee. Computer models of elections are a surprisingly old idea, first implemented around 1960 or even earlier, but the direct ancestors of the work in this dissertation are simulations of elections that Walter designed in the context of election fraud detection. Walter directed energy, money, and a great deal of time towards developing computer models that use game theoretic ideas about voter behaviour to generate simulated election results, and from there it was a short road for me to build those models out in different directions and apply them to actual elections. Walter's influence is visible in every step of that work. All the parts of this dissertation are better not just because of his attentive suggestions and meticulous feedback, but because of the example he sets as a researcher. For all of that, I am very grateful.

The research that follows, and my approach to it, was also heavily shaped by a series of projects I was lucky enough to undertake with Allen Hicken. From Allen I learned a great deal about good research — both abstract lessons about political institutions and the forces that shape them, and concrete lessons like the sheer amount of toil that sometimes goes into filling just one cell in a very big dataset. But I learned just as much about good mentoring. Allen modeled the lesson that, even when research is terribly difficult and exacting, there is no reason to approach it with anything less than generosity and good humour. I will pay this lesson forward whenever I have the chance.

To Iain Osgood I am grateful for a tremendous volume of sage professional and life advice over the years, for deep and insightful comments about every model I've ever shown him, and for helping me present my work in a way that other humans can hopefully understand. I had the privilege of being Iain's teaching assistant for two semesters; the only other person whose classes I assisted with twice is Scott Page. I am very grateful to both of them for demonstrating how to be an engaging, compassionate, professional, and highly effective teacher.

And to Scott Page I am also grateful for, time and time again, asking the sort of modeling questions that make me have to go into a room, shut the door, and stand very still in front of a white board for a few days ... and for somehow always knowing how to take what comes out of that process and connect it in the most direct way to a claim about the social world.

Thanks are also due well beyond my committee. I have tried to keep them brief, and sort them by time and theme. Lists of people are alphabetized by last name.

Much of this project was completed while I was on a Rackham Predoctoral Fellowship. The simulations in this dissertation were run on computers that I was able to purchase thanks to a fellowship from the Michigan Institute for Computational Discovery and Engineering. Some of the work that this dissertation builds on — notably the setup that first exposed the problems that I study in Chapter II and the original derivation of Equation 1 — was done with Walter Mebane and Fabricio Vasselai under NSF award SES 1523355. Some specific ideas, identified by footnotes, were included in this dissertation thanks to feedback on past work from Lucia Motolinia, Lori Thorlakson, and George Tsebelis. I am grateful for very helpful feedback on earlier versions of various parts of the dissertation from Rob Franzese, Mai Hassan, Michael Lerner, and Jim Morrow, and I thank Lawrence LeDuc for sharing valuable insights about electoral system reform in Canada. Without Corina Simonelli my brain would have turned into pudding before I figured out how to make the maps in this dissertation. Mike Thompson-Brusstar saw more potential than I did in one of my most bizarre research ideas, which pushed the work in a much more interesting direction.

Would I ever have started doing this sort of work without Fabricio Vasselai? How far would I have gotten? I'm very glad I don't know the answers to those questions.

The path to a dissertation like this one could only have led through the Complex Systems Center, and I thank Mita Gibson, Patrick Grim, Linda Wood, and the CSAAW community for making that such an interesting place to loiter. Charles Doering, after a few minutes with a blackboard, knew more about the continuous-time version of the problems in Chapter II than I know even now about the discrete-time version. His insights are badly missed. I also learned many important lessons about simulating scientific phenomena from fiddling around with Modules for Experiments in

Stellar Astrophysics at the encouragement of Fred Adams, and the models that follow are better because he had the patience to discuss stellar evolution models with a dilettante.

Years ago Kenichi Ariga let me help gather electoral system data and Neil Nevitte let me help research vote choice. I'm very grateful for the lessons they taught me (though when I glance at this dissertation while reflecting on the fact that they had me working on electoral systems and voter behaviour all those years ago, it does make me worry about the limits of my own imagination). And nowhere have I learned more than at Spacesim with Jim Magwood.

I am also, of course, very grateful to my family. Among many other things, I have to thank my parents, Jay Baltz and Rachel Eugster, for demonstrating that people can be deeply affected by even (especially?) the really boring-sounding parts of politics ... like "zoning bylaws", say, or "electoral systems".

Finally, this dissertation is the most involved manifestation yet of an itch I've had for a long time: that humans and our political systems are simply terrible at really long-term planning. The source of that itch deserves recognition too: I got it by reading *Cosmos* by Carl Sagan. It was while reading *Cosmos* that I — a nerd who liked science, reading as I walked home from my chemistry summer classes — realized that the problems I wanted to do something about were not actually engineering or even natural science problems, but rather political problems. Of course that was step #1 on the path to this dissertation. If you haven't read any Carl Sagan, then I encourage you to put this dissertation aside for a moment and go read some — say, for example, chapter 22 of *Pale Blue Dot*. The work that follows will make more sense in that context.

PREFACE

How many people have complained to me over the years that votes are counted in a way that *just does not make any sense*?

I came of age between the unification of the Canadian right and the electoral system hubbub of 2015, while Americans too were waking up to the sheer clownishness of some of their electoral institutions, and I have heard the complaints too many times to count. In a certain cross-section of North American geekdom, everyone has a pet electoral system that they know might not fix everything but at least would make things better — *right*? And it's easy to understand why people feel so strongly: electoral systems in themselves might not be the most important issue governments face, but they are the ur-problem, the one that affects whether and how the country can address any problem in the future. I do have to agree that the basic complaint makes a ton of sense: wouldn't it be shocking if the very simplest ways to run elections, the ones people came up with so many hundreds of years ago, were also the best possible ways? But no matter how many times I have this conversation, even as I have slowly made the study of electoral systems my day job, I never seem to persuade anyone that once you get past truly minimalistic counting rules, we have very little idea of what *exactly* a certain electoral system might do. I have an almost perfect failure rate, across so many conversations that I've lost count, in persuading people that as much it might *really seem like* some electoral system would have such-and-such effects, even the most obvious-seeming consequences may or may not come true when we plunk an abstract system down into the bubbling goo of real politics. These decisions are massive, and we cannot afford to be vague.

So for the many people who have been understandably unconvinced by my handwavey arguments over the years about what exactly these systems might do, concretely, in the language of a specific country's politics at a particular moment in time — what follows is finally some specificity. I would be honoured if you would go ahead and judge it for yourself.

TABLE OF CONTENTS

CHAPTER

LIST OF FIGURES

FIGURE

LIST OF TABLES

TABLE

LIST OF APPENDICES

ABSTRACT

This dissertation addresses problems in the computer modeling of elections, and then uses these models to simulate electoral system reforms. Computer simulations of how people vote often do not produce just one fixed election result, and instead get stuck in unending cycles. To motivate why these models are nevertheless worth studying, I introduce an extremely simple voting heuristic that correctly matches the winner of 95% of the districts in Canada's 2019 federal election. I show that the vote cycles in a broad variety election models obey simple constraints, and that, for some substantively important rulesets, where the model is stopped within these cycles does not strongly affect the aggregate results. I then apply these models to the question of election reforms. Politicians agree to change election rules when they think a change would benefit them, but there is no way to know how a specific party would perform under a specific alternative system. Combining game theory, past election results, and public opinion data, I estimate those effects falsifiably and accurately. The model suggests that an election reform attempt in Canada did not need to fail, and would have helped a supermajority of parliament. In Britain, it instead shows smaller effects from election reform than expected. The model also supports the idea that a convoluted new electoral system in Alaska was designed to protect a specific incumbent senator, but was probably not the best electoral system for the job. Election reform efforts frequently become mired in uncertainty and fail, and this method offers a way to transparently set expectations for different systems' effects.

CHAPTER I

Introduction

This dissertation investigates how to use the methods of scientific modeling to simulate large democratic elections. The type of computer models considered here have three characteristics: they represent just one specific election, they explicitly simulate a realistically large number of individual electors, and those electors periodically consult how much support each party has and update their intended vote accordingly.[1] The output of these models is a simulated vote count, intended to be as close as possible to the real vote count that might occur under the conditions being modeled. Even though the use of computers to simulate elections — and even to simulate individual voters — goes back to the middle of the last century, this framework has appeared in only a small number of publications. Two obstacles have prevented it from being used more widely.

The first problem is that some very basic methodological questions have never been thoroughly addressed. The most important arises from a gap between the method of study and the object being studied. There is no guarantee that the outcome of an iterative computer model will converge to one fixed result. It might instead change across iterations, perhaps forever.

If a question is very concrete, then there may be a straightforward way to interpret these changes in values. For example, in an iterative model of light traveling from the sun to the Earth, we might break the model down into individual seconds. Each second is one iteration of the model: we simulate a fixed number of photons, and every iteration we increment each photon from the sun to the Earth by the distance that light travels in one second. In every iteration, then, we know the flux of the sun's light on the Earth's surface — the meaning of this outcome is clear, and the fact that

---

[1] By "electors", I mean anyone who has a right to vote.

it changes between iterations is not a problem. Alternatively, if the question is fairly simple, then the problem of varying outcomes might just never arise. In a model of skilled players facing off in a game of Tic-Tac-Toe, the players might play an optimal move every iteration until the model is done according to the rules of the game.

Elections represent a very different challenge. There is no natural increment of time that can be connected to the iterations in a computer model. But the iterative analogy is tempting: some sorts of information about elections do become available at discrete intervals. If a popular polling aggregator updates the average of recent polls at 10 am every day during an election cycle, then electors have new information about the state of the race every 24 hours. If electors update their vote intention as they observe that some parties are rising and others falling in the polls, then that process will itself appear in subsequent polls. But there is no reason that it should ever converge. In fact, scenarios with repeated voting famously tend to produce extremely complicated cycling behaviours, with one collection of votes leading to another collection of votes which leads to another, and so on. We can imagine many different ways that electors use the information that they encounter every time-step, and many interesting and plausible ways of responding to polls in iterative election simulations will turn out to have these chaotic cycles.

The first obstacle leads to the second. Because the process of electors arriving at a vote choice is so challenging to simulate, it is difficult to see how they can be concretely applied to politics. For political scientists interested in empirical claims about political phenomena, the bare minimum expectation of a model of an election needs to be that it produces a set of expected votes. When these votes are generated by a model that is designed to resemble a real election from the past, the simulation can be assessed by comparing its vote totals to the real election results.

Chapter II of this dissertation will address the first challenge, developing ideas for assigning just one vote total to an iterative election game that endlessly bounces between several different vote counts — a situation that I call "iterative equilibrium". I motivate these computer simulations with an empirical success. I demonstrate that an extremely simple model of strategic voting, which only requires electors to know what percentage of votes they expect each party to win, correctly predicts about 95% of the electoral district results in a real election. I also establish that, in an application to a real election, very different rules for picking a single vote total in each electoral district produce very similar results. The minimalistic voting heuristic used in this application is a type of iterative election game in which electors deterministically connect each expected vote total to one particular vote choice, rather than randomly picking a vote choice when they see a certain vote total. I show that in every game like this, the number of vote choices that will repeat forever is finite, and in rulesets that model realistic strategic voting, the number of repeating vote choices tends to be small. I also show that the complexity of the changes in the vote totals tightly constrains the complexity of electors' changes in their expected vote choices.

Chapters III and IV address the second obstacle: how to use models like this to answer big questions about political phenomena. I consider the problem of electoral system reform. Electoral systems — the sets of rules that translate votes in an election into seats in a legislature — are famous for changing very rarely. Many major democracies have spent centuries without ever seriously altering their electoral systems. The consensus about why these rules are so hard to change is that reform would require powerful groups with very different interests to agree to adopt a new system. Usually in a democracy, a given electoral system is perceived as better for some parties and worse for others; since all parties are competing over the same seats in legislature, it is rare for an electoral system to be perceived as beneficial for enough powerful groups that the system actually gets changed.

However, there is also a consensus that electoral systems are so complicated that it is extremely hard (if not impossible) to predict the effects of any alternative system other than the simplest and oldest ones. We have an extremely detailed understanding about how electoral systems tend to affect many different features of a country's politics, and from these broad tendencies we can try to infer roughly how a given system might help or hurt a party that plays a particular role in a specific party system. As easy as it is to find researchers engaged in this kind of approximate reasoning by analogy, it is just as easy to find electoral system scholars who assert that it is unavoidably foolhardy to guess how a *specific* complicated electoral system might affect a *specific* political party. The result is not just that election scholars are unable to explain or predict electoral system changes (or, more often, the lack thereof), but that political parties and their leaders also do not have the tools to thoroughly understand their own self-interest. If nobody knows how exactly a particular alternative electoral system would affect a specific political party, how should that party decide whether or not it prefers that system to the status quo? In fact, there are examples of political parties making mutually contradictory judgments about whether or not to support some alternative electoral system.

I argue that the types of computer models studied in Chapter III offer a way out of this bind. In Chapter III, I construct an iterative model of election campaigns that explicitly simulates how a realistic population of individual electors arrive at a vote choice. I compare these models to real past election results, and I show that the simulated seat totals are similar to the real election results. Then, by varying just the electoral system, the model produces an estimate of how each party would be affected by a switch to some counterfactual electoral system. In Chapter III, I extend existing models of pivotal voting to cover electors' strategic considerations in a type of Ranked Choice Voting election. At the close of Chapter III, the model is applied to the case of the real electoral system change in New Zealand in 1996, and the model results under the new electoral system closely resemble the results of that election.

In Chapter IV, I apply the model to the very real prospect of electoral system change in two long-standing Westminster democracies: Canada and Britain. In both cases I verify that the simulation produces similar seat totals to the results of their 2019 elections, and then I simulate what might have happened in that election if it had been held under either Proportional Representation or the Borda count style of Ranked Choice Voting. The results of the model in Canada suggest that a supermajority of parties had the opportunity to switch to at least one alternative electoral system that would have been mutually beneficial for all of them, but that in the absence of a tool to estimate their seat counts under different systems, they let the opportunity pass by. In Britain, the model shows much smaller effects, suggesting that the Conservative wave of 2019 really cannot be attributed to the electoral system and that other systems might not immediately change British politics dramatically.

Finally, I apply the model to a very different case: the state of Alaska, where a new and highly unusual type of hybrid electoral system has been implemented, but not yet used in a real election. The new system staples together a plurality vote and an Instant Runoff round. The effects of this kind of hybrid are extremely unclear, and the complexity of the system makes it difficult to answer even very basic questions about what kinds of effects it might have. Nevertheless, there is a strong conventional wisdom about the system. It is seen as so favourable to centrists that many political commentators suspect it was specifically designed to protect Lisa Murkowski, one of Alaska's representatives in the United States Senate, where she is famously one of the most centrist members and has a history of casting pivotal votes on major legislation. The model supports the claim that the new system will be better for Murkowski than the old system, but it also suggests that certain types of Ranked Choice Voting would have been better for Murkowski than this hybrid system will be.

These substantive results are more than illustrations of a tool. They deepen our understanding of why proposed election reforms so often fail. By making transparent predictions about how those reforms could affect the people who have the power to make them happen, they unearth incentives that the players themselves might misjudge. They do so falsifiably, and the predictions they make about history match historical events.

However, these claims about electoral systems are not the main feature of this dissertation. The primary objective of what follows is to help unlock computer simulations of elections for broader use in political science, and to show how they can be used to better understand political phenomena.

CHAPTER II

Iterative Election Games with Simultaneous Updates

"I do not think there is any permanent equilibrium in the political affairs of any nation. It is always a moving equilibrium."

— Pierre Elliott Trudeau, *House of Commons Debates*, 23 March 1981, vol. 8, p. 8519

Election campaigns are a learning process. Voters, and those who plan not to vote, receive a stream of news about how many people support each candidate for office, and strategic voters use this information to decide how to vote. But as people adjust their votes, the number of people who plan to vote for each candidate will change. If the mechanism of that change is polling data, then these changes are made possible at discrete intervals. When a new poll is released, or a polling aggregation website updates their models, every voter has access to new information about other voters' intentions.

In this chapter I study voting games that proceed iteration by iteration: voters publicly signal their intended vote choice, then they update their intended vote choice based on the intention that other electors signaled, then voters update again based on that new update, and so on. When voters can map a given expected vote total onto any one of several different vote choices, this describes a probabilistic voting game, which is a well-studied type of iterated voting. But if voters map any given expected vote total onto just one vote choice, the situation is quite different. To classify the possible outcomes of these sorts of games, I introduce the idea of an iterative equilibrium. This definition expands the idea of a fixed static equilibrium to include situations in which electors respond with a sequence of vote choices that never settles down to just one election result. I introduce a barebones strategic voting rule called threshold pivotality, which voters can compute with almost no attention to politics and almost no intellectual effort, yet which hews closely to the theory of how real people vote. I then show that reasonable applications of this very simple heuristic closely match election results with real polling data and real political preferences in the case of Canada's 2019 federal election. Finally, I study the broader question of all voting games where electors deterministically map each possible expected vote total onto just one vote choice. I show that all such voting games reach either a fixed equilibrium or an iterative equilibrium, as do many probabilistic voting games, and I show an upper bound on the length of the cycles at iterative equilibrium. I also show that the electors in these games behave extremely predictably when the behaviour of the aggregate system is known.

To illustrate the setting, consider the problem confronting a voter in a multiparty election — say, for example, a voter in Ontario during Canada's 2019 federal election. This voter is not highly invested in politics, but every week or so they see a new poll on the evening news, or they check the average of recent polls, and they get an up-to-date sense of who is ahead in their area. The way this voter uses that information is very simple. Ideally they want to vote for the Green Party of Canada, or failing that for the New Democratic Party, but they understand from the polling averages that the Green Party has a very low chance of winning and the New Democratic Party is only competitive in some places. So when they check the polling average, if the Liberal Party and the Conservative Party are the only parties near the front of the pack, then they figure they will compromise and vote for the Liberal Party. But if it looks like the New Democratic Party is doing well enough, then they plan to support them.

Checking the polling average in Ontario on the day the election is called, the elector learns that the Liberal Party has a sizable lead, and the only party close to them is the Conservative Party, which has more than double the provincewide support of the New Democratic Party or the Green Party.[2] Now the elector has a choice to make: is the New Democratic Party close enough to support,

---

[2] This was the real situation at the start of the 2019 race: see the values for September 10, 2019, at the CBC poll tracker, available at https://newsinteractives.cbc.ca/elections/poll-tracker-archive/canada-2019/.

or is that just a waste of a vote if the contest will probably play out just between the Liberals and the Conservatives? Luckily, this is just the start of an election campaign. If enough voters decide to support a third party, then the next poll will likely capture a rise in the strength of that party, and it will become competitive by election day. If instead most electors decide to play it safe from the start, pooling only on the Liberal Party and the Conservative Party, then every other party is locked out.[3]

Under the classic approach to modeling elections, this story is hard to capture. Usually, models of strategic situations like large democratic elections abstract away any idea of the time between the start of an election and its end. Since the 1950s, social scientists have attempted to model the value of each vote choice facing an elector, to generate a theory about why people vote the way they do (Black, 1958, Downs, 1957). The conventional setting is game theory, which has embedded in its deepest foundations that a solution to a strategic interaction is "a set of rules for each participant which tell him how to behave in every situation which may conceivably arise" (von Neumann and Morgenstern, 1953, p. 31). An individual who solves a game therefore "contemplates, in one comprehensive view, everything that lies before him", understanding "the range of alternative choices open to him, not only at the moment but over the whole panorama of the future" (Simon, 1983, p. 13). Under the logic of best responses, the relevant quantity is how likely a given vote choice is to actually change the result; if every elector is playing the best response vote strategy to every other elector, then how should all of them be expected to vote? In the context of a large democratic election with thousands or even millions of other electors, answering this question requires predicting the decisions of every other elector, and calculating the value of each vote choice under every combination of votes. That would truly be a feat for "prodigious mathematical prodigies" (Binmore, 1990, p. 61).

The situation was dramatically simplified when researchers took advantage of the fact that it

---

[3] Throughout this dissertation I will usually refer to voters choosing between parties, not candidates. There are two reasons for this wording. The first is theoretical: the outcome of interest will always be how many seats each party obtains in a legislature, so what is really at issue is how many votes go to each party in a given district. The second is practical: even in countries with large polling industries, there is often no available data about peoples' opinions regarding every relevant candidate in every electoral district, but there are high quality data about the opinions that people in every electoral district have about the major political parties.

does not matter *who* votes each way, only *how many* people settle on each vote choice. With a few extra assumptions, this means that relatively straightforward statistical distributions can be used to estimate the vote totals of each alternatives, and therefore the probability that an elector will be able to create or break a tie between any two candidates (Myerson, 1998, Palfrey and Rosenthal, 1985).

The power of this approach for studying elections is that the incentives that voters have to support different candidates can be quantified and compared. But the assumptions required are such famously ill-fitting gloves when applied to very large elections that they have given rise to one of the most famous paradoxes in political science: unless great care is taken to avoid nonsensical results, this approach famously results in almost no electors voting (Wuffle, 1984).

Clearly people do tend to coalesce around more popular parties. Duverger (1951, Book 2, Ch. 1, §1) attributed the impact of electoral systems on the viability of political parties to two separate types of pressure: one direct, through the mechanics of the system, and the other psychological, in which the system constrains which parties people think are viable. Measures of strategic voting, though they strongly differ in how exactly they define the idea (Aldrich et al., 2018), find at a minimum that a substantial proportion of the population votes strategically (Alvarez and Nagler, 2000, Kawai and Watanabe, 2013). Other research has shown that voters do perform strategic computations, up to a limit: people vote strategically so long as the strategic computations are not too demanding, but after a certain point they revert to either using simple heuristics or just voting sincerely (der Straeten et al., 2010).

There is a middle ground in modeling as well. An increasingly popular approach has been to model electors following logic which is not a precise or full solution under the traditional definition. Some authors choose heuristics which model real-world behaviours of interest, while others focus on adaptive behaviours to leverage the fact that elections happen regularly and can inform electors' vote choices as they observe repeated votes (Bendor et al., 2011, Laver and Sergenti, 2012). One strain of the adaptive voting idea goes further, actually merging pencil-and-paper formal theoretic models with dynamical systems ideas, to model how elections might evolve in real time (Acemoglu et al., 2009, Demichelis and Dhillon, 2010, Fey, 1997, Golder, 2003, Konishi and Ray, 2003, Linzer and Honaker, 2003, Montoro-Pons and Puchades-Navarro, 2001, Sieg and Schulz, 1995). Electors updating their vote intentions continuously in time is not a model of repeated election polls *per se*, so much as model of electors constantly receiving updates in many different ways about the state of politics (say, by seeing lawn signs, talking to their friends, or reading the news throughout the day). And it remains for the modeler to choose how exactly the electors are learning about politics continuously in time, how they form beliefs about other electors' intended vote choices, and also how they update those beliefs over time.

In this chapter I study a simpler but closely related framework, already in use in some applications, that makes extremely minimal assumptions about electors' knowledge and computational investment in the game. I show that a very simple heuristic can capture the learning process that occurs in the lead-up to an election, in a way that can closely reflect reality. This model of vote choice is a type of strategic voting that retains the basic logic of rational choice game theory by carrying out best responses in a series of discrete time steps.[4]

The model makes three substantive assumptions:

- ↪ Each elector has complete, intransitive preferences over the candidates
- ↪ Electors pay attention to polls
- ↪ Electors update their best response poll-by-poll

When our Ontario voter finds the Liberal Party and Conservative Party in the lead at the start of the election, they know that this is only one signal out of many signals that will come. If they, and others like them, decide to signal support for a different party, then that party will have more strength when the pollsters call people again and broadcast another update about the parties' levels of support. So the situation is iterative: electors need not extrapolate from pre-campaign polling to solve the riddle of how every other elector in their region will vote, when instead they can simply wait for the next poll to come out and update their intended vote choice then.

---

[4] In this chapter I use "best response" loosely, to mean the Myopic Best Response to the vote total just presented (Tal et al., 2015). This is not necessarily a global mutual best response.

So what is the end state of such a model? Electors need some rule for best-responding to a poll, which does not necessarily involve predicting every single other elector's best response and finding a mutual best response that will be optimal on election day. It is easy to imagine voting rules that will quickly settle down to simple fixed states (for example, all electors might pick the rule "I vote for the first party alphabetically"), and it is also easy to imagine rules that would never settle down to a fixed state (if everyone perversely decides, say, "I will always support the party that received the fewest votes in the most recent poll"). But if electors use some sort of wasted vote logic, in which they are likely to vote for candidates in proportion to how popular each candidate is, the result is less obvious. Will this type of voting rule always pull electors in to fixed election results, or can it unleash chaos and never produce simple vote choices?

If a scheduler can tell the electors to update in a specific order until no elector wishes to update, then under the right conditions the voters will reach a fixed mutual best response (Kavner and Xia, 2021, Lev and Rosenschein, 2012, Meir et al., 2010). There is a growing literature on this type of sequential iterative voting, but the conceit of a scheduler is most appropriate for a different category of voting problems. In an actual election, no scheduler instructs every citizen to answer polls in a certain order until everybody has best-responded to their heart's content. If updates about the electors' current vote intentions are published for all electors at the same time, giving them one synchronized signal about what impact their vote might have — the situation when a poll is conducted in a large democratic election — then a difficulty immediately becomes apparent: there is no obvious moment when the electors will all stop updating their beliefs. And what little formal work has focused on this setting has shown that under some voting rules, this issue is at least as complicated as it initially appears (Kloeckner, 2020).[5]

The problem can be understood with reference to one of the oldest counterintuitive findings in voting theory: Condorcet's Paradox. de Condorcet (1785) considers three electors ϕ_1, ϕ_2, ϕ_3 who rank three candidates c_1, c_2, c_3 as follows:

---

[5] Though the few studies of how people practically behave in such a situation suggest that the abstract difficulties do not necessarily prevent people from choosing simple behaviours and sticking with them when presented with games of this structure (Tal et al., 2015).

$$\begin{cases} \phi_1: & c_1 \succ c_2 \succ c_3 \\ \phi_2: & c_2 \succ c_3 \succ c_1 \\ \phi_3: & c_3 \succ c_1 \succ c_2 \end{cases}$$

de Condorcet (1785) imagines a sequence of pairwise contests between the three candidates: if c_1 faces off against c_2, then c_1 earns the votes of ϕ_1 and ϕ_3 and is victorious, but similarly c_2 would defeat c_3, and c_3 would defeat c_1. The famous result is that the group's expressed preference ranking of $c_1 \succ c_2 \succ c_3 \succ c_1$ is intransitive, even though each individual's preference ranking is transitive.

Another consequence is that an iterative cycle is present in the Condorcet preference structure. Imagine that each of the three electors is determined to vote for a candidate, but they want to know how best to spend their vote. Each elector adopts a very simple learning rule: if I can boost a candidate other than my least-preferred option into first place by switching my vote, I will do so.

The electors apprise the situation by counting the number of votes they expect each candidate to receive. A sensible starting point is for each elector to publicly declare their top choice (imagine responding to a poll when you have no information about the competitiveness of any candidates — you might as well just name the alternative that you most prefer). Then, with v_i the expected vote total of candidate i, the vector of expected votes is

$$\begin{bmatrix} v_1 & v_2 & v_3 \end{bmatrix} = \begin{bmatrix} 1 & 1 & 1 \end{bmatrix}$$

Seeing this total, ϕ_1 reasons that by switching their vote, they can at least boost their second-favourite option over their least-favourite option, so they switch their vote to c_2. Similarly, ϕ_2 switches to c_3 while ϕ_3 switches to c_1. And again each candidate is expected to receive one vote. These electors, using this learning rule, will never stop switching.[6]

---

[6] Of course, one assumption of this simple example is particularly unreasonable: why should electors prefer switching rather than retaining their vote choice? This was just a consequence of the illustrative decision rule, which was simple to a fault.

Substantively, this might actually be a desirable property: clearly voters do continue to evaluate political parties and groups throughout the whole duration of an election, and even after an election ends. And the idea that cycles appear in iterated sequences of votes is a very robust and wide-reaching finding, though historically researchers have focused on chaotic sequences of cycles that voters adopt in repeated but static equilibria, rather than iterative cycles that occur during a single iteration like one election or one campaign (McKelvey, 1976, Schofield, 1978).

Methodologically, such models are also rapidly becoming necessary. Computer models have been used to study elections since the year 1960 or even earlier (Lepore, 2020, ch. 6). In the 1990s, a series of papers used iterative models to understand the dynamics of party competition (Kollman et al., 1992, 1997, 1998). This topic of study has remained small, but computational election models have been published with increasing frequency throughout the last two decades as large-scale computing power has become more available (Kim et al., 2010, Quinn and Martin, 2002, Sadiraj et al., 2004, 2006). In the last few years, several groups of political scientists have either applied scientific modeling methods to party competition and voting games (Bendor et al., 2011, Laver and Sergenti, 2012, Mebane et al., 2019) or begun to study the methodology of how to recover and interpret the output of large-scale computational models of political processes (Bendor et al., 2011, Siegel, 2018).

But if this phenomenon is substantively reasonable and methodologically interesting, it is still not clear how to handle it. In this chapter I define a notion of iterative equilibrium, and I show that every model must reach either a fixed or iterative equilibrium so long as electors map all of the vote totals involved onto a specific vote choice with certainty. I illustrate the idea of iterative equilibria by formalizing the simple heuristic voting logic described above. This type of game has fairly tame end states, and using the example of Canada's 2019 federal election, I show that it might not matter how iterative equilibria are resolved in the case of a real election. No matter how these equilibria are resolved, the heuristic model produces results that are very similar to the real seat counts. I also show that, so long as voters deterministically map vote totals onto vote choices, the results must contain fairly simple patterns. In particular, in a model where we observe the vote totals bouncing around in a certain way, we can infer very strict limits on how much the electors can be bouncing around as well.

Iterative equilibria

In this section we will formalize the strategic interaction sketched out in the previous section, using the following model of single-winner elections. An electorate is composed of electors who hold a rational preference ordering over the set of candidates. In an iterative process, each elector aims to pick one best-response vote following classic voting logic:

$$u = p \cdot b - c$$

where u is the utility obtained from casting a certain vote, p is the probability of changing the result of the election with that vote (called the "pivotal probability"), b is the benefit of causing that change, and c is the sum of all other factors influencing the vote, including cost of doing so and the electors' sense of duty to vote (Riker and Ordeshook, 1968).

To seed the iterative system, every elector announces their sincere top preference. Every elector learns the distribution of sincere preferences in the electorate, in the form of a poll. Electors then use this poll to pick a best-response strategy, which is announced in the next iteration.

So the structure of a model is:

$\hookrightarrow$ Initial expected vote totals are announced (e.g. sincere preferences or past results)
$\hookrightarrow$ Electors consult a poll or census of the expected vote count
$\hookrightarrow$ Electors calculate their expected pivotal probabilities
$\hookrightarrow$ Electors calculate the expected utility of each pure strategy
$\hookrightarrow$ Electors signal the best response vote choice
$\hookrightarrow$ Repeat from the second step

I will call this an **iterative election game**. The outcome of an iterative election game is a sequence of vote totals $\mathbf{v}$, indexed by the iteration t that produces the model: so I denote the vote total after electors have updated in iteration t as $\mathbf{v}_t$. If the model is never stopped, this repetition of vote totals will be countably infinite. The initial expected vote totals, which may be drawn from electors' sincere preferences or from the results of a previous election, are signaled before any updates have been made, so we might denote this $\mathbf{v}_0$. When electors update in response to that expected vote total, together they produce the results of the first iteration, $\mathbf{v}_1$. Electors' simultaneous responses to $\mathbf{v}_1$ in iteration 2 will produce $\mathbf{v}_2$. And the same is true for every subsequent iteration.

The phrase "iterative election game" is meant to capture the idea that these are simultaneous updates of opinions that take place during the course of a single election, to distinguish this type of game from sequential voting setup with a scheduler in the literature on "iterative voting" (Meir et al., 2010) as well as from the iterated elections setups of some classic models with cyclic voting patterns (McKelvey, 1976). An example of an application of an iterative election game is Mebane et al. (2019). Kloeckner (2020) has studied this same setting, but in the context of approval voting, and also with a focus on the abstract social choice properties of the system rather than the practical aim of obtaining a single election result from a system which, as Kloeckner finds, is full of "bad cycles and chaos". Indeed results in those papers, and simple intuition, show that a conventional fixed equilibrium is not guaranteed in iterative election games. But if we wish to use them to model elections, they have to produce some expected vote choice. So when do iterative election games stop?

I propose a definition of the end-states of iterative election games, which I call an "iterative equilibrium":

Definition 1. *The vote totals* $\mathbf{v}$ *in an iterative election game reach* ***iterative equilibrium*** *in iteration* t *if there is some natural number* $k > 1$ *such that* $\mathbf{v}_{t'} = \mathbf{v}_{t'+k}$ *for any* $t' \geq t$.

I refer to vote totals at iterative equilibrium as being "k-cyclic" or "k-periodic".[7] Definition 1 characterizes any end-state of a model where, after a certain number of iterations, the vote totals begin to repeat in every multiple of some number. If for example t is 10 and k is 3, then the vote total in iteration 10 is the same as the vote total in iteration 13, which in turn is the same as the vote total in iteration 16, and so on. Similarly the vote total in iteration 11 is the same as the vote total in iteration 14, which matches the vote total in iteration 17, and so on. More directly, the model is at iterative equilibrium in iteration t if there is a number k so that, whenever the difference between two iterations beyond t is a multiple of k, the vote totals in those iterations must be equal: iterative equilibrium is the situation where there is some natural number k larger than 1 so that for any pair of naturals t', t'' satisfying $t \leq t' < t''$, if $(t'' - t') \equiv 0 \pmod{k}$ then $\mathbf{v}_{t''} = \mathbf{v}_{t'}$. Note that the restriction $k > 1$ separates iterative equilibria from the meaningfully different case of fixed equilibria, in which the model settles to one vote total for all iterations. I will show that broad classes of games do fit the iterative equilibrium notion, and that vote switches are guaranteed to follow certain patterns that are fairly orderly and match a reasonable intuition about what an equilibrium is.[8]

In the next section I illustrate the idea of iterative election games and the iterative equilibria that can arise by using a simple decision rule that I call "threshold pivotality". In threshold pivotality, if a poll shows that there are multiple candidates within a certain threshold of winning a seat, then every elector will vote for their most-preferred candidate among those options. Otherwise, they vote for their sincere top choice. Imagine picking up the newspaper a few weeks before an election and finding that your favourite candidate is far behind the other options, but your second-favourite candidate is only 5% away from beating your least-favourite candidate. Under threshold pivotality, you would compromise and vote for your second-favourite candidate. But if there is only one candidate who has any reasonable chance of winning the election, then you might as well just vote sincerely or just stay home.

Iterative election games with threshold pivotality are an example of what I will call "determinis-

---

[7] The fact that k is not unique is a wrinkle in terminology that I will iron out in a later section.

[8] Thanks to Chris Hawthorne for thoughtful advice about how to communicate these ideas.

tic election games": iterative election games in which all electors deterministically map a given expected vote total onto a unique vote choice. After illustrating threshold pivotality, I generalize the results, and show that similarly simple results apply to all deterministic voting games.

Illustrative model: Threshold pivotality without abstention

In this section we study the illustrative voting heuristic of threshold pivotality. In a given iteration, an elector surveys the state of the world by checking whether or not any two candidates were within a certain threshold τ of first-place. For example, τ might be 5% in a recent opinion poll. I will call that situation a "τ-tie", and I use S to denote the set of candidates whose expected vote total is within τ of the highest expected vote total. If the number of candidates in the τ-tie is two, so $|S| = 2$, then there are two candidates whose expected vote total is within τ of the highest expected vote total. If $|S| = 1$ then the only candidate within τ of first place is the candidate who is expected to attain first place, and we will say that there is no τ-tie. I will also index S with the iteration that S refers to, so the set of candidates involved in a τ-tie in iteration t will be written S_t.

If there is a τ-tie in a poll, then every elector takes the expected utility of voting for one of the candidates in the τ-tie to be proportional to the sincere utility of voting for that candidate, and they assign 0 expected utility to each of the candidates not involved in the τ-tie. If there is no τ-tie, then electors vote for their sincerely most-preferred candidate.[9] In this section I do not consider abstention, though I will discuss that extension later.

This is a type of strategic voting ruleset, and can be phrased using the classic calculus of voting. In the equation

$$u = p \cdot b - c$$

electors using threshold pivotality assign the following values. If there is more than one candidate

---

[9] I am grateful to George Tsebelis for suggesting this decision rule.

involved in a τ-tie, so $|S| > 1$, then the electors set the pivotal probabilities p of all the candidates in S to be equal, while assigning all other candidates a pivotal probability of 0. If instead $|S| = 1$, then electors assign every candidate a pivotal probability p of 0, except for their sincerely-most preferred candidate, to which they assign some pivotal probability ϵ such that $\epsilon > 0$. Finally, the value of c must be less than the greatest value of $p \cdot b$ for any candidate, so that the elector always votes. This setup will produce threshold pivotality voting behaviour in the classic calculus of voting framework.

The threshold pivotality ruleset aims to directly capture the psychological mechanism by which electoral systems constrain party systems, in which most people only vote for parties that are within a certain range of winning. This idea is many decades old. In one of the original discussions of how different forces in a political system can work to constrain the number of political parties, Duverger (1951, Book 2, Ch. 1, §1) described a process by which individual psychological pressures can whittle the number of parties from three down to two. All that Duverger assumes is that electors can differentiate between two competitive parties and a third, less competitive party, and he famously argued that in a majoritarian electoral system many electors will transfer their votes from the third party to whichever party they find less objectionable among the other two. The three assumptions laid out in the previous section (that electors can rank how much they like the parties, that they know which parties are ahead and which are behind, and that there is a process over time by which they can shift their support from one party to another) were chosen to match this extremely simple voting logic, without asking any more of electors than is absolutely necessary. Threshold pivotality is a minimalistic way for electors to sort parties in the categories of “competitive” and “not competitive”, and then transfer their votes from non-competitive parties to competitive parties; it is a direct operationalization of Duverger’s original sketch of the psychological influence of wasted votes.

I am not aware of this exact ruleset having been formally explored and applied to the empirical case of a real election. Threshold pivotality does have some resemblance to the Laslier Leader Rule, but that rule concerns approval voting rather than single-vote elections, and unlike in Laslier’s rule, voters in threshold pivotality insist that the alternative they vote for must also be competitive

(Laslier, 2009). Other threshold models for generating candidate support have also been studied, but with different ideas of what a threshold is; see for example the modeling by Stiles et al. (2020).

An immediate challenge is that in threshold pivotality, cycles will be ubiquitous, and it is not immediately obvious if there is a cap on how long they can be. In a cycle of length 2, the mutual best response of all electors to the first vote total produces a second vote total, and the mutual best response to that second vote total returns to the first result. In this 2-cycle, two distinct vote totals repeat forever. We can also imagine a 3-cycle, where one vote total leads to another which leads to another which again produces the first. Indeed it does not immediately seem obvious that there is any cap on how long a cycle can be.

Before we begin any technical study of this ruleset, a brief interlude is required regarding the structure of preference orderings in the population of electors. I do not assume strict preference orderings, so the sincere utilities of any two candidates may be equal. In this case, tie-breaking is random. However, I will assume that whenever ties are broken randomly, the candidate that they favour does not change from iteration to iteration; relaxing this assumption would break the idea of a deterministic mapping from vote totals to vote choices. Instead, the elector performs a random tie-break the first time they encounter a particular tie, and in all future iterations they will always break that tie the same way. This means that we can freely speak about an elector's most-preferred candidate, since it is whichever candidate is involved in a first-place tie in their preference ordering and is chosen to break ties among that group of candidates. For example, if an elector ranks candidates A, B, C as $A \sim B \succ C$, then they arbitrarily choose either A or B to break ties, and whichever is chosen will be considered their most-preferred candidate.

There is one exception to the tie-breaking rule: if all of the candidates in S are tied for last place in an elector's sincere preference ordering, it is incoherent for the elector to vote for either one on strategic grounds, so in this case too they vote for their sincerely most-preferred candidate. But this also introduces a contradiction with the rule for breaking ties to obtain a most-preferred candidate: what happens if an elector has all candidates tied in last place? For the moment, while studying the abstract properties of threshold pivotality, I will assume that no elector ever has such a pathological preference ordering, or equivalently that any electors who have such an ordering must conduct arbitrary tie-breakers when all parties are tied for last. Either assumption suffices for the results in this chapter. But in the empirical examples of threshold pivotality that follow, I do include electors who have ties and I do not force them to perform tie-breaks, so the results of those examples may not always exactly match the abstract claims in this chapter.

So what are the possible end-states of this model? Let us assume that each elector begins by signaling their sincere vote, producing the initial vector of votes $\mathbf{v_0}$, or which we may distinguish from other ways of initializing the system (say by beginning with past election results) by denoting $\mathbf{v}_s$ to emphasize that it is the sincere vote totals. Either there is a τ-tie in this vector of sincere votes, or there is not. Use m to denote the number of candidates in the election. The following two theorems describe what happens in each case.

Theorem 1. *Suppose there is no τ-tie in the sincere vote total. Then in every iteration electors cast their sincere votes and there is never a τ-tie.*

Proof. By the definition of threshold pivotality without abstention, electors begin by casting their sincere votes, and by assumption there is no τ-tie. If there is no τ-tie in some iteration $k \in \mathbb{N}$, electors vote sincerely in iteration $k+1$, and again there is no τ-tie. So by induction electors always vote sincerely and no iteration contains a τ-tie.

$\square$

So if the model begins without a tie, then it immediately reaches a simple fixed equilibrium.

If there is no τ-tie in some iteration t, the electors will respond with their sincere votes, and if those votes have no τ-tie then Theorem 1 paints a simple picture of the result: an immediate fixed equilibrium. But if there is a τ-tie in the sincere vote totals, then the situation is more complicated. When there is a τ-tie in some iteration t, there may or may not be a τ-tie in iteration $t+1$. The voters might consolidate their votes on just one party in S and push it more than τ points above the nearest competitor, or the highest expected vote total might remain within τ of the next-highest total. The following theorem concerns what happens in the more complicated case that there is a τ-tie in the sincere vote totals.

Theorem 2. *Suppose there is a τ-tie in the sincere vote totals of a threshold pivotality game. Then the game reaches iterative equilibrium after at most $m - 1$ iterations, where m is the number of candidates.*

To prove this theorem we will use the following lemmas:

Lemma 1. *Denote the vote totals in iterations t and r of a threshold pivotality game by $\mathbf{v}_t$ and $\mathbf{v}_r$, and suppose that in $\mathbf{v}_t$ and $\mathbf{v}_r$ there is a τ-tie between the same set of candidates. Then $\mathbf{v}_{t+1} = \mathbf{v}_{r+1}$.*

Proof. By the definition of threshold pivotality, an elector responds to a tie between k candidates by selecting their most-preferred vote choice from among those k candidates. So all electors will cast the same vote in $t + 1$ as in $r + 1$.

$\square$

Lemma 2. *Suppose that more than one candidate in a threshold pivotality game is involved in a τ-tie in iteration t, so $|S_t| > 1$. Then $S_{t+1} \subseteq S_t$.*

Lemma 2 follows directly from the definition of the voting rule.

Now we can prove Theorem 2.

Proof. There is a τ-tie in $\mathbf{v_s}$, so electors who switch their votes will switch only to candidates in that τ-tie. At most m candidates are involved in that τ-tie. There are two cases: either $\mathbf{v_2}$ has a τ-tie or it does not.

First suppose there is no τ-tie in $\mathbf{v_2}$. Then electors return to $\mathbf{v_s}$, and the system is at 2-periodic iterative equilibrium.

Now suppose there is a τ-tie in $\mathbf{v_2}$. There are again two cases: either the τ-tie in $\mathbf{v_2}$ again contains all m candidates, or it does not. If m candidates are again included in a τ-tie in $\mathbf{v_3}$, then by Lemma 1, all future vote totals will equal $\mathbf{v_2}$. And if $\mathbf{v}_3$ includes a τ-tie between a set of k candidates, where $1 < k < m$, then, by Lemma 2, $\mathbf{v}_4$ will contain either the same k candidates or fewer. If it contains the same k candidates, the game is at iterative equilibrium. If it contains fewer, then in every following iteration the vote total can either enter iterative equilibrium in the same way, or the number of candidates in the τ-tie must decrease by at least 1. After $m - 1$ such decreases there will be only one candidate remaining, so $m - 1$ is the upper bound on the number of iterations for $\mathcal{M}$ to reach iterative equilibrium.

□

The crucial point in the proof of Theorem 2 is that, in a threshold pivotality game, k-cycling is necessarily the result of $k - 1$ consecutive τ-ties.

So threshold pivotality models reach iterative equilibrium quickly. But how complex can those iterative equilibria be? The following theorem shows that they must be quite simple.

Theorem 3. *If a threshold pivotality game has m candidates, then it can have an iterative equilibrium of at most cyclicity m.*

It is useful to first observe that any iterative equilibrium in a threshold pivotality game must include the vector of sincere votes.

Lemma 3. *If a threshold pivotality game is at iterative equilibrium in iteration t, then there is an iteration $t' > t$ such that $\mathbf{v}_{t'} = \mathbf{v}_s$.*

Proof. Suppose $\nexists t' > t$ such that $\mathbf{v}_{t'} = \mathbf{v}_s$. Then there is a cycle of $k \in \mathbb{N}, k > 1$ vote totals at iterative equilibrium, denote them $\mathbf{v}_1^\tau, \mathbf{v}_2^\tau, \ldots, \mathbf{v}_k^\tau$, such that $\mathbf{v}_i^\tau \neq \mathbf{v}_s \; \forall i \in \{1, \ldots, k\}$ and at some iteration r, $\mathbf{v}_r = \mathbf{v}_k^\tau \implies \mathbf{v}_{r+1} = \mathbf{v}_1$. Because none of these vectors equal $\mathbf{v}_s$, each must contain a τ-tie. If there is some $m \in \mathbb{N}$ such that each $\mathbf{v}_i$ contains a τ-tie between m candidates, then by Lemma 2 it cannot be true that $k > 1$. So instead suppose that the number of candidates in each τ-tie are not all equal. Then by the proof of Theorem 2, the number of candidates in the τ-tie at $\mathbf{v}_k^\tau$ must be less than the number of candidates in the τ-tie of $\mathbf{v}_1^\tau$. But then it cannot be true that $\mathbf{v}_{r+1} = \mathbf{v}_1$ if $\mathbf{v}_r = \mathbf{v}_k^\tau$, since that would require an elector to switch their vote to a candidate which was not involved in the τ-tie. So at least one of $\mathbf{v}_1^\tau, \mathbf{v}_2^\tau, \ldots, \mathbf{v}_k^\tau$ must equal $\mathbf{v}_s$.

□

Now we can use Lemma 3 to prove the theorem.

Proof. By Lemma 3 the vote totals at iterative equilibrium must include $\mathbf{v}_s$, so the cyclicity of the equilibrium is the number of iterations required for $\mathbf{v}_s$ to recur. By the proof of Theorem 2, at most m iterations are required. □

The preceding theorems add up to a picture in which, in the aggregate, the vote counts of a threshold pivotality game are quite orderly. But the electorate is composed of an arbitrarily large number of electors. Are the electors themselves similarly orderly, or could they be changing their behaviour in much more complicated ways that somehow do not complicate the aggregate vote counts?

The following theorem shows that they, too, cycle in a very simple way.

Theorem 4. *If a threshold pivotality game is k-cyclic, then all electors cycle with period k or period 1.*

Proof. Since the threshold pivotality game is k-cyclic, the vote totals at iterative equilibrium follow a repeating sequence $\mathbf{v}_s \to \mathbf{v}_2 \to \ldots \to \mathbf{v}_{k-1} \to \mathbf{v}_s$. An elector following threshold pivotality will always cast the same vote in response to some aggregate vote total. So, denoting elector r's vote choice in iteration t as v_t^r, any elector's sequence of vote choices in $\mathcal{M}$ must follow a repeating sequence $v_s^r \rightarrow v_2^r \rightarrow \ldots \rightarrow v_{k-1}^r \rightarrow v_s^r$.

It remains to be shown that k is the least cyclicity than an elector can have. Since every elector's vote choice repeats with period k, such a cyclicity would have to divide k. So, suppose an elector has cyclicity $j \in \mathbb{N}$ so $1 < j < k$ and $k \mod j \equiv 0$. Then $\exists\ n, m, q \in \mathbb{N}$ with $1 \leq n < q < m < k$ such that $v_n^r = v_m^r$ while $v_n^r \neq v_q^r$. Denote the set of candidates involved in a τ-tie in these iterations by S_h for $h \in \{n, q, m\}$. By Lemma 2, $S_n \supseteq S_q \supseteq S_m$. Denote the elector's vote choice in iteration n as A, so $v_n^r = A$, and denote $v_q^r = B$. Since $B \in S_n$, the elector must have $A \succ B$. So $v_n^r = B \implies A \notin S_q$. But $S_m \subseteq S_q$, so also $A \notin S_m$. This is a contradiction with the assumption that $v_m^r = v_n^r$.

□

We might have imagined that under a placid surface of unchanging aggregate vote totals, electors could perform arbitrarily chaotic switching that would not be reflected in the aggregate vote totals. But in threshold pivotality games that is not possible. Theorem 4 shows that when the aggregate vote totals are simple, every elector also behaves simply.[10]

What does it mean that the elector cycles with exactly period k or period 1, and cannot have a cyclicity that divides k? The following example illustrates the problem.

Example 1. Imagine a four-way race between parties $A, B, C,$ and D. Suppose that in the sincere vote totals there is a τ-tie between all four parties, and that the game reaches a 4-cyclic iterative equilibrium. By the proof of Theorem 2, the set of candidates S involved in a τ-tie begins with cardinality four and shrinks by one for three successive iterations, before returning to the sincere vote totals. So (labeling the candidates according to how many iterations they remain in S) we have the following sequence of τ-ties:

---

[10] On a technical note, the proof might appear to assume strict preferences, but strict preferences can be replaced with favouring one candidate in a pre-determined tie-breaker without loss of generality.

$$\begin{cases} \mathbf{v}_s: & S = \{A, B, C, D\} \\ \mathbf{v}_1: & S = \{A, B, C\} \\ \mathbf{v}_2: & S = \{A, B\} \\ \mathbf{v}_3: & S = \{A\} \\ \mathbf{v}_s: & S = \{A, B, C, D\} \end{cases}$$

In this election, what possible behaviours could electors follow? An elector may be 1-cyclic — for example, any elector who most-prefers A will always vote for A. An elector may also be 4-cyclic: for example, an elector with the preference ordering $D \succ C \succ B \succ A$ will vote in the sequence $D \rightarrow C \rightarrow B \rightarrow A \rightarrow D$.

We know that an elector cannot be 3-cyclic, because then they are not 4-cyclic; that would violate the rules of threshold pivotality because the elector will always respond the same way to the same vote totals, and every vote total is the same as the vote total 4 iterations before it. But an elector that is 2-cyclic will vote the same way every 4 iterations. So could an elector be 2-cyclic?

Suppose first that the elector begins by voting for A. Then the elector prefers A to every other alternative. So they will vote for A in every iteration, and have cyclicity 1. Now suppose instead they begin by voting for B. Then they prefer B in every iteration, including in response to $\mathbf{v}_s$, so they are 4-cyclic.

If instead an elector most-prefers C, they will vote for C in response to $\mathbf{v}_s$ and $\mathbf{v}_1$. If they hold $A \succ B$ then they will support A after $\mathbf{v}_2$, but after $\mathbf{v}_3$ they revert to C, so in this case they are 4-cyclic with the repeating sequence $C \rightarrow C \rightarrow A \rightarrow C$. Similarly if they hold $B \succ A$ they are 4-cyclic with $C \rightarrow C \rightarrow B \rightarrow C$. And if $A \sim B$ with both least-preferred, then they are 1-cyclic with $C \rightarrow C \rightarrow C \rightarrow C$. So an elector who most-prefers C will be either 1-cyclic or 4-cyclic.

Finally, can an elector who most-prefers D be 2-cyclic? If $D \succ A \succ \cdots$ then they are 4-cyclic with the repeating sequence $D \to A \to A \to D$. If $D \succ B \succ \cdots$, then they are 4-cyclic with $D \to B \to B \to D$. If $D \succ C \succ B \succ A$, then they are 4-cyclic with $C \to B \to D \to D$. And if $D \succ C \succ A \succ B$, then they are again 4-cyclic with $C \to A \to D \to D$. The same holds whenever there are ties in their preference ordering that are resolved by tie-breakers. Introducing two-way ties in last place (which is only relevant in the latter two preference orderings) changes the vote sequences but still produces 4-cyclic votes, and if $D \succ A \sim B \sim C$ then the elector is 1-cyclic, only ever voting for D. So in this final case, the elector again cannot be 2-cyclic.

Example 1 is not a special case that only applies to 4-way races or 4-cyclic vote totals: this was an example of the problem identified in the proof of Theorem 4. In a later section of this chapter we will study in a more general setting some related constraints that can be inferred about individual-level vote choices from aggregate vote totals.

One particularly interesting feature of the illustrative threshold pivotality setup is its two possible types of rest points. One rest point is the repetition of the sincere vote totals, which by Theorem 1 occurs if the sincere vote totals do not contain a τ-tie. But if there is a τ-tie in the sincere vote totals, it is still possible for the model to reach a stable equilibrium, as in several of the cases in the proof of Theorem 2. A rest point that does not contain the sincere vote totals is any vote total which contains a τ-tie, and which is also produced by electors' reponses to it. For example, imagine that the sincere vote total in a 4-way race contains a τ-tie between two candidates. Then every elector (except those who least-prefer both of the alternatives in the two-way τ-tie) will vote for one of those two candidates. And suppose that the response to this two-way τ-tie again produces a two-way τ-tie. Then that same vote total will repeat forever.

This leads to an important observation about the system: it captures the crucial dynamic that arises from the wasted vote logic of Cox (1994), that the only strategic rest point of the system is a tie or near-tie.

Remark 1. *The only strategic rest point of threshold pivotality is the classic wasted vote result: a* *set of candidates are involved in a tie or a near-tie, and nearly every elector supports one of those candidates.*

Cox (1994) identified simple conditions under which the Duvergerian pattern of two-way competition arises from wasted vote logic in a competition over one seat (and further that competition over M seats will give rise to $M + 1$ competitive candidates). Remark 1 shows that a similar result follows from the extremely simple voting rule of threshold pivotality: even the absolute barest wasted vote logic drags electors to support the candidates who are involved in ties. So Remark 1 is a replication of the Cox (1994) result in an iterative voting framework. In this setup, there is no reason to think that two-way ties are special, but even the wasted vote logic of Cox (1994) does not rule out non-Duvergerian results in which more than two candidates are included in the tie.

In the next section we will evaluate how often these different types of end states occur in threshold pivotality logic, using the example of a real election.

Threshold pivotality application

In practice, how reasonable a model of voting is threshold pivotality? And how complicated are the equilibria that arise when a model is run of an actual election?

I use the example of the 2019 Canadian election to check the cyclicities of equilibria in the model of a real election. Canada holds Single-Member District Plurality (SMDP) elections, in which candidates compete across 338 electoral districts in single-vote elections, and whichever candidate receives the most votes in a given district wins a seat in parliament.[11] Running a threshold pivotality game on this election requires two types of information: what is the distribution of electors' preferences over the actual political parties in that election, and what are the expected vote totals of the different parties?

During the 2019 election, the Canadian National Election Study asked tens of thousands of

---

[11] The number of seats will rise from 338 to 342 following the seat allocations that were announced in October 2021.

respondents to rate the major political parties on a thermometer scale, and the electoral district of every respondent was identified (Stephenson et al., 2020).[12] This means that the set of responses to thermometer scores gives a complete preference ordering over the parties, though the preference ordering can include ties, since respondents can give the same rating to multiple parties. To turn these preference orderings into sincere utilities, I take the conventional approach of normalizing the utilities between people: each voter has 0 utility attached to their last-ranked party, 1 utility attached to their second-last ranked party, and so on. In this way, in every district, we can take from the election study the sincere utility that dozens of randomly sampled citizens would obtain from the victory of each party that has a history of winning seats in parliament.

But an election does not take place between dozens of electors; real districts contain many thousands of electors. To get a realistic number of voters in a given district, I randomly sample preferences from the reported preference orderings in that district until there is one preference ordering for every elector who cast a vote there in the previous election.[13] The assumption is that the number of voters in each district does not tend to vary dramatically much between elections. This sampling is uniformly random, but with one constraint: for every vote that was cast for each party in the previous election, I sample a preference ordering that does not have that party in last place. This will turn out to be empirically unimportant (varying this assumption does not substantially change the resultant preferences), but it is enforced as a matter of internal self-consistency: strategic voting can never compel someone to vote for their last-ranked party, so without this constraint, the model would sometimes generate a population of electors that could not logically have produced the previous election's results, which contradicts the method's assumption that the voters in a district are rarely quite different from the voters in that district a few years ago.

In Chapter III, I will assess the accuracy of this method by comparing the sampled preferences to real election results. In Figure 8 of Chapter IV and in the surrounding text, I will sample

---

[12] Preference orderings are inferred using the "cps19_party_rating" battery, which is an individual thermometer rating for the major parties. From this I infer a ranking of the parties which must be complete (that is, omitting any respondents who did not respond to one or more thermometers) but may or may not be strict. The riding variable used is "constituencyname".

[13] For the sake of runtime, the models reported in this dissertation actually use $\frac{1}{10}$ the actual vote total, except where noted otherwise.

preferences according to the reported preferences in the national election study, and then count the expected election results based just on sincere preferences by summing electors' top preferences district-by-district. That figure shows that in the Canadian 2019 election, the preferences matched the winner in 96% of model runs, and the median estimated seat totals were within about 10% of every party's real total, and usually much closer. Importantly, though, that is not really a validation that the *preferences* were sampled correctly, because in between preferences and actual votes stands strategic considerations (among many other variables that might drag election results away from true preferences). But the reported preferences' similarity to the real election results is an informal check that they are at least reasonable approximations of voters' actual preferences.

This sampling strategy generates the first necessary piece of information. But to run a threshold pivotality game also requires that the electors have guesses about how much support each party is expected to receive. As a simple estimate that would be readily available to even minimally attentive electors, I use the region-by-region polls from the day that the election began.[14] Because polls at the electoral district level are rarely publicly available in Canada, it is reasonable that electors might assess the current distribution of support by party in their local district by checking the regional polling averages.

This setup provides electors with (a) fixed preferences, and (b) an initial guess about the competitiveness of the parties. But once electors have updated their vote choices, imagine that another poll is conducted which asks every elector to report their current intention. By consulting this poll as the new best measure of the parties' competitiveness, adjusting their vote choices accordingly, and then repeating this process, electors engage in an iterative threshold pivotality game where the initial vote total, instead of being their sincere vote choices, is based on empirical data about the parties' real competitiveness.[15]

---

[14] These figures are taken from the Canadian Broadcasting Corporation's 2019 poll tracker on September 10, 2019, available at https://newsinteractives.cbc.ca/elections/poll-tracker-archive/canada-2019/. The provinces of Alberta, British Columbia, Ontario, and Québec are individually tracked there. In that poll tracker the three maritime provinces as well as Newfoundland and Labrador are averaged together, and Saskatchewan and Manitoba are also averaged together. Regional polling is not available for Canada's three territories, each of which has only one electoral district, so in the place of regional polls for those three districts I substitute the nationwide polling averages.

[15] With this assumption we depart slightly from the properties of threshold pivotality in the previous section, since there we assumed that the system is initialized with the sincere vote total $\mathbf{v}_s$, but the differences are very subtle. If the

Throughout this section I will use just one random number seed for every run of the model, so the preferences are always the same, and I will vary certain parameters without varying the other details of the model. Other seeds produce substantively similar results.

So if electors hold the preferences that Canadians reported in their 2019 National Election Study, and assess the initial competitiveness of the parties by checking the public broadcaster's poll tracker in their region, how often do they settle down to one fixed election result? Figure 1 shows how common each cyclicity is for every possible threshold τ, where here τ represents the percentage of votes that need to separate a party from the highest vote-getter in order for electors to consider those parties to be in a τ-tie.

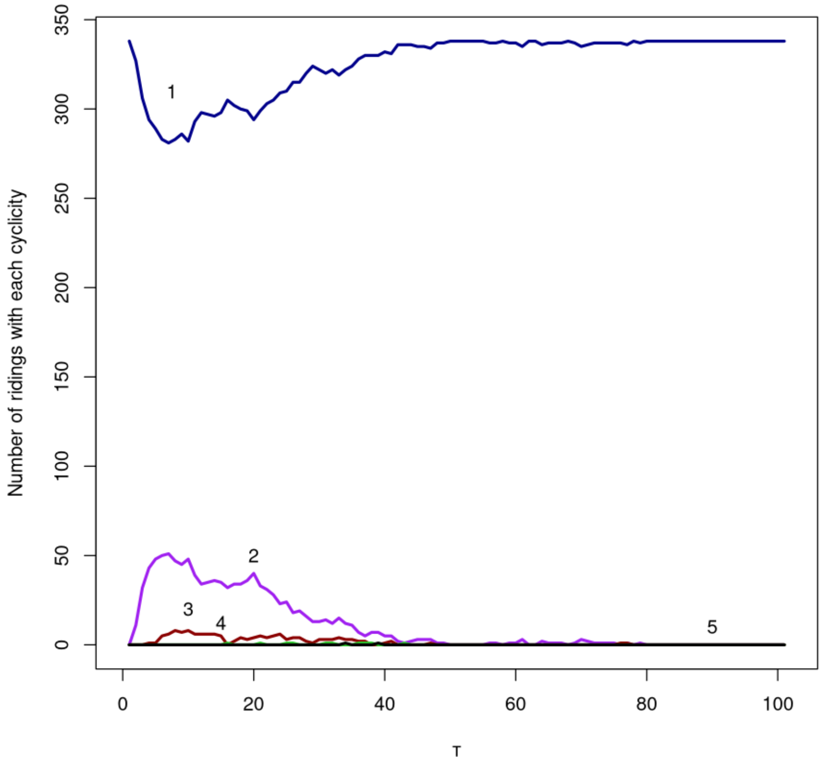

Figure 1: The number of ridings that result in equilibria with each possible cycle length. Each curve represents the number of equilibria that have a certain cycle length (labeled on the plot), for the τ specified on the x-axis.

---

sincere preferences are used instead of polling data, the results that follow are not substantively different.

When the threshold is exceptionally small, electors will almost always cast their sincere vote and then stick with that vote for every iteration. At the small-τ limit of $\tau = 0\%$, a τ-tie is an exact tie, and no exact tie is ever observed in this election. So at $\tau = 0$ there are no τ-ties. But incrementing τ up to 1 creates a τ-tie in any riding in which multiple parties are within 1% of the highest expected vote total of any party.

With τ in the range of about 5% to 10%, values which substantively may be a reasonable range for electors to use to evaluate whether or not more than one party is actually competitive, the number of 2-cycles rises to more than 50 out of 338 districts. In this range of τ there are also about a dozen 3-cyclical districts.

Districts with cyclicity of 4 or 5 remain very rare throughout the entire range of τ, but both can happen.

Table 1 gives examples of 1-, 2-, 3-, and 5-cyclic equilibria with various choices of τ.

τ	Riding	Iteration	Type	Con	Lib	NDP	Green	Bloc	Eq
0%	Ottawa Centre								
	(ON)	0	First poll	32.9%	38.8%	14%	10%		
		1	Sincere	17.7%	47.2%	24.5%	10.6%		✓
1%	Esquimalt								
	(BC)	0	First poll	32%	30.2%	17.3%	16%		
		1	Sincere	23.9%	22%	**26.7%**	**27.4%**		✓
		2	Heuristic	0%	0%	48.8%	51.2%		✓
5%	Skeena								
	(BC)	0	First poll	**32%**	**30.2%**	17.3%	16%		
		1	Heuristic	44.6%	55.4%	0%	0%		
		2	Sincere	**31.9%**	**29.9%**	**31.7%**	6.5%		✓
		3	Heuristic	**33.4%**	**37.7%**	28.9%	0%		✓
		4	Heuristic	44.6%	55.4%	0%	0%		✓
38%	Montcalm								
	(QC)	0	First poll	**21.4%**	**35.7%**	**9.6%**	**10.1%**	**18.9%**	
		1	Heuristic	**10.5%**	**24%**	**10.5%**	8.1%	**47%**	✓
		2	Heuristic	11.5%	**25.1%**	**13.1%**	0%	**50.4%**	✓
		3	Heuristic	**24.1%**	16.8%	0%	0%	**58.1%**	✓
		4	Heuristic	30.9%	0%	0%	0%	69.1%	✓
		5	Sincere	**9.5%**	**26%**	**11.3%**	**7.9%**	**45.2%**	✓

Table 1: Examples of 1-, 2-, 3-, and 5-cyclic results, varying τ and in various ridings.
Bolded numbers represent parties that are in a τ-tie, and the Equilibrium column identifies the iterations that are part of an iterative equilibrium. Esquimalt is the riding Esquimalt-Saanich-Sooke, Skeena is Skeena-Bulkley Valley, and the province is identified below each riding (due to space constraints).

In Table 1, Ottawa Centre with $\tau = 0$ is an example of how τ-cyclicity very frequently produces stable (1-cyclic) equilibria. In the last real Ontario-wide polling average taken before the election began, there was no exact tie between any two parties. So, by the rules of threshold pivotality, the electors respond with their sincere vote choices. Since those also do not contain an exact tie, this is a static equilibrium.

But raising τ to 1% opens the door to cyclic equilibria like the example of Esquimalt-Saanich-Sooke. In this case the Conservatives lead the province-wide initial polls and there is no party within 1% of their expected vote total, so electors respond with their sincere preferences. But there is a near-tie in their sincere preferences: the NDP and the Green Party are within 1% of the lead. Because no voters in this riding hold the NDP and the Green Party tied for last place in their sincere preference ordering, every elector is capable of choosing a preferred option between these two. When they do, however, the Green Party leads the NDP by more than 1%. So the electors next revert to their sincere choices, producing a 2-cyclic equilibrium.

Among Canada's 338 electoral districts, 78 have 5 parties that compete in their districts and were included in the battery of questions about voter preferences in the 2019 Canadian National Election Study: those are the 78 ridings of Québec, which are contested by the regional Bloc Québécois as well as the four other parties in the national party system. To sweep 101 values of τ across the 78 ridings of Québec meant running 7,878 models with the same distributions of preferences but varying values of τ. In all of these runs, there were only two cases that produced 5-cyclic equilibria. Table 1 shows one of those: Montcalm, with a very large τ of 38%. The initial regional polling average includes a τ-tie between every party, and then in each of the next 4 iterations, one party falls out of the τ-tie until ultimately the Bloc have accrued so much support that there is no longer any party within 38% of them, which prompts electors to return to a sincere vote distribution that also contains a 5-way τ-tie. The delicacy of this situation shows why it is exceptionally rare. But the fact that it happens for multiple ridings also demonstrates that very complicated equilibria do arise even in the very simple setup of a threshold pivotality game.

A central outcome of interest in the study of voting systems is how many competitive parties are supported. We have seen, for example in Lemma 2, that there is a tight connection in threshold pivotality games between the cyclicity of an equilibrium and the number of parties involved in a τ-tie. Another result, Remark 1, connects the rest points of threshold pivotality to Duverger's classic result that in a single-member plurality contest two parties should be expected to field competitive candidates. In the empirical example of Canada's 2019 election, how common was this Duvergerian outcome?

As a first analysis, we separately check those ridings in which the final vote total was induced by the heuristic, rather than just being a repetition of the sincere vote total. The reason for separating out heuristic vote counts is that it is exceptionally unlikely that there will be any riding where a party will receive exactly 0 sincere votes; we should always expect these ridings to have 4 competitive parties outside Québec and 5 competitive parties in Québec.

For every possible value of τ, Figure 2 shows the proportion of all ridings that result in a Duvergerian two-party equilibrium (blue), and the proportion of just those ridings in which electors cast heuristic votes rather than just sincere votes (red). So the blue curve is the proportion of all districts in which threshold pivotality logic reduces the number of vote-getting parties down to just two, whereas the red curve is the proportion of non-sincere vote counts that have just two parties.[16]

How do we get single vote counts out of iterative equilibria? For the moment, I simply count the parties in the first vote total that was a repetition of a previous vote total: that is, the vote count in the moment that iterative equilibrium was identified. The implications of this decision will be explored shortly. If instead we drop the iterative equilibria and only consider ridings with fixed equilibria, the curves look nearly indistinguishable.

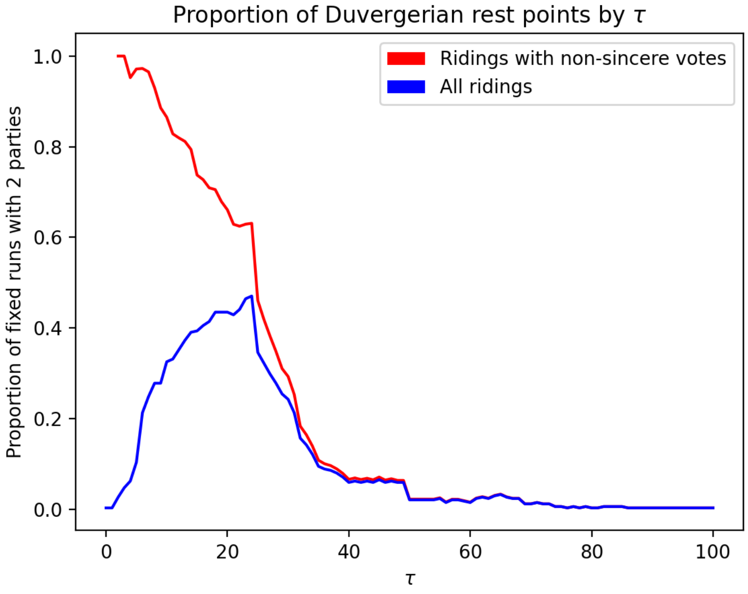

Figure 2: Proportion of ridings in the 2019 Canadian election that produce Duvergerian equilibria. The model uses threshold pivotality logic, with each possible value of τ. The red curve excludes the ridings that result in fixed repetitions of sincere vote totals, while the blue curve includes all ridings.

---

[16] This may be a harsher requirement than necessary: the curve actually represents the proportion of ridings in which two parties receive more than 0 votes. This is a conservative decision, because it excludes any ridings where electors have two parties tied in last place and therefore vote for another party. So it may understate the number of ridings that are effectively Duvergerian, where more than two parties technically receive some votes, but only two parties receiving meaningful vote totals.

When $\tau = 0$, there will never be a τ-tie, because in no district is there an exact tie in the regional polling data, so every fixed equilibrium is just repetition of the sincere vote totals. For reasonable values of τ (from very small values up to about 25%), the proportion of Duvergerian equilibria among districts with non-sincere vote counts falls from every riding having a Duvergerian two-party outcome down to about $\frac{2}{3}$ of ridings. In the same range, the proportion of ridings overall that produce a Duvergerian result rises to nearly $\frac{1}{2}$. In the limit as τ rises to 100%, every riding features a τ-tie between every party, so the heuristic vote totals approach the sincere vote totals as threshold pivotality induces electors to never switch away from their sincerely most-preferred party, and the proportion of Duvergerian equilibria falls to 0.

For reasonable values of τ, threshold pivotality produces a Duvergerian outcome in a substantively believable number of districts. Imagine that, on the day the 2019 Canadian election begins, voters consult the polling average for their region to see whether or not multiple parties are within 10% of first place. In many ridings where there are multiple parties within 10% of first place, the electors settle on an equilibrium in which everyone supports one of those competitive parties, but not necessarily their sincerely most-preferred party. In about 90% of such districts, this heuristic process produces competition between only two parties. But it only produces this outcome in about half of all districts. Canada does not have a two-party system, and the initial polling data in this election show how far it is from only two parties enjoying any support. And yet, there are only two parties in the modern party system which have any history of forming government. For reasonable values of τ, threshold pivotality — which is perhaps the simplest way to nudge electors towards preferring more competitive parties — closely matches the real substantive situation. In Figure 2, about half of all districts are Duvergerian, while nearly all of the districts in which electors are prompted to think strategically are Duvergerian.

As a basis for comparison, the Effective Number of Electoral Parties in a typical run of the heuristic model with $\tau = 5\%$ is 2.76, and the value in the real election was 2.77. Just as the number of Duvergerian districts is believable, so is the Effective Number of Electoral Parties.

This has showcased a new way of producing Duvergerian results, and it also emphasizes a new angle on the famous idea: perhaps there is nothing special about two-party competition in a race for a single seat, but rather that when electors follow extremely simple and minimalistic strategic voting heuristics, two-party competition is simply a much less delicate and therefore a much more common outcome. See, for example, how easy it is to sustain an iterative equilibrium between two parties in Table 1 compared to an iterative equilibrium that involves 5 or 4 parties.

At the same time, this finding also forced us to confront a limitation of the setup so far: we adopted only a very simple rule for obtaining a single vote count from iterative equilibria. From Lemma 2 we know that in threshold pivotality, the number of competitive parties is not constant between iterations in an iterative equilibrium of cyclicity greater than 1. And crucially, this is the same problem that is holding us back from another empirically interesting question: how closely do the results of threshold pivotality resemble the real election results? As it stands we cannot confidently tally up the number of ridings won by each party because we have no tool for identifying the winner of a district at iterative equilibrium. So how can we take a riding in which there is endless switching between different vote totals and identify the number of parties that would be competitive in that district?

There is no one correct answer: since there is no fixed equilibrium in these ridings, there cannot be one static result that is the most correct. So I propose the following five different stopping rules.

↪ **Election day:** Stop every riding after an arbitrary fixed number of iterations.

↪ **Random stopping:** Stop every riding after a random number of iterations.

↪ **Recurrence:** Check the party that wins most often across the vote totals in the iterative equilibrium and declare that party the winner.

↪ **Last candidates standing:** Pick the last iteration before the sincere vote total.

↪ **Default to sincere:** In any iterative equilibrium, simply choose the sincere vote total as the election's final result.

After quickly discussing each rule's motivations, we will study some of their expected properties, and then examine how they affect the model results in the example of this real election.

The election day rule is meant to straightforwardly mimic reality. If we imagine that there is some mapping between iterations of the model and the real time in an election (which is a tempting analogy, since iterations are sequential, but also highly tortured whenever we try to assign specifics — there is no obvious way to answer questions like whether one iteration is more like a day or a week or a month, or even if each iteration must represent the same amount of time) then we can imagine the election stopping rule as being a fixed date on which votes are cast. Whatever voting intentions electors happen to hold on election day, those votes will be the election results. It is important to note that the election day rule does not interrupt voters while they are not yet at an equilibrium, so it never returns out-of-equilibrium vote choices. Instead this rule waits for iterative equilibrium to begin, then waits for a fixed number of iterations more, and then holds the election. The crucial point is that each vote total within an iterative equilibrium is an equilibrium result under the rules of the system; the only question is which part of the equilibrium will serve as the fixed result.

The next two rules are each aimed at addressing a different property of the election day rule that might not be desirable.

The random stopping rule, instead of waiting a fixed number of iterations in every riding before holding the vote and declaring a fixed result, stops at a random moment. Will this ever be systematically different from the election day rule? To see why it might, imagine that the election day rule waits an arbitrary 10 iterations after the beginning of iterative equilibrium before it stops and reports the current vote total. Using for example $\tau = 10\%$, a glance at Figure 1 shows a severe downside to this idea: most iterative equilibria have cyclicity 2, and some have cyclicity 3. So suppose that the election day rule identifies an iterative equilibrium immediately after the sincere vote total has repeated. In a 2-cyclic iterative equilibrium, by Lemma 3 there can only be two vote totals: the sincere vote total, and a heuristic vote total in which there is no τ-tie. Since 2 divides 10, the election-day rule will just report the heuristic vote total for every 2-cyclic equilibrium. If instead the election day rule reports the result *on* the 10th iteration (so, *after* 9 iterations), it would always report the sincere vote total for every 2-cyclic iterative equilibrium. Consulting Figure 1 with our arbitrary example of $\tau = 10\%$ shows that this could have a major effect on the model results. So an odd feature of the election day rule is that, depending on the value chosen, it might either very closely resemble the sincere vote totals, or not resemble them at all. The random stopping rule avoids this obstacle, but with the downside of introducing extra randomness into the model's output.

The recurrence rule takes the opposite approach to tackling the problems with the election day rule. The results of the election day rule will vary considerably depending on the number of iterations chosen before the model stops. The recurrence rule uses a more stable idea: it checks which parties are winning in each iteration at iterative equilibrium, and then it assigns the riding to the candidate that leads in the largest number of iterations. Note that in the limit as the number of runs increases, the random stopping rule should converge to the results of the recurrence rule. There are two downsides of the recurrence rule. First, candidates might lead in an equal number of runs, in which case a tie-breaker is required. An obvious tie-breaker is to break ties uniformly randomly, and I will use that when applying the recurrence rule. But the second downside is more damaging: the recurrence rule can identify a most common winner, but there is no obvious way to pick vote totals. So the recurrence rule can only be used in situations where all that is needed is to identify a single winner per riding.

The last two rules represent two different principled approaches to resolving the iterative equilibrium.

The last candidates standing rule, inspired by the patterns observed while proving Theorem 2, interprets the winnowing of candidates from iteration to iteration in the iterative equilibria of threshold pivotality games as a sequence that itself selects the election winner.

Defaulting to sincere represents a belief that iterative equilibria do not include a single best response or fixed winner, and so in those cases the sincere vote totals serve as a backstop.

These rules all stand in contrast to a major category of stopping rule that is not considered here: anything which introduces a vote total that does not appear in the iterative equilibrium. One might for example imagine taking the set of vote totals at iterative equilibrium and averaging them to produce the result of the election. This sort of rule is not considered because election models are motivated by a focus on equilibrium results. If the electors were presented with a vote total that is not part of the iterative equilibrium, it would presumably push them to a new equilibrium (fixed or iterative) which may or may not include that other vote total — so this other vote total may not even be part of any equilibrium. Because we are only concerned with vote choices at equilibrium, the stopping rules only consider vote totals that are part of the iterative equilibrium.

So, how do these voting rules play out in the example of a real election? Now that we have methods for identifying a winner in every riding, we can also answer the question: how closely does threshold pivotality logic mimic real election results?

Before answering those questions, it is worth taking a moment to consider the role of questions like this throughout this dissertation. For checking the validity of models in this chapter as well as at the end of Chapter III and throughout in Chapter IV, I aim to check what Whicker and Sigelman (1991, p. 66) called "event validity": how much the results of a computer model of a political process resemble the real results of that process. Some models are not intended to have substantial event validity, and some models of politics are so abstract that it is hard to even imagine what "event" the event validity would refer to. There have been subtle and longstanding debates over whether all models or indeed any models should have explicit empirical implications, whether there are models that can get away without having any such thing as "validity", and how to check the validity of other models. But the models in this dissertation are of a simpler kind. I do aim to make stylized points about elections, and of course these models are simplifications of reality. But in this chapter I explicitly aim to check how accurate the minimalistic voting rule of threshold pivotality is when used to model the example of a real election, and in future chapters I will ask the same questions about more complicated models. So the appropriate way to assess the performance of the model is to ask: how close does the model come to the real election results? In Chapter IV I will describe similar reasoning for the questions asked in that chapter.

So, the crucial feature in this validation is the following. Threshold pivotality is an exceptionally simple and stripped down model of strategic voting. We should expect there to be a difference between the sincere vote counts and real election results, because real election results include strategic behaviour (among other things). So if threshold pivotality models strategic behaviour, then the results of a simulation that includes threshold pivotality should (for reasonable values of τ) be not merely close to the real election results: they should also be closer to the real election results than a purely sincere voting model.

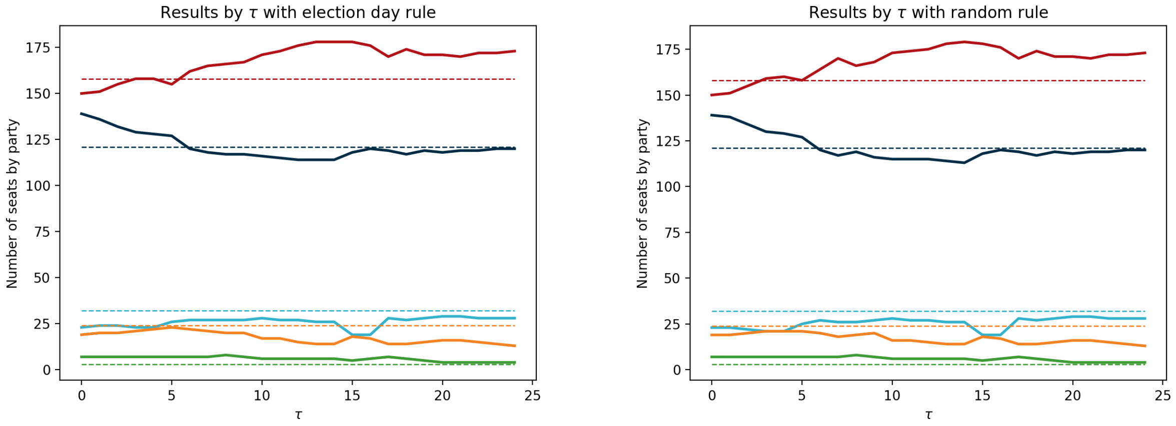

(a) Election day rule (b) Random stopping rule

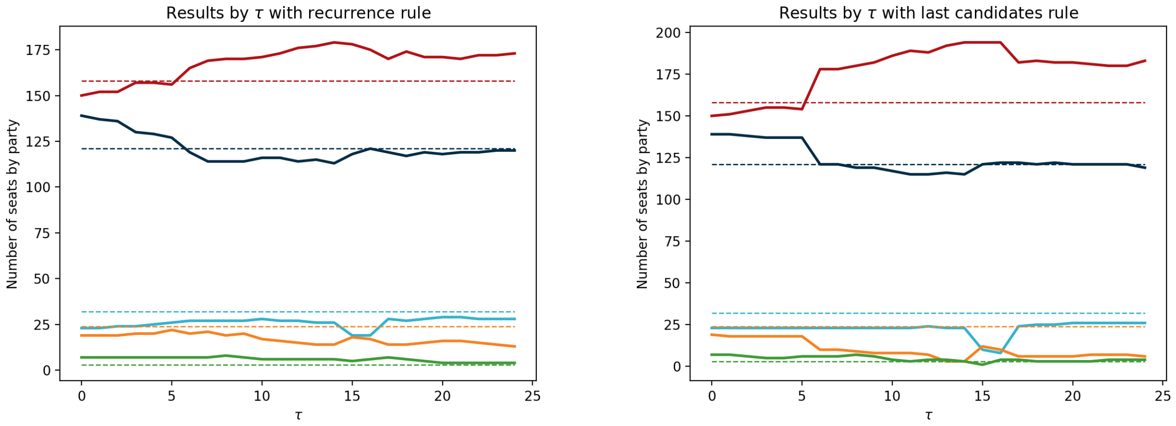

(c) Recurrence rule (d) Last candidates standing rule

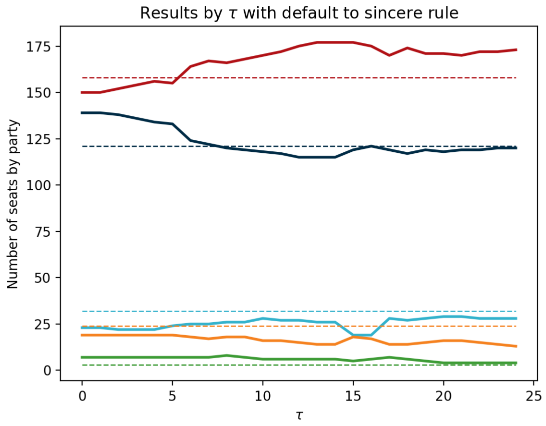

(e) Default to sincere rule

Figure 3: The results of a simulated 2019 Canadian election using threshold pivotality. τ ranges from 0% to 25%. Dotted lines represent the real election results for each party, and solid lines show the simulated results, for each of the five parties in parliament: Liberals (red), Conservatives (dark blue), Bloc (light blue), NDP (orange), and Green (green).

Figure 3 shows that, for a reasonable range of τ, the voting rule does indeed drag the simulated results very close to the real election results. When $\tau = 0$, there is never a τ-tie, so the results of every district are just the sincere vote totals. As I will show in Chapter III, the party preferences reported in the 2019 Canadian National Election Study are quite close to the real election results if they are sampled with replacement and used as the vote totals for a perfectly sincere election, but they do overstate the strength of the Conservative Party and understate the strength of the Bloc Québécois compared to the real election results. Several of the stopping rules do correct for this to some extent.

As τ is increased to between about 3% and 5%, under every stopping rule, the threshold pivotality vote totals for the Liberal Party are extremely close (and sometimes identical) to the real Liberal Party seat total. In the same range, the Conservative Party seat totals fall towards their real seat totals, though still produce a larger total for that party than it had in reality. However, above $\tau \approx 6\%$ every stopping rule produces simulated seat totals for the Conservative Party that are close to its real seat totals. Every stopping rule and value of τ underestimates the Bloc and the NDP while overestimating the Green Party, but the errors appear to be smallest in the range around $\tau \approx 5\%$.

Does the closeness of the simulated vote totals depend on which stopping rule is used? Figure 3 shows that, in the example of the 2019 Canadian election, stopping rule is emphatically not important. The curves of even dramatically different stopping rules closely resemble each other. The most different stopping rule is the last candidates standing stopping rule; because this is the only stopping rule that picks iterations based on the vote totals of those iterations themselves, it appears that this selection method does systematically track with the specific party that is in the lead (and it typically benefits the Liberal Party, and does so more for larger values of τ). But the other stopping rules look highly similar, and even this stopping rule never appears to depart from the trend of the other stopping rules by more than a few percent of the total number of seats. With reference to Figure 1, the reason that all of the rules produce more or less similar results is straightforward: in each case only about $\frac{1}{6}$ or fewer of the seats have iterative equilibria. Since there is also no obvious reason that the presence of iterative equilibria in a district would correlate with one specific party leading, it is not clear that the resolution of these iterative equilibria should be systematically related to the party-by-party election results.

One property of interest that is not obvious by looking at the party-by-party seat totals in Figure 3 is which value of τ comes closest to the real election results. Figure 4 shows the error for each value of τ and all five decision rules. This number is the sum of the absolute differences between the simulated seat totals and the real seat totals, divided by two since each mistake is zero-sum (this is the count of how many seats in parliament were mis-allocated by the model, away from one party and towards another, which would be double-counted by just summing the party-by-party seat differences). So the number of mis-allocated seats shown in Figure 4 is the number of seats that were won by a party in the real election but assigned to a different party in the simulation, presented for all five stopping rules and by every value of τ.

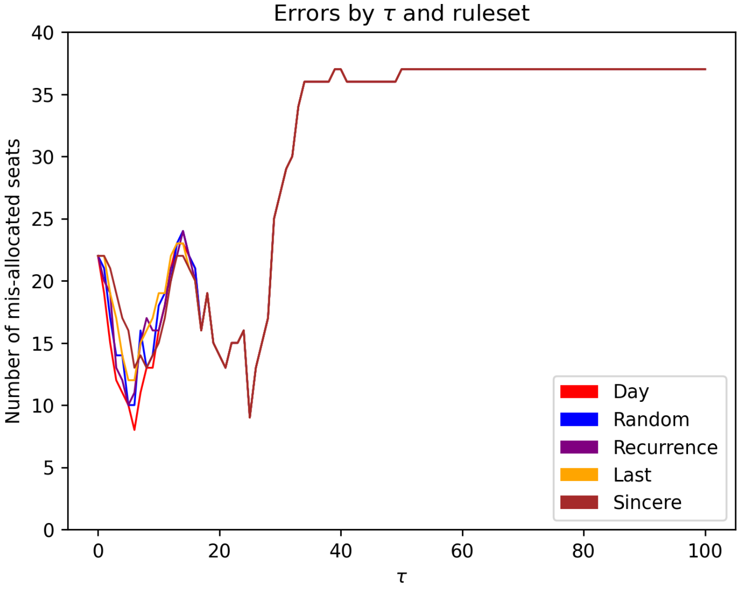

Figure 4: The number of mis-allocated seats under threshold pivotality.

The globally best model is at $\tau = 6\%$ under the election day rule, with the number of days set to zero (so, the model result is the iteration in which iterative equilibrium is first identified). This rule only mis-allocates 8 seats in parliament when $\tau = 6$. As seen in Figure 3a, this is approximately the value of τ at which the simulated seat counts of the two largest parties match their real seat counts, the NDP is less under-estimated than it is for higher values of τ, and the Bloc and Green party are close to their real results.

Below $\tau \approx 20\%$, the errors made by each stopping rule are distinct, though often similar. But above about $\tau \approx 15\%$, the errors made by every stopping rule cancel each other out, and all of the stopping rules begin to follow almost identical trajectories. They converge around $\tau \approx 20\%$, by which point there are very few iterative equilibria (as seen in Figure 1) and consequently very few opportunities to resolve iterative equilibria differently.

Of course, running five stopping rules for 101 values of τ means that the plot includes the results of 505 models, so maybe it was to be expected that at least one of those 505 models should come very close to reality. But Figure 3a contains a far more important result: the figure shows that for reasonable values of τ (between 1% and 20%, say), no rule mis-allocates more than 25 of the 338 seats in parliament. In fact, every rule manages a similar feat to the election day rule around $\tau \approx 6\%$, reliably producing errors below about 10 or 15 seats. Further underscoring threshold pivotality's added value is that the "default to sincere" rule is the very worst-performing rule. So picking a substantively reasonable τ, between about 1% and 15%, along with any stopping rule, yields a simulation result that matches the real election result in about 95% of the seats in parliament.

The threshold pivotality voting idea, when implemented with a small value of τ and using real reported preferences and polling averages, can be extremely successful in replicating election results. The nearest model — with arguably the most obvious simple stopping rule, and a substantively reasonable value of τ — successfully matched the election winner in 330 out of 338 ridings, and other stopping rules in neighbouring values of τ performed nearly as well.

We have compared threshold pivotality to just using the sincere preferences from the Canadian National Election Study, and we found that it did add value over just using the sincere preferences. But there is one more detail that needs to be checked. This model actually had three ingredients: the preferences from the election study, threshold pivotality, and the regional polling data that we used to initialize the system. So before concluding that threshold pivotality itself is an extremely successful model of the election, we should ensure that the regional polling data was not secretly driving its success. To do so, I simply multiply the regional polling proportions for each party by the number of seats in each region. The result of this naïve allocation is very different from the real election results. In the real election, the Liberals received 158 seats, the Conservatives 121, the Bloc 32, the NDP 24, and the Greens 3. When allocating seats according to the polling averages on the first day of the election the Liberals win 119 seats, the Conservatives 119, the Bloc 15, the NDP 48, and the Greens 37. Evidently the success of τ pivotality is not driven by the polling averages, but is a success of the heuristic itself.[17]

Having examined the abstract properties of threshold pivotality and demonstrated its success in modeling strategic voting using the example of the 2019 Canadian election, we next turn to the properties of a family of games that threshold pivotality is just one example of: iterative election games where electors map a distribution of votes onto some vote choice with certainty.

Deterministic election games

Threshold pivotality models are one example of a broader type of iterative election game, described by the following definition:

Definition 2. *An iterative election game is* ***deterministic*** *if it maps any aggregate vote count onto exactly one aggregate vote count in the next iteration.*

As a quick terminological note, I use “election game” and “voting game” interchangeably to describe the types of models under consideration here, and occasionally I preface either term with the word “iterative”. These should all be understood to be the same thing, and refer to the sorts of iterative computer models that we are concerned with in this chapter. This should not be confused,

---

[17] We of course do not know how the support in the regional polling averages breaks down by electoral district, so in the spirit of D’Hondt allocation I round down and then allocate remainders in decreasing order of vote share. This affects about two dozen seats, though it cannot replicate the process by which district by district elections will always push (for example) the Green party seat share below its regional polling performance.

for example, with the idea of a "deterministic voting game" as defined by (Coughlin, 1992, p. 21) to capture situations in which candidates are able to perfectly predict the response that their strategic maneuverings will inspire in voters. In cases where that confusion would be particularly plausible, such as when I specifically discuss deterministic iterative voting, I will use "deterministic election game" to distinguish it from both Coughlin's "deterministic voting game" and from **probabilistic iterative election games**, which can with positive probability map some aggregate vote count onto any one of several different aggregate vote counts in the next iteration. The specific contrast with the latter is whether or not any elector randomizes their vote. If every elector will always respond to a given vote distribution by selecting the same pure strategy (that is, some vote choice with certainty), then the vote total in the following iteration is also deterministic, because it is a sum of deterministic choices. So for vote totals to be mapped probabilistically from one iteration to the next, at least one elector must be choosing randomly.

In what situation will an elector choose randomly? Under the calculus of voting, they should only choose to randomize between vote choices if no vote choice dominates the other: there need to be at least two vote choices with the same expected utility. This means that in a large electorate with varied preferences, iterative election games should never be truly deterministic. For example, suppose there is an elector in an iterative election game who has a preference ordering in which candidates A and B are tied, and are preferred to at least one other candidate. Then the iterative election model cannot be deterministic, because we can construct vote totals in which the elector has no dominant pure strategy: this will be true for any vote total in which A and B are expected to receive the same number of votes, and voting for any other party has a lower expected utility than voting for A or B. And of course, electors with strict preference orderings can also pursue mixed strategies, whenever a difference in vote counts balances out a difference in sincere utility to produce equal expected utilities among the pure strategies that dominate all others.

However, it less clear how practically important this is. In the public opinion data that was used for the application to the 2019 Canadian election in the previous section, no parties are ever expected to receive the same number of votes. In subsequent iterations, when simulating electorates with thousands of electors, it presumably remains extremely rare for parties to have identical expected vote counts. And it is also substantively reasonable to expect that vote counts should rarely exactly balance out differences in sincere utilities to induce mixed strategy behaviour in electors with strict preference orderings (or, at least, no ties among the candidates who have the highest expected utility for that elector).

So while almost no iterative election model should ever be completely deterministic, in cases where electors are explicitly calculating exact pivotality probabilities, we should expect accurate election models to mostly be deterministic models with some noise induced by a few electors occasionally pursuing mixed strategies. In the extreme case in which all or nearly all electors pursue mixed strategies, the situation is not so dire, because that is a much better-studied realm: that of probabilistic voting games, about which a great deal is already known (Coughlin, 1992). But for simpler heuristics, the model might be exactly deterministic; threshold pivotality models, for example, are perfectly deterministic iterative election games.[18]

Along these lines, Theorem 5 underscores the fragility of Deterministic iterative election games, which cannot include even one probabilistic voter.

Theorem 5. *If an iterative election game includes a vote distribution that an elector maps onto multiple different vote choices with positive probability, then the aggregate vote total does not necessarily cycle.*

Proof. At iterative equilibrium the vote totals satisfy $\mathbf{v}_t = \mathbf{v}_{t+k}$ for some $k \in \mathbb{N}$ and any t at equilibrium. Suppose that one of the vote totals prompts at least one elector to choose between multiple different vote totals with positive probability. Then it is not necessarily true that $\mathbf{v}_{t+1} = \mathbf{v}_{t+k+1}$. This is a contradiction with the claim that the model cycles.

□

---

[18] So long as each electors breaks a given tie in the same way in every iteration that they confront it.

Of course, the proof of Theorem 5 also makes it clear how a deterministic iterative election game can cycle even when there are electors who vote probabilistically: cycling is possible so long as they respond deterministically to all of the vote totals that appear in the cycle. This is a tremendously less restrictive requirement.

Corollary 1. *An iterative election game maps any aggregate vote count onto exactly one vote count when it is at iterative equilibrium if and only if all electors respond deterministically to every vote total that appears in the model's iterative equilibrium.*

It is also the case that, in an electorate that does not reach iterative equilibrium, it is possible for nearly all electors to vote deterministically. Consider for example an iterative election game of N electors and m candidates in which $N-1$ electors vote deterministically and one votes uniformly randomly. Suppose that the model reaches a 2-cycle: candidate A leads in even iterations, and candidate B leads in odd iterations. There are deterministic voting rules under which this pattern could continue without the vote total ever exactly cycling, since one elector is voting randomly, and indeed in which all other electors may adopt cycling strategies in which they switch 2-periodically from supporting either A or B to supporting some other candidate depending on which candidates are in the lead (threshold pivotality is an example of a voting rule that could produce this situation). So every elector except one is voting deterministically and cyclically, but the vote totals never cycle.

But this example also illustrates why the insights we draw from considering deterministic iterative election games can be quite relevant to the study of probabilistic voting games. In the situation where $N-1$ electors are cycling and one elector votes randomly, the vote totals are almost exactly the same as if 0 electors voted randomly, and that model's results will be very close to the iterative voting game's results. One elector voting randomly could also prevent other electors from cycling and make the results very different from the iterative voting game, but so long as that is not true, the insights of deterministic election games are useful for very slightly non-deterministic election games.

Or imagine an electorate in which very few people are ever truly indifferent between more than one highly competitive party. In that polity, most of the time people will respond to expected vote totals in a way that is highly predictable to anyone who knows their sincere preference ordering, but a small number of people will choose randomly between at least two alternatives. This is essentially a cycling electorate with a small amount of noise, which by the construction of the example is irrelevant to the relative performance of the candidates. So deterministic iterative election games — depending on what the voting ruleset is exactly, and on the number of electors and the number of candidates — can produce similar aggregate vote counts even when there is a small number of electors voting non-deterministically. Substantively, this might be a fair description of electorates where there are not very many swing voters between any two highly competitive parties.

Now that we have developed a sense of the scope of deterministic iterative election games, we next investigate their properties. Theorem 6 is a simplifying fact, demonstrating the applications of the iterative equilibrium idea in Definition 1.

Theorem 6. *Every deterministic election game with a finite number of voters and candidates will reach either a fixed equilibrium or an iterative equilibrium.*

Proof. Enumerate the r possible values that the aggregate vote counts can adopt. In iteration $r+1$, by the pigeonhole principle, there must be at least one vote total that appears in both iterations x and y, where $1 \leq x < y \leq r+1$. By Definition 1 and Definition 2, the game has an iterative equilibrium of cyclicity $y - x$.

□

The number r is also known, in terms of the number of voters and candidates. Because it only matters how many voters are voting for each candidate, but the candidates themselves are distinguishable, this is an example of the classic "stars and bars" problem, solved for example in Planck (1901) and Feller (1968, p. 38).

Corollary 2. *The cyclicity k of a deterministic iterative election game is bounded above by*

$$k \leq \frac{(N+M-1)!}{N!(M-1)!}$$

Of course, this potentially enormous number is mainly offered for the psychological comfort that deterministic election games must have some kind of an end state. It dramatically overstates the real bound of any realistic election model. For an idea of how to find smaller bounds for more specific types of models, consider the basic property of any deterministic strategic voting model: electors will only update towards more competitive parties, away from less competitive parties. Pick any vote total $\mathbf{v}$, and consider the set L of vote counts in which the best-performing party has fewer votes than it has in $\mathbf{v}$. L can also be enormous.[19] In a deterministic strategic voting model, $\mathbf{v}$ can never lead in one iteration to any vote count in L. But here we have to be careful: there could be a chain of vote counts that lead from $\mathbf{v}$ to some vote count in L in more than one iteration, if other parties are boosted so far ahead of the best-performing party that voters begin to flee from it. But this line of reasoning suggests that a few more constraints on the electors' behaviour might rule out large categories of vote totals.

One fact that we have not yet directly addressed is that iterative election games can have multiple cyclicities. The next theorem addresses this question, showing that deterministic election games have several different cyclicities only when all of the cyclicities are multiples of the smallest cyclicity.

Theorem 7. *Suppose that a deterministic election game is j-cyclic, and is not also h-cyclic for any $h < j$. Then it is also k-cyclic, with $k > j$, if and only if j divides k.*

---

[19] If the best-performing party is expected to win a large share r of the votes, then there are many more possible vote counts in which that party has a larger vote share than than r compared to vote counts in which it has a smaller share than r. Then the general question is: how large is L as a function of r?

Proof. First we prove that the voting game can be both j-cyclic and k-cyclic only if j divides k. Consider a k-cyclic deterministic iterative voting game that is also j-cyclic, for some $j \in \mathbb{N}$. Then $\forall n \in \mathbb{N}$, there is an iteration t such that $\mathbf{v}_t = \mathbf{v}_{t+n\cdot k}$ and also $\mathbf{v}_t = \mathbf{v}_{t+n\cdot j}$. So the game reaches both the repeating sequence $S = \mathbf{v}_t, \mathbf{v}_{t+1}, \ldots, \mathbf{v}_{j-1}, \mathbf{v}_j$ and the sequence $T = \mathbf{v}_t, \mathbf{v}_{t+1}, \ldots, \mathbf{v}_{k-1}, \mathbf{v}_k$, and $|S| = j$ while $|T| = k$. Since j is the least cyclicity of the system, $\mathbf{v}_j$ does not appear more than once in S. Then consider the greatest $m \in \mathbb{N}$ such that $m \cdot j \leq k$. By the determinism of the system, $\mathbf{v}_{mj}$ appears in S. If j does not divide k, so $mj < k$, then there are $w = k - mj$ elements between $\mathbf{v}_{mj}$ and $\mathbf{v}_k$, and $w < j$. Since the system is k-cyclic, $\mathbf{v}_{mj+w+1} = \mathbf{v}_t$. But since mj is a multiple of j, it is also true that $\mathbf{v}_t = \mathbf{v}_{mj}$. And since $w < j$, and because j appears only as the jth element in S, then $\mathbf{v}_{mj+w} \neq \mathbf{v}_j$ and consequently $\mathbf{v}_{mj+w+1} \neq \mathbf{v}_t$. This is a contradiction, so a deterministic election game can be both j-cyclic and k-cyclic only if j divides k.

Next we show that if a deterministic election game is j-cyclic, then it is also k-cyclic whenever j divides k. Since $k = nj$ for some $n \in \mathbb{N}$, consider the sequence formed by n repetitions of S. This is a k-cycle in the vote totals.

□

All this talk about cyclicity might not be very useful if a game could be wantonly 3-cyclic and 5-cyclic and 11-cyclic without any limitations. Theorem 7 simplifies the situation: the cyclicities of a game all share a common divisor greater than 1, and that common divisor is the smallest cyclicity of the game. So a game that is 3-cyclic must also be 6-cyclic, 9-cyclic, and so on, because if the vote totals repeat every 3 iterations then they will also repeat every 6 iterations and every 9 iterations. But that game cannot also be, say, 5-cyclic, because 5 is not a multiple of 3.

We already knew that there will always be a smallest cyclicity of any game at iterative equilibrium: if a game is 1-cyclic, it has reached a fixed equilibrium, so the smallest iterative equilibrium is 2-cyclic. Theorem 7 shows that any game at iterative equilibrium actually has infinite cyclicities: it is not just k-cyclic, but also repeats in every multiple of k.

How should we talk about the cyclicity of games, then? An important point is that while (say) a 2-cyclic game is also n-cyclic, where n is any even number, it is not the case that every

n-cyclic game is also 2-cyclic. If a game has the repeating vote totals $\mathbf{v}_1 \rightarrow \mathbf{v}_2 \rightarrow \mathbf{v}_1 \rightarrow \cdots$, then every vote total is the same as the vote total 2 iterations later (2-cyclic) and every vote total is the same as the vote total 4 iterations later (4-cyclic). But if a game has the repeating vote totals $\mathbf{v}_1 \rightarrow \mathbf{v}_2 \rightarrow \mathbf{v}_3 \rightarrow \mathbf{v}_4 \rightarrow \mathbf{v}_1 \rightarrow \cdots$, then every vote total is the same as the vote total 4 iterations later (4-cyclic), but it is not true that every vote total is the same as the vote total 2 iterations later (2-cyclic). So a reasonable convention is to call a deterministic game j-cyclic, where j is the smallest cyclicity in that game's iterative equilibrium.

One subtlety worth noting in the proof of Theorem 7 is that, if it is possible for $\mathbf{v}_j$ to occur more than once within S, then the final step in the proof of the "only if" claim — asserting that the number of repetitions of the sequence is impossible — no longer produces any contradiction. Imagine for example that a deterministic election game is 2-cyclic. Without the constraint that j must be the *smallest* cyclicity of the game, we now know that a 2-cyclic game will also be 4-cyclic and 6-cyclic, so we might pick $j = 4$ and $k = 6$. 4 does not divide 6, but the game is both 2-cyclic and 6-cyclic, because in the construction of the set S above it would be two repetitions of the two-cycle, and it would line up correctly with the vote totals in the set T.

Theorem 7 provides a useful stepping stone for a valuable connection: using the cyclicity in the aggregate vote totals to identify constraints on the behaviour of individual electors. When N is large this can be extremely practically useful, since checking the properties of a $1 \times M$ vector of votes is easier than checking the behaviour of N individual electors. But the tight connection between vote totals and vote choices enforced by the idea of a deterministic election game makes these models theoretically much tidier as well. In Theorem 4 we saw that, in the context of the simple system of threshold pivotality, electors had to cycle with cyclicity 1 or k. The following theorem gives the slightly more complicated analog for the broader case of all deterministic election games.

Theorem 8. *A k-cyclic deterministic election game contains only electors who cycle between vote totals with cyclicities that divide k.*

Proof. Suppose that the aggregate vote totals in a deterministic election game are k-cyclic, but that there exists an elector whose set of equilibrium vote choices S has cyclicity h, where h neither divides k nor is a multiple of it. Let the model reach iterative equilibrium in iteration t, label the elector's vote choice in the following iteration v_{t+1}, and label w the number of iterations after t before the elector first responds with a different vote choice as in iteration t. Since $h > 1$, such an iteration must exist.

Since $\mathbf{v}_t = \mathbf{v}_{t+k}$, by Corollary 1 and Definition 1, $v_{t+1} = v_{t+nk+1} \;\; \forall n \in \mathbb{N}$. Now consider the remainder between the model's k-cycle and the elector's h-cycle: $\exists r \in \mathbb{N}, r > 0$ so that $k \equiv r \pmod{h}$. Then $wk \equiv wr \pmod{h}$. This means that iteration $t + wk$ and $t + wr$ are separated by h iterations. Because wk is a multiple of k, it prompts the same response in the elector as iteration t did, so $v_{t+1} = v_{t+wk+1}$. And because the elector's votes are h-cyclic, iterations wk and wr must prompt the same vote choice: $v_{t+wk+1} = v_{t+wr+1}$.

After h repetitions of the cycle, it follows again from Definition 1 that $v_{t+hwk+1} = v_{t+hwr+1}$. But because the elector's vote totals repeat in every multiple of h, $v_{t+w+1} = v_{hwr+1}$, and therefore $v_{t+w+1} \neq v_{t+1} \implies v_{t+hwr+1} \neq v_{t+1}$. This is a contradiction, so no such elector can exist.

□

There is a very straightforward interpretation of this result: a system composed of multiple types of cyclers should be understood as k-cyclic, where k is the least common multiple of all of the different cycles that exist in the system. Indeed, what we are really saying when we call a system k-cyclic is that k is the smallest natural number so that all of the electors in the system are k-cyclic. So when we use k to identify the smallest cycle length of a deterministic election model's iterative equilibrium, that value k is actually the least common multiple of all the cycles of the cycling electors in the system. That means that if k is prime, then any deterministic election game will have the same constraint on individual cycling that threshold pivotality always has (see Theorem 4):

Corollary 3. *A deterministic election game that is k-cyclic, where k is prime, can contain only electors whose shortest cycle has either exactly period k or period* 1.

Another quite simplifying theorem follows immediately from the definition of a deterministic election game (Definition 2):

Theorem 9. *In any deterministic election game, if the aggregate vote counts do not change from one iteration to the next, then no elector changes its vote choice.*

Theorem 9 contains the simplifying observation that an elector who is conducting pivotality calculations does not take into account how specific other electors are voting; all that matters is the total number of electors pursuing each strategy.

In the following section, I consider three ways of extending the work in this chapter to include other situations of interest.

Three brief extensions

In this section I briefly discuss how to incorporate abstention into deterministic election games, and how to connect the theory of deterministic election games to the theory of probabilistic election games. The latter topic includes suggesting some alternative notions of iterative equilibrium.

Nothing about deterministic election games is inconsistent with modeling voter abstention. We can use the example of threshold pivotality to see how it can be worked into the framework, and what sort of difference it makes. In threshold pivotality, if there is no τ-tie, then electors revert to their sincere preference orderings. A very natural alternative approach might be to instead have the electors conclude that, since there is one obvious election winner and nobody is close to catching them, it is not worth voting. This replaces the sincere voting round in threshold pivotality with an abstention round.

Imagine for a moment that all electors follow this logic. Then, after the abstention round, all parties are expected to receive 0 votes, since every elector is expected to abstain. This is a

τ-tie for any value of τ, and will prompt the sincere vote totals again. So in this model there would be two cases. If the sincere vote totals contain a τ-tie, then electors will vote for their most-preferred candidate within that τ-tie, which could produce an iterative equilibrium of cyclicity up to m. Otherwise, if the sincere vote totals do not contain a τ-tie, there will be a 2-cyclic iterative equilibrium which vacillates between the sincere vote totals and complete abstention.

Of course, this is an extreme example. A more realistic approach would be to give electors varied costs of voting, and set their expected utilities so that when there is no τ-tie, some electors abstain, while others simply revert to their sincere vote totals. Varying the number of electors who abstain would alter the τ-ties that are produced and could change the cyclicity of electoral districts. But overall this idea is quite close to the threshold pivotality game explored in this chapter, and would have similar properties; it just introduces another type of repeating iteration into the mix.

I have mentioned that probabilistic voting games are well-studied. The focus has traditionally been on the setting in which candidates have the capacity to manipulate their (perceived or real) position in some ideological or issue space, and the candidates have some distribution of probabilities over the number of voters who will support them at any given position. A tremendous amount has been written about these sorts of models; see for examples the thorough summary of classic work in this field by Coughlin (1992, p. 18-20).

Any amendments to the models in this chapter which incorporate probabilistic voting will pull it closer to this very large literature. But I stress four embedded differences that any such extensions should retain.

First, the setting of public polling means that electors have a clear signal about the state of the race. If a much stronger mechanism were introduced for electors to use polls to predict future states of the race, this would presumably mean sacrificing the iterative nature of the system, and would run into the problems of game theory models of elections. And in this setting, the only electors who need to vote randomly are electors who have no dominant pure strategy.

The second point is related: I have argued that, in the spirit of these models, it should never be the case that a very large number of electors vote probabilistically. A natural extension of this chapter would be to re-run the application of threshold pivotality to the 2019 Canadian election with no tie-breaking, and random voting instead. My expectation is that the results would not be very different. Analogously, the derivations regarding deterministic iterative election games can only be so wrong when there is some small population of probabilistic electors. For the most part I expect that they can be rephrased with bounds on the claims to account for how many probabilistic electors there are. The exceptions are situations where a small amount of probabilistic voting can chaotically throw off a result, like by randomly pushing a candidate out of a fragile tie and dramatically changing cycling behaviours. But it should be possible to account for such cases by instead phrasing the bounds as a combination of the difference in vote totals between the most competitive candidates compared to the number of probabilistic voters.

Third, there are three structural differences between the models in this chapter and classic probabilistic voting models that are worth pursuing in their own right. The literature on probabilistic voting is highly interested in the strategic behaviour of candidates, but in this dissertation I set aside that extremely well-studied topic in favour of the much more neglected area of voter strategy. The classic probabilistic voting literature also nearly always includes some notion of spatial utilities, whereas here I only assume preferences. And the social choice properties of these games, though of interest, have not been of central interest in this chapter; my concern is much more with how much they resemble empirical reality, and how we can use them to inform computer models that describe reality (or make conjectures about counterfactual tweaks to reality). Each of these motivations could be varied in the future, but the result would be a very different type of model than the ones discussed in this chapter.

Finally, probabilistic voting extensions that aim to have applications to discrete computer models must retain some idea of iterative equilibrium. However, we have shown something very damaging for the idea of iterative equilibria in any setting where there is random voting: iterative equilibria will never happen. But there is a simple fix. Definition 1 is an extremely restrictive idea, since it only considers vote totals to have reached iterative equilibrium if they repeat exactly in a regular pattern for all time. An extension to probabilistic voting could instead require that the party-by-party vote totals at a probabilistic iterative equilibrium repeat within only a small difference between them. The difficulty would be that, in many settings, probabilistic voting could be enough to bounce one nearly-repeating sequence of vote totals away to another nearly-repeating sequence of vote totals. One way of handling that might be to allow probabilistic iterative equilibria to be temporary, so that a model can enter and leave iterative equilibria.

With these thoughts I have attempted to outline a type of probabilistic voting that can be incorporated into these models to make them more realistic and more useful models of real elections, but without dramatically overhauling the setting of the chapter.

Conclusion

This chapter outlined a model of elections in which electors iteratively signal their intended vote choice, simultaneously update their intended vote choice based on that signal, signal their updated choices, and so on. These iterative election games have begun to appear in political science research, and some of their social choice properties have been studied under specific election rules. But models of elections need to produce election results: what are the vote counts on election day? And in these iterative election models, there is no obvious answer to that question.

This chapter began with the idea that we can say that a result has effectively been reached when vote choices begin to repeat. I called this idea iterative equilibrium, which consists of all of the vote totals that will continue to appear so long as the model runs. I showed that in deterministic iterative election games, in which one vote total produces some other specific vote total with certainty, there will always be an iterative equilibrium. The number of possible vote totals that can be in that iterative equilibrium also has an upper bound, although that bound is very large.

I illustrated iterative election games with an idea called threshold pivotality, in which electors check to see if multiple parties are within some threshold τ of the largest number of votes that any party is expected to receive. If there is such a near-tie, then electors vote for whichever party they most-prefer in that near-tie; otherwise, they simply vote expressively for their sincerely most-preferred party. I applied this model to real public opinion data from the 2019 Canadian election, and I showed that with reasonable values of τ threshold pivotality correctly assigns about 95% of the seats in parliament to the party that actually won that district. I also show that the number of parties involved in near-ties is consistent with major theories about party systems, with about half of districts in threshold pivotality reducing to exactly 2-way races. I also tried out five different ways of resolving iterative equilibria in the threshold pivotality application, and all five ways produced similar results.

Finally, I proved a series of constraints on the behaviour of electors in deterministic iterative election games. I showed that the vote totals in such a model at iterative equilibrium will only repeat in a certain number of iterations, or in multiples of that number. I also showed that electors in a deterministic election game can only have a cycle of vote choices with a length that divides the length of the cycle in the overall vote counts, and consequently that if the length of the cycle in the vote counts is prime, then electors whose vote choices cycle can only do so with that same length. In these games there is no chaos bubbling under a placid surface: the electors are nearly as orderly as the overall system. Finally, I showed that these claims apply to models which are deterministic just in the vote totals that arise during iterative equilibrium, and also that several of these findings are at least somewhat tolerant of a small amount of non-deterministic voting even in response to the vote totals at iterative equilibrium.

With the end of this chapter, threshold pivotality has served its purpose, though it will make a brief return in Chapter IV as a point of comparison for a more complicated model. But the general setting of this chapter is exactly the setting of the rest of the dissertation, and the simplifying theorems we have seen will be helpful tools for teasing results out of larger and more involved simulations.

CHAPTER III

A Direct Approach to Understanding Electoral System Reform

As the rumble of election reform has grown to a roar in many longstanding democracies, why have so few electoral systems actually changed? The problem is that politicians will only support electoral systems that they think will serve their interests (Ahmed, 2012, Boix, 1999, Cusack et al., 2007). But if party leaders evaluate alternative systems according to “a virtual iron law of political self-preservation”, how do they know how different electoral systems would affect them (Nagel, 1994b, p. 525)?

To understand why party leaders support or oppose a change in electoral systems, we need to know how they might answer the crucial question: what would happen to our party if we changed systems? A century of research has revealed a great deal about how electoral systems affect broad features of a country’s politics, and these broad tendencies can be used to make guesses about how different systems might affect different parties. But these tendencies miss the sort of concrete evidence needed to engage with the heart of election reform debates: how would a switch to a *specific* electoral system affect the *specific* parties that have the power to make it happen (Benoit, 2004, 2007)? Summarizing the state of our knowledge about the complex effects of electoral systems, Taagepera (2007, p. 36) warns policymakers that “no one can predict their actual workings — and you kid yourself if you think that you can”.

This gap forces party leaders to make a momentous policy decision under extreme uncertainty. Because we cannot confidently estimate how a specific party would fare under a specific electoral system, party leaders often simply do not know how different systems might affect them. This opens the door for politicians to take stances on electoral system reform that are inconsistent with their own self-interest, making it possible for parties to oppose system reforms that might actually help them. The implication is that many opportunities for reform might actually be compatible with the interests of enough party leaders for the reform to be possible, if only the system's likely effects were common knowledge.

In rare cases, reform efforts get far enough that we can observe politicians' actual reasoning when facing a possible election reform. The clearest example in recent years began when the Liberal Party of Canada, reeling from its worst election ever, promised that the 2015 election would be "the last federal election using" the country's Single-Member District Plurality (SMDP) system.[20] But even after the three leftist parties that supported reform captured more than two thirds of the seats in parliament, they could not agree on *which* alternative system to adopt, and no election reform was passed. Did these parties really conclude that SMDP is the best possible system for them? Or did they overlook a system that would benefit all three leftist parties at the expense of the right? In the United Kingdom, similar questions were raised by a 2011 referendum to replace the original Westminster electoral model with Alternative Voting, and by the vacillating it inspired among Labour Party leaders in particular. These cases make it hard to believe that party leaders confidently know exactly which systems would serve their interests.

Of course, sometimes electoral systems do change, but often they do not obey a simple narrative about politicians seizing an obvious opportunity to accrue power. In Alaska, an extremely convoluted new electoral system seems suspiciously well-designed to benefit incumbent senator Lisa Murkowski, and many observers allege that it was crafted by her supporters to maximize her re-election chances. And yet, real consensus on the system's effects has proved elusive: however it appears, commentators argue that such a complicated system "scrambles" elections and may produce "surprising results",[21] affecting campaigns "in unpredictable ways".[22] And in one of the great success stories in the history of electoral system change, the situation was somehow even more confused. The first move to replace the SMDP system that New Zealand had used for a century

---

[20] Tania Kohut. 2016. "What Trudeau said: A look back at Liberal promises on electoral reform", *Global News*.
[21] Kerry Picket. 2021. "New voting system scrambles Alaska 2022 Senate race", *Washington Examiner*.
[22] Ed Kilgore. 2021. "Will Ranked-Choice Voting Help or Hurt Lisa Murkowski?", *New York Intelligencer*.

and a half was a promise by the Labour Party leader David Lange, during a live debate in 1987, to hold a binding referendum on whether to adopt “a modified form of proportional representation”. That promise reversed his party’s position, which he later recalled was actually to “consider holding a referendum; in other words, avoid the issue”. Why reverse one of your party’s most important positions on live television? Recalling the event in 1992, Lange explained, “I thought this was our policy. It said so in black and white on a piece of paper in front of me.” But “there was a mistake in the notes”.[23]

Labour’s desperate struggle to retain SMDP after this enormous gaffe suggests that they expected a massive loss of power after the system change, so there is a final irony to the fact that Labour’s election result in 1987 is eerily similar to its election results in 2020: both times they turned 48% of the (constituency) vote into a narrow majority of seats.

When it comes to electoral systems, party leaders might not really know their own interests. And just as politicians “prefer those institutional formulas and procedures that can consolidate, reinforce, or increase their relative strength” Colomer (2004, p. 3), it turns out that voters in election reform referenda also tend to support whichever system they expect will help their favourite party (Heller, 2021, Riambau et al., 2021). So how can we make sense of when election reforms succeed or fail? What is needed is a method for understanding the range of seats that a particular party is likely to win under a specific alternative system. To that end, this chapter develops a computational formal model that combines game theory, previous election results, and public opinion data to estimate how a specific party would have performed in a specific election, if that election had instead occurred under some alternative electoral system. The model is falsifiable and accurate, and it can be applied to any fair election where there is data about voters’ party preferences.[24] Because people might not vote the same way under a different electoral system, the method explicitly models the strategic

---

[23] The admission by Lange, in a speech on July 15, 1992, was reported the next day in *The Press* by Peter Luke; when I quote Lange’s description of the event it is from this article, which I was made aware of by a footnote in Jackson (1993). For help identifying and acquiring these materials I am indebted to the Document Delivery staff at the University of Michigan and the University of Wisconsin-Madison, and especially to the Tuakiri | Identity floor at Tūranga Library, who had the relevant article in their microfiche collection.

[24] I do not model the creation of original legislative districts from scratch, so in this dissertation I address the special case of switches away from SMDP.

voting incentives under different systems, and the proportion of electors who engage in strategic behaviour can be set to any level.

Some approaches already exist for estimating party fortunes after an electoral system change. After discussing how the model builds on these approaches, I outline its structure and assumptions. I discuss how the model starts with voter preferences and produces simulated election results, and then I verify that the model makes accurate predictions about a historical case of electoral system change. In the following chapter I will apply the model to simulate proposed electoral system changes in Canada and Britain, and an electoral system change in Alaska that has already been made but has not yet been used in a statewide general election.

Modeling party fortunes: two problems and a solution

What approaches exist for estimating the strength of specific parties after an electoral system change? Even though established democracies have overhauled their electoral systems only a few dozen times in modern history (Colomer, 2004, p. 54), political scientists have very clear evidence on the broad tendencies of different electoral systems.

For a century, political scientists have studied electoral systems by observing countries before and after rare reforms, comparing countries that use different systems (Norris, 1997), or using experimental approaches (Blais et al., 2012, Blumenau et al., 2017) and natural experiments (Bechtel et al., 2016, Eggers, 2015, Sanz, 2017). The basic division is between what Colomer (2004, p. 3) calls "inclusive" and "exclusionary" electoral formulae: these determine how closely the distribution of seats should match the distribution of votes, and often provide a basic idea of who is accountable to whom (Loewen, 2017). The number of parties in a political system is also closely connected to the electoral system, following the famous observation by Duverger (1951, ch. 2.1) and its many refinements (Lijphart, 1994, ch. 5). The number of parties that tend to compete is now known with exceptional precision, as a function just of the size of a legislature and the magnitude of electoral districts (Shugart and Taagepera, 2017, ch. 4). And similarly consistent regularities connect electoral systems and voter turnout. It has long been suspected that turnout is higher under Proportional Representation (PR) than SMDP (Gosnell, 1930, Pollock, 1930), and a small positive difference has consistently been identified (Cox, 2015, Eggers, 2015). And perhaps most importantly, electoral systems are one of the crucial determinants of the share of women in a legislature, and can lessen or reinforce the exclusion of ethnic, racial, and religious minorities (Krook, 2018, Lublin and Bowler, 2018).

Given all of this information about the broad tendencies of electoral systems, perhaps the easiest way to respond to a party leader wondering whether system A will help them more than system B is to guess how the party system might change, and then make assumptions about where their party might fall in that altered party system. For example, if we expect the party system to shrink, then we might make informed guesses about which parties will no longer be competitive (Nagel, 1994a). But broad tendencies of electoral systems offer only vague guidance; as Cochrane (2017, p. 42) emphasizes, a question about the effects of electoral systems should be "an empirical question".

Some sophisticated methods have been developed to estimate the results of elections under alternative electoral systems. Gelman and King (1994) developed a statistical approach tailored to the context of American elections, while Ndegwa (1997) showed how the local vote results in a national election can be used to make guesses about peoples' preferences in different regions, which in turn can be used to estimate the results of an election under different election rules. For the specific case of "reforms that involve changes in the mean district magnitude" of a D'Hondt system, Flis et al. (2018, p. 37) derive a seat share prediction model that can use a variety of pre-election parameters as inputs, which Evci and Kaminski (2021) applied (along with proprietary software based on that model) to the case of a 2018 election rule change in Turkey. There are even models of how parties might adapt and respond to electoral system changes (Baker and Scheiner, 2007).

But the approach that makes up the bulk of estimations is much more simple and much more direct. Many researchers — understanding electoral systems as rules for turning votes into seats — have simply taken the votes that were cast in a previous election, applied a different electoral system to those votes, and treated the resulting number of seats as an estimate of election results under that system. For example, to estimate how American elections would look under a perfectly proportional system, a researcher might take the Democrats' and the Republicans' legislative vote shares in the 2020 election, and simply assign each party that share of legislative seats. This is the main way that researchers have made conjectures about how an electoral system might affect the specific parties and candidates in a country.

This simple approach has been used for decades, applied to countries in Africa, Asia, Europe, North America, and South America, and published in academic articles, policy briefs, and public commentaries. And indeed, its straightforwardness is quite appealing. Table 2 lists examples of the simple approach by author, year, whether the approach was used in a venue for public communication or in an academic publication, and what was being estimated. It also notes their framing of the assumption required to pursue this simple approach, in cases where they explicitly discussed that assumption. The examples in Table 2 were selected to emphasize the breadth of these applications around the world over the last few decades.

Context	Type	System	Citation
Canada, 1997	Academic	SMDP → PR	LeDuc (1999, p. 76)
	↪ “assuming that votes cast [after reform] were similar [...] Of course, such an assumption is not entirely reasonable. [Reform] is certain [...] to encourage strategic voting.”		
Chile, 1989	Academic	Varying PR rules	Navia (2003, ch. 4) and Zucco (2007)
	These assume just one alternative receives the same vote proportion.		
Hungary, 1990	Academic	SMDP → others	Benoit and Schiemann (2001, p. 170)
Kenya, 1992	Academic	SMDP → others	Ndegwa (1997, table 3)
Myanmar, 2015	Policy	SMDP → PR	Mun (2020)
Thailand, 2016	Public	MMM → MMA	Hicken and Pundit (2016)
	↪ “The biggest assumption [...] is that voters would behave the same way under MMA.”		
UK, 2015	Public	SMDP → PR	Blumenau and Hix (2015)
	↪ “our analysis makes the (completely unreasonable!) assumption that voters’ and parties’ behaviour would remain constant if the electoral system was changed.”		
UK, 2019	Public	SMDP → PR	Brandenburg (2019)
	↪ “Any exercise to model an alternative election outcome comes with the major caveat that we don’t know how a different electoral system would affect voting behaviour.”		

Table 2: Selected examples of the simple approach, recalculating past results under new rules. The type column specifies whether the analysis was part of an academic publication, a policy briefing for a government or non-governmental organization, or an article intended for the general public. The electoral systems are Single-Member District Plurality (SMDP), Proportional Representation (PR), Mixed-Member Majoritarian (MMM), and Mixed-Member Apportionment (MMA); in one case, just the type of PR is varied. Where the assumption that voters would not adjust their votes is explicitly discussed, an excerpt from that discussion is quoted.

The simple approach has become an extremely popular tool among political journalists in Westminster countries, as well as for advocacy organizations that seek to promote election reform. Because it is often used with the explicit goal of advocating for some system, among pieces that were published in the popular press I only include pieces written by academics. But I do include and identify pieces that were published by academics in popular press outlets, because their prevalence illustrates the fact that the simple approach is not just an academic concern: it is actively used in important public policy discussions.

It is also important to note that examples for sub-national legislatures were not included, though in the United Kingdom particularly there are rich examples of the simple approach being used to re-calculate the results of local elections.[25] Also excluded are the huge volume of claims about how American elections would look under PR. Because there are only two competitive parties, here the simple approach is often invoked almost unconsciously, with authors sometimes covering in less than a sentence the seat-vote gap between the two major parties; at the same time, because no other parties seriously compete in the country's elections, the direct approach cannot straightforwardly gauge how successful other parties would be. These are both such well-studied topics (on sub-national legislatures, and on the effects of systems in America) that there is at least one article that fits both criteria: Latner and Roach (2011).

Finally, while I have focused on selecting examples of the simple approach as applied directly to SMDP elections, I have also selected a few that vary the type of PR, to emphasize that the same tool has often been used in this type of investigation.[26]

Unfortunately, this widely used tool has one major limitation that makes it frequently impossible, and even when it is possible, it requires a major assumption that is known to be wrong. In this section I discuss both of these limitations, and describe how they can be overcome, to make the simple approach more defensible but no less direct.

The biggest problem with the simple approach is that it can only be applied to some alternative electoral systems. Some alternatives to SMDP, like Ranked-Choice Voting or Mixed-Member Proportional, ask voters for more information than just their first choice. Unless we know electors' preferences beyond their top choice, there is no formula we can apply to the results of a single-vote election in order to estimate the results of a hypothetical election in which voters are asked to make

---

[25] See for example Richard Ault, "How Buckinghamshire Council would have looked under proportional representation", *Buckinghamshire Live*, 2021.

[26] I am grateful to Lucia Motolinia for pointing this out to me.

multiple choices. The second serious problem is that the method treats electors as rigid constants, unresponsive to the context in which they vote. If the electoral system suddenly changed, would some voters not change how they vote?

For example, consider the problem faced by Vowles et al. (1995, ch. 11) when they devoted a chapter to predicting how the switch from SMDP to Mixed-Member Proportional would affect New Zealand's parties. In 1993, voters were only asked to choose a local representative. But in 1996, they would be asked to vote for a local representative *and* a political party. With no obvious way to estimate the party votes based just on constituency votes, Vowles et al. (1995, ch. 11) could not resort to the simple approach, and instead set previous election results aside and based their predictions only on survey experiments. LeDuc (1999) encountered a similar challenge when using real votes for the Canadian House of Commons in 1997 to estimate the seat distribution of a hypothetical Mixed-Member Proportional Senate, but he answered it by simply assuming that everyone would vote for the same party as they voted for in their local constituency.

But the assumption that electors would have voted the same way under a different electoral system has only ever been an assumption of convenience. LeDuc (1999, p. 76) called the assumption that voters would act exactly the same under Mixed-Member Proportional as they had under SMDP "not entirely reasonable". Blumenau and Hix (2015) hit a strikingly similar note nearly two decades later when re-calculating a British election under more proportional rules: to them, the assumption that voters would not adjust to the new system was "completely unreasonable!"

The idea that electoral systems act on the psychology of voters, constraining their idea of the number of parties that they can viably vote for, was identified by Duverger (1951) and has been a mainstay of the study of electoral systems and strategic voting since then (Lijphart, 1994, p. 70). And research designs that aim at identifying the psychological effect of electoral systems on the vote shares of parties, as a distinct quantity from any mechanical way that electoral systems might directly constrain the number of parties that can be present in the legislature, have found that voters do change their vote choice as a result of the psychological effects of different electoral systems (Blais and Carty, 1991, Fiva and Folke, 2014).

So there are two gaps that need to be bridged before we have a method that successfully estimates party fortunes under alternative electoral systems. First, this method will need to account for electors' full preference orderings, so that something can be said about votes they might cast *beyond* their simple top choice. Second, it must be possible for electors to adjust their vote choice to account for the new system. But how can we determine electors' preferences beyond the first party, and how can we use those preferences to estimate vote choices under some alternative electoral system?

Brandenburg (2019) described the crucial bridge, while re-calculating the British 2019 election under PR: to Brandenburg, "the most obvious flaw" with the simple approach is that under a different system, there are "very different needs for tactical voting". Indeed, only strategic voters should be systematically affected by changes in electoral systems, because non-strategic voters always choose their sincere top preference. An electoral system is just a modification of the probability that a party will win a seat when it receives a certain number of votes; in the language of strategic voting, a new system changes the expected utility of the possible vote choices. That is irrelevant to a sincere elector's vote choice.[27] Of course, this does not mean that everyone votes strategically: some portion of the population will vote sincerely under any system, and another group might change their vote if the electoral system changes.

Three pieces are therefore required, in order to build on the simple method to cover systems which ask for more than one choice, and allow electors to respond to the system change. First, we need to estimate every elector's full preference ordering over the parties. Second, we need a strategic decision rule that electors can use to turn those preferences into votes. And third, we need to be able to tune the level of strategic voting so that it is empirically reasonable.[28]

---

[27] By strategic voting, I mean any process in which electors consult the behaviour of other electors to try to cast an optimal vote. In this section I will also identify two caveats to the idea that only strategic electors should switch votes.

[28] In developing this method, I only consider switches away from SMDP. Colomer (2004, p. 54) found that 76% of modern electoral system changes were away from majoritarian systems and towards more proportional systems, and the approximately 60 democracies that use SMDP are unignorable: they include the world's largest democracy, the world's biggest economy, the largest democracy in Africa and many of Asia's largest, and the original Westminster country. Vasselai (2021b) has studied how to study election reforms in which original districts are drawn, as in the American redistricting literature (McGhee, 2020).

Model structure

The direct approach to understanding the effects of electoral system changes has two main components: to model electors' expected votes under a new system, we need to know their preferences over the parties, and we need to know how they turn those preferences into a vote choice. We also need to compare their vote choices under different electoral systems. First, the model generates a list of preferences which represents the preferences that real people hold in a certain electoral district. Counting the electors' most-preferred parties gives the results of an election in which everyone votes sincerely. Second, electors use the results of the previous election in their district to iteratively select their best vote choice. Counting votes at this stage simulates an election under the real SMDP electoral system, so it provides an empirical check on the model: the results should be close to the real election results. Finally, an electoral system change is introduced by modifying how electors judge the competitiveness of the parties. Figure 5 summarizes the model's structure.

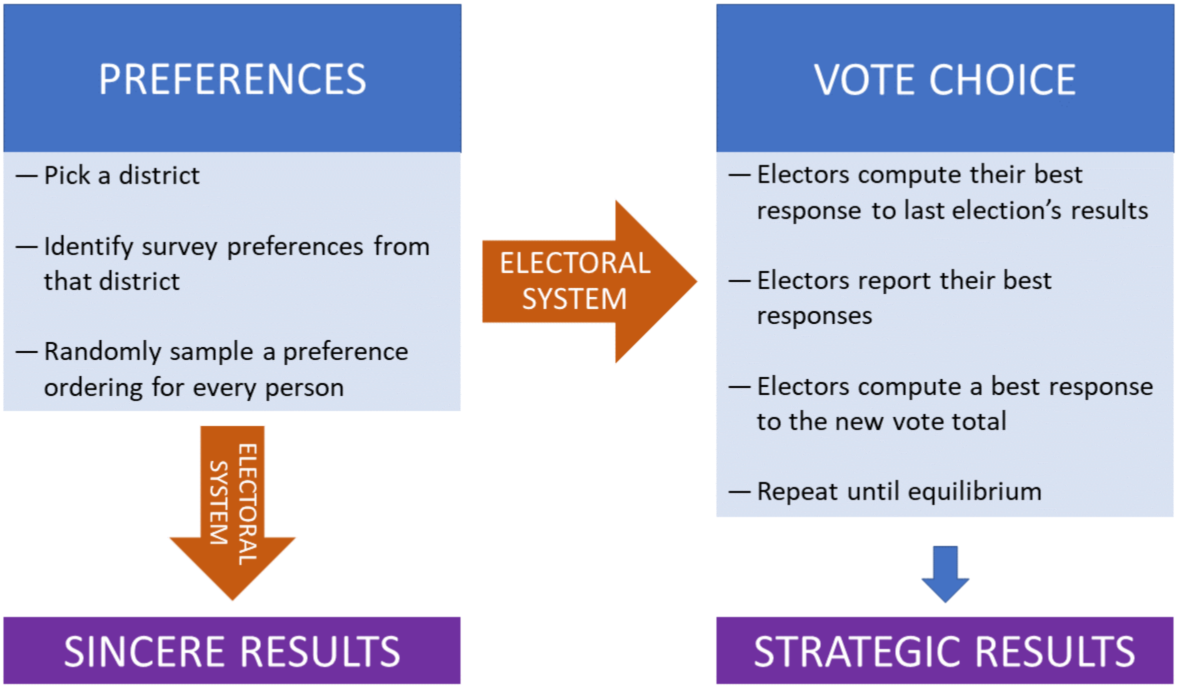

Figure 5: The two stages of the model, and the results that each stage produces.

The third structural point is that the method (when applied district-by-district to a national election) only considers the local competitiveness of the candidates, not their national standing. You could imagine, for reasons similar to the ones that Kedar (2014) discusses, that an elector should take the parties' national competitiveness into account when calculating the expected utility of their available vote choices. For example, suppose that I am a Canadian voter who strongly dislikes the Liberal Party, and my district is a close race between the Liberal Party and the New Democratic Party, with the Conservatives a close third. But nationally, suppose that the Liberals and the Conservative parties are expected to win about the same number of seats, while in the national seat count the New Democratic Party will be a distant third. A strategic voting calculation that includes only the standings of the parties in my district might suggest that I should vote for the NDP. But a more complete picture would include the information that if I vote for the Conservatives, I have a chance at breaking a national tie, and not just a local tie. I do not consider these details, but this would indeed be a natural extension of the model.

Generating preferences

In Chapter II, I described and used a method for generating preferences.[29] I applied the idea described in that chapter to the 2019 Canadian election, but similar election studies have asked the same battery of questions to a similar number of respondents in recent elections in Canada, Britain, and New Zealand. These data are all particularly high-quality, and will be the main source of preferences throughout the dissertation. But in Chapter IV I will also subject the model to much harsher conditions. Polling is scarce in Alaska, and I will construct preference orderings just from publicly released cross-tabs of Likert scales for major Alaskan political figures broken down by partisan identification. First I generate electors with specific partisan identifications, in the proportion that they exist in the real Alaskan electorate. Then I sample ratings of the major political

---

[29] The passage in question is the one on page 27 that begins "During the 2019 election, the Canadian National Election Study asked tens of thousands of respondents to rate the major political parties on a thermometer scale" and concludes with "quite different from the voters in that district a few years ago".

figures in the proportion that they are reported in survey data. This gives a preference ordering for every elector.

I assess these different ways of generating preference orderings with empirical tests. Even without accounting for any strategic voting, and instead just counting the imputed sincere preferences, produces results that are reasonably close to real election results. Much more challenging is the model's second stage: how do electors turn these preferences into votes?

Voting logic and system changes

There is no single correct way to model strategic voting, as half a century of empirical paradoxes and red-hot debate have shown. But the purpose of this chapter is not to advance a new model of pivotal voting; it is simply a requirement for one stage of the model that electors judge how their vote in an election might affect the results of the election. Crucially, all that matters is that the strategic voting model is good enough to capture how voting patterns might change between SMDP and some other system.

I build the elector logic on game theoretic models of elections, adapting a classical pivotal voting approach to use in the computer. There are two reasons that the model uses pivotal voting logic, even though previous computer models of elections have favoured heuristic voting (Bendor et al., 2011, Laver and Sergenti, 2012). First, since an electoral system change is a modification in the probability that a party gets a certain number of seats when it gets a certain number of votes, the natural way for electors to incorporate such a change is to actually calculate the probability that their vote helps to elect a candidate. Second, if electors have a specific pivotal voting rule under one electoral system, then there can be a unique correct way to derive their voting rule under another electoral system. The aim of this dissertation is to study what happens to election results when *only* the electoral system changes, so it is essential that electors have the same rules for voting under both systems; with heuristics, it is not straightforward to pick one heuristic under SMDP and an exactly matching heuristic under, say, Ranked-Choice Voting.

Importantly, in a return to the threshold pivotality heuristic later in the analysis, I will demonstrate that the particular details of the strategic voting rule are not at all important. All that matters is that some electors are inclined, more or less strongly, to support more competitive candidates. A future section compares the voting rules that I use in the main model (because of their unique extension to other electoral systems) to that dramatically simpler heuristic, and finds that the results are quite similar.

With that caveat in mind, the model's strategic voting logic is as follows. There are two situations in which electors can be pivotal in a single-vote election: the case in which they create a first-place tie, and the mutually exclusive case in which they break a first-place tie. The probability p of being pivotal in causing some party i to win the seat over another party j is the probability that the vote totals of the parties, v_i and v_j, are exactly equal, and that i and j both have more votes than any other party. So an elector's expected utility for casting a vote for party i is

$$u(i) = \sum_{\substack{j=1 \\ i \neq j}}^{M} \left(\mathbb{P}(v_i = v_j \wedge v_j > v_k) \cdot (u_i - u_j) + \frac{1}{2} \cdot \mathbb{P}(v_i = v_{j+1} \wedge v_j > v_k) \cdot (u_i - u_j) \right) \tag{1}$$

for all other parties k, where u_i and u_j are the utility obtained from the victory of i or j respectively and v_i and v_j are their expected vote totals, and M is the number of parties in the election. The factor of $\frac{1}{2}$ in the second term assumes that ties are decided by a coin toss, so that creating a tie is half as desirable as breaking one.[30]

As discussed in Mebane et al. (2019), the computer differs from the asymptotic arguments common in game theoretic models of strategic voting, because in the computer there is only a finite population of electors. This is a virtue: there are finitely many electors in the real world too. Mebane et al. (2019) set up a Poisson voting game among a finite number of electors in the framework of Myerson (2000), which means that the difference between expected vote totals is a difference between Poisson random variables, which in turn is Skellam distributed. And this expression is not just valid for a single-vote election; the same logic holds for any Single Non-Transferable Vote

---

[30] This expression ignores ties between more than two parties, which I follow previous authors in taking to be ignorable in elections with realistic turnout.

system (Cox, 1994, Myerson, 1998).

But it is essential that we can simulate elections where electors make more than one choice. So next I derive the utility that an elector can expect to obtain from casting a multi-choice ballot in a Borda count election: that is, an elector's top choice gets a certain number of points, the second choice gets one less point than that, and so on, and the winner is whichever candidate has the most points. Then a ballot of length l has the potential to boost any one of the l parties into or out of a first-place tie. So Borda count elections can be assembled out of the same logic as the SMDP model: electors simply perform exponentially more comparisons to evaluate the expected utility of every possible ballot (all of the length-l permutations of the parties), and choose the most favourable bundle of votes. To find the utility of casting a particular ballot, start by assuming that all ballots must be complete; later I will return to that assumption, and start with an expression of the value of giving r points to party i. Once we have that expression, we can compute the value of all possible profiles of votes, since giving r points to a given party comes with constraints about how many points can be given to each other party on the same ballot. In the spirit of the single-vote setup, we split the problem first into the probability of bringing party i into first place replacing another generic party j that is already in first place, or alternatively bringing party i into a tie with the party or set of parties that are currently in first place. Denote the probability of bringing party i into a clear first place as $P(i \uparrow)$ and the probability of bringing i into a first place tie as $P(i \rightarrow)$. I assume no multi-way ties, and begin by writing out the set of all cases in which we create a first-place winner by assigning 1 point, then consider assigning more than 1 point. The probability of bringing i into first place over the current first place party j by giving 1 point to party i can be expressed as

$$P(i \uparrow j | i_1) = P(v_i = v_j \wedge v_j > v_k \ \ \forall k \neq i, j)$$

where k is any other party label in the set of all parties $\{1, 2, \ldots, M\}$, and v_i is the vote total of party i. Similarly,

$$P(i \rightarrow j|i_1) = P(v_i + 1 = v_j \wedge v_j > v_k \;\; \forall k \neq j)$$

First, we seek a full expression for the value of assigning r points to party i, where $r \in \mathbb{N}$ such that $1 \leq r \leq M$, with M the number of parties (and therefore also the number of votes that can be assigned on a ranked choice ballot). Then,

$$P(i \uparrow j|r_i) = P(((v_j - v_i = 1) \vee (v_j - v_i = 2) \vee \ldots \vee (v_j - v_i = r_i - 1)) \wedge (v_j > v_k \;\; \forall k \neq j))$$

Because these are mutually exclusive events,

$$\begin{aligned}
P(i \uparrow j|r_i) =& P(v_j - v_i = 1 \wedge v_j > v_k \;\; \forall k \neq j) \\
&+ P(v_j - v_i = 2 \wedge v_j > v_k \;\; \forall k \neq j) \\
&\vdots \\
&+ P(v_j - v_i = r_i - 1 \wedge v_j > v_k \;\; \forall k \neq j)
\end{aligned}$$

Or, more compactly,

$$P(i \uparrow j|r_i) = \sum_{\rho=1}^{r-1} P(v_j - v_i = \rho \wedge v_j > v_k \;\; \forall k \neq j)$$

Correspondingly, there is only one case in which assigning r points to i creates a first-place tie:

$$P(i \rightarrow j|r_i) = P(v_j - v_i = r_i \wedge v_i > v_k \;\; \forall k \neq j)$$

But this is only the expression for one alternative party j. Still assuming that no multi-way ties can occur, we need to consider the potential for ties with i and every possible party j. There are $M - 1$

parties j (since only $j \neq i$). The situations in which each possible competitor j is the first-place party are all mutually exclusive events. So, we may add those events as follows:

$$P(i \uparrow j|r_i) = \sum_{\substack{j=1 \\ j\neq i}}^{M} \sum_{\rho=1}^{r-1} P(v_j - v_i = \rho \wedge v_j > v_k \ \ \forall k \neq i)$$

Finally, to get an expression for the utility of these outcomes, notice that the expected utility of causing party i to defeat party j is simply the utility obtained from the victory of party i over party j, denote it $u_i - u_j$, and to produce expected utility we need only multiply this by each term in the summation. Also denote the probability that i will win a runoff against j as $P_f(i, j)$ (a simple choice is $P_f(i, j) = \frac{1}{2}$, to represent systems in which ties are decided by a coin flip). Then, with $U[r_i]$ the expected utility of assigning r points to party i, we find

$$U[r_i] = P(i \uparrow j|r_i)(u_i - u_j) + P(i \rightarrow j|r_i)(u_i - u_j)(P_f(i, j))$$

$$\begin{aligned} U[r_i] =& \sum_{\substack{j=1 \\ j\neq i}}^{M} \sum_{\rho=1}^{r-1} P(v_j - v_i = \rho \wedge v_j > v_k \ \ \forall k \neq i) \cdot (u_i - u_j) \\ &+ \sum_{\substack{j=1 \\ j\neq i}}^{M} P(v_j - v_i = r_i \wedge v_i > v_k \ \ \forall k \neq j) \cdot (u_i - u_j) \cdot P_f(i, j) \end{aligned}$$

These are the full expected utilities for assigning r points to party i when we assume the non-existence of multi-way ties. What, then, is the expected utility of a given point allocation on a full ballot? Let U be a function which maps from an allocation of points r to corresponding parties ρ_i, onto the set of real numbers. To clarify what I mean by allocation, say that my options are to assign $S = \{1, 2, \ldots, l\}$ points to a party such that I must choose each element in S and apply it to a party exactly once, where l is the length of the ballot, $l \leq M$. The utility that I get for picking party g_i, labelled arbitrarily such that $i \in \mathbb{N} : 1 \leq i \leq M$, and giving it r points, is, calling it the Ballot Profile β,

$$U(\beta) = U(1, g_1) + U(2, g_2) + \cdots + U(l, g_l)$$

$$U(\beta) = \sum_{i=1}^{l} U(i, g_i)$$

We can decompose each of these terms individually using the expression obtained above for the utility of assigning r points to party i.

$$\begin{aligned} U(\beta) =& \sum_{i=1}^{l} \Bigg(\sum_{\substack{j=1 \\ j \neq i}}^{M} \sum_{\rho=1}^{i-1} P(v_j - v_i = \rho \wedge v_j > v_k \ \ \forall k \neq i) \cdot (u_i - u_j) \\ &+ \sum_{\substack{j=1 \\ j \neq i}}^{M} P(v_j - v_i = i \wedge v_i > v_k \ \ \forall k \neq j) \cdot (u_i - u_j) \cdot P_f(i, j) \Bigg) \end{aligned}$$

And now this statement contains a redundant sum which we can collapse, yielding a full expression of a voter's expected utility for casting a Borda count ballot:

$$\begin{aligned} U(\beta) =& \sum_{i=1}^{l} \sum_{\substack{j=1 \\ j \neq i}}^{M} \Bigg(\sum_{\rho=1}^{i-1} \big(\mathbb{P}(v_j - v_i = \rho - 1 \wedge v_j > v_k \ \ \forall k \neq i) \cdot (u_i - u_j) \big) \\ &+ \frac{1}{2} \cdot \mathbb{P}(v_j - v_i = i \wedge v_i > v_k \ \ \forall k \neq j) \cdot (u_i - u_j) \Bigg) \end{aligned} \tag{2}$$

Note that at no point does an elector estimate the probability of a candidate occupying a particular ballot position on some elector's ballot; an elector only needs to estimate the probability that the difference between vote totals of multiple candidates will be of a certain magnitude, regardless of how that vote total was obtained in terms of individual multi-option ballots. That problem has already been addressed in the context of approval voting (Kloeckner, 2020, Myerson, 1998). I have mentioned that the assumption that all ballots must be complete is nonessential. Could that be feasible? In an election with M parties and ballot length l, with $M > l$, and if all ballots are complete, then the number of possible complete ballots is the number n_c of l-permutations of the M parties:

$$n_c = \frac{M!}{(M-l)!}$$

Consider the example of Québec, a polity with 5 major parties (and the largest party system I consider in this dissertation). If all ballots are complete, and picking arbitrarily a ballot length of 3, the number of ballots that electors have to consider is 60. Now imagine if instead electors could leave any number of spaces blank on their ballots. Then the number n_b of possible ballots with any number of spaces left blank is

$$n_b = \sum_{i=0}^{l} \left[\binom{l}{i} \frac{M!}{(M-(l-i))!} \right]$$

In Québec, with ballot length 3, $n_b = 136$. This is only double the number of possible complete ballots, so since electors' computations are independent but model runs are parallel, it can be computed in somewhat more than half the time. This is feasible. However, what could we expect to learn by modeling it? Leaving a ballot location blank will always yield utility 0, so it will only be strategically viable when it replaces a vote that would be forced to be spent on a vote with negative expected utility. Unless ballot length is very large compared to the number of parties, this should be a rare occurrence.

Iterative voting, model scope, and model structure

This section describes iterative voting in the model, identifies some limitations on the model's scope, and finally discusses three embedded assumptions in the model's structure.

Once an elector computes the expected utility of each vote choice, though, it may wish to update its vote. A fundamental feature of the model is that electors update their vote choices iteratively; if an elector wishes to update its vote choice, it is not at equilibrium yet, and it is not clear how to obtain meaning from out-of-equilibrium votes. So while electors calculate their initial expected utilities based on the real results of a previous election, one iteration later they instead judge expected utilities based on the total number of supporters that each party has in the simulated electorate. This is analogous to the process of reading opinion polls during an election cycle to judge which parties are locally competitive; before the first election poll is conducted, a reasonable way of judging the competitiveness of the parties is to recall the results of the last election.

With full expressions for electors' expected utility from casting a vote choice, can we now derive a Nash Equilibrium for the voting game? This is an example of the situations that we studied in Chapter II: In this case, too, it is not so straightforward. In fact, because the expected difference between two Poisson variables with given parameter is not itself a random variable but is rather simply a real number, an elector will always map a given vote total onto just one pivotal probability with certainty under both Equation 1 and Equation 2. That means that a given vote choice always prompts exactly the same expected strategic utility for an elector with a given sincere preference ordering, so a given vote total will always cause them to vote in the same way. That means that the model outlined above is a deterministic iterative election model, and has all of the properties derived for that type of model in Chapter II.

In Chapter II we saw that it is possible for deterministic iterative election models with simultaneous updating to produce fixed best response equilibria, but that this does not always happen. A simple counter-example is shown in the first chapter of this dissertation, in the study of the sorts of iterative equilibria that are possible in voting games in which electors simultaneously update based on information about the current expected vote count. There I show that the ubiquitous situation in these sorts of models is that the best response to vote totals is often a repeating cycle. So, if electors view a vector of expected votes $\mathbf{v}_1$, and the mutual best response of all electors to that vector is some other vector $\mathbf{v}_2$, it will often happen that the best response to $\mathbf{v}_2$ is the vector $\mathbf{v}_1$ again, so that the expected result of the election will bounce between $\mathbf{v}_1$ and $\mathbf{v}_2$ without end. However, the end state of these models is always either a fixed vote choice or a finite cycle; there cannot be an equilibrium that wanders for infinitely many iterations, and indeed in Chapter II we found an upper bound for the number of iterations that a cycle can last for, as a function of the number of candidates and electors.

In the models in this chapter and the next one, it is important that elections produce winners, not an endlessly alternating cycle of different winners. So when reporting one result within a district, I first ensure that the model is at either a fixed equilibrium or an iterative equilibrium. In either case, the results represent a mutual best response to the expected vote count in the previous iteration. Then, if the equilibrium is a moving equilibrium, I arbitrarily pick an iteration in which to end the model and decide the winner. In the following models I always pick a default value of 0 iterations. This is the "election day" stopping rule outlined in Chapter II, which was motivated by the idea that real elections are held on some particular day, even though people may still be changing their minds about how to vote. Averaging results over many model runs will then feature winners in districts that have moving equilibria exactly in proportion to how often each candidate has the highest vote count at the start of iterative equilibrium.

As I argued in Chapter II, these complications are a minor cost for a powerful tool. The method introduced in this chapter should be understood as simply a formal model that is able to generate examples based on real data, taking advantage of the potential of computer models to generate and solve examples of game theoretic interactions in political phenomena (Siegel, 2018). The use of computers has two major benefits. First, the computer can sample thousands of actual preference orderings to use in the model, which connects the game theoretic framework to actual empirical information about the election. This allows the model to be fitted to the situation in a particular country at a particular time, and then tested against reality. Second, the computer can explicitly compute examples of elections, which means that it gives specific quantitative estimates rather than just general tendencies.

Some discussion of the model scope is also warranted. If to try to predict the effects of a new electoral system is to "kid yourself" (Taagepera, 2007), the literature holds much harsher warnings still for anyone who attempts to predict the long-term *status quo* under a new electoral system.

For Polga-Hecimovich and Siavelis (2015, p. 272-273), it is actually "a fundamental limitation of electoral design" that "although institutions can be manipulated to benefit one group in the short-term, longer-term behavior is much more unpredictable". Enough uncertainty abounds when political actors try to predict what will happen in the first election after an electoral system change. But a much higher wall stands between us and simulating the long-term future of elections under a different system. So far, I only offer a model of the first case: what will happen *shortly* after the adoption of a new electoral system, in just the first few elections after the change is made?

I have offered no model of how larger institutions might rearrange themselves after an electoral system change, such as whether or not new parties will form. But I argue that, empirically, this does not seem to be a particularly serious limitation on its ability to estimate how parties might perform in their first election under a new electoral system. It is rare for a party that did not have a large presence in the legislature to immediately obtain a very large seat share under a new electoral system. For the case that the bulk of this chapter was devoted to analyzing, Stephenson (2016) argues that, because "so many parties already exist in Canada, the question is not so much whether more options will crop up under a new PR system, but whether voters will shift their support to some of the parties that are already available under different electoral rules." In the Alaskan example too, one of the most difficult parts of assessing the effects of the new electoral system will prove to be choosing from the perennial contenders enough options for the second-round threshold of 4 to actually affect a simulated election result. So it seems plausible that Stephenson's analysis would apply there too, and for the same reasons. So while the model is not necessarily limited to surprising or unexpected system changes, it should be understood as a model of what happens shortly after the change, and the model as it exists should be applied only to cases where that is an interesting question.

There are three reasons that this limitation was necessary, but not too limiting.

First, the model is an attempt to clarify the actual results that party leaders should expect for their political organization after an attempt to change a party system. In what situations will the model give results that do not reflect the real payoffs to the parties? Two things would need to be true:

party leaders would need to be focused on the long-term results for their parties, and those long-term results would have to be different from the short-term results. Whether and in what conditions party leaders are focused on the long-term fates of their parties rather than just the near-term fates of their parties is not a question that has a simple answer, but surely it is not terribly limiting to argue that the model should specifically apply in those cases where party leaders are at least slightly concerned with what happens in the next election, rather than in the far future. The model results should be of interest to leaders who are *at all* interested in the near-future of their parties after an electoral system switch, which should nearly always be the case, especially since that includes the re-election prospects of their caucus members. It also seems apparent that for minor parties who are largely locked out of parliament under SMDP (take for example Canada's Green Party), long-term trends after a system change can hardly be a major concern, since even unfavourable conditions under a more proportional ruleset should usually be expected to break in their favour more than any SMDP election — for some parties, there is nowhere to go but up. And for the leaders of already sizeable parties who are interested in their party's long-term future, the model still provides a signal about what a post-change election might look like. So for example, even if the New Zealand Labour party expected that the first election under MMP would be more favourable than subsequent elections, the size of the change in that first election should matter — the fact that they only lost a few seats in the first MMP election might have put to rest concerns that they would be permanently unable to recover any majority status, concerns that were permanently banished in 2020.

But the question of the long-term fates of parties under various electoral systems remains interesting and worth addressing in its own right. So a second point about the model as it stands is that it is indeed a logistical first step towards answering that question. A model of the long-term fates of parties would almost by necessity include a model of their short-term prospects, so one can envision that as a later, separate addition to the model outlined here, although the questions it addresses and some aspects of its design would need to be very different.

Third, however, in this dissertation I highly prioritize falsifiability. It is crucial that we can perform certain validation checks, like applying the SMDP model to an SMDP election, or applying the model to the first election after an electoral system change. It is obvious how the model could succeed or fail at these verifications. And if either Canada or Britain does eventually change its electoral system, the model with either succeed or fail at predicting the results of those elections. Already it makes certain observations about the new Alaskan electoral system that can be checked in 2022. But how can we falsify a model of what happens 30 years after an electoral system change? It is hard enough, with thousands of polls and decades of sophisticated research, to predict an American presidential election the day before it happens. Surely a model of the long-term fate of parties under various electoral systems cannot make a point estimate of how many seats a particular party will hold N years in advance. But if the expectations are distributional (say, that a model in 1996 would have predicted that by 2020 it would be highly unlikely but not totally impossible for Labour to form a majority government), how can we hope to test them when dealing with one of the most notoriously data-poor questions in political science, where we have just a few dozen examples in the whole history of modern democracies?

It is for these reasons that the model is tightly focused on how parties should expect to perform in the near future following a change from SMDP to some other electoral system.

Finally, I close this section with three important structural points about the model itself.

First, there are two caveats to the claim that only strategic voters will change their vote between electoral systems. The first caveat is that one can imagine electors with voting rules like "under SMDP, I vote for party X, but under PR, I vote for party Y", even without any strategic motivation for doing so. But there is no reason to think that there are enough electors with perverse voting rules like this that it represents a *systematic* connection between electoral systems and vote choice. The second caveat is that the utility that electors expect to obtain from the victory of different parties under different electoral systems can depend on something other than their winning probabilities: as Kedar (2014) argues, the utility that an elector obtains from a party winning depends on what exactly "winning" means, which is partly determined by the rules of the political system. This is a separate type of strategic thinking that I do not consider in this dissertation; here I will take voters' preferences over the parties to be fixed.

The second structural point is that I ignore abstention. The goal of this method is to produce empirically grounded estimates of what might actually happen after an electoral system change. Famously, the inclusion of abstention in voting models, except with very specific and often *post hoc* parameters, causes them to depart hugely from reality (Wuffle, 1984). But ignoring abstention also comes with a substantive complication, since it is one variable that is famously affected by electoral systems. However, as summarized in the introduction, empirical estimates of how much turnout varies across electoral systems have tended downwards since they were first estimated to be on the order of tens of percentages about 100 years ago, and turnout is now believed to be on the order of a few percentage points higher in some systems than others. Of course, the vote decision of a few percentage points of a voting population can be substantively important – bearing in mind that in a large polity like the US or India, that might be hundreds of thousands of people. But turnout *itself* should not be expected to dramatically alter election outcomes.

To see why, take the most extreme case, where an extra several percent of the voting population would abstain in SMDP but would vote entirely for one party under PR. That is only the equivalent of adding a few percentage points to the total of whatever party you expect would benefit. So if turnout is expected to go up by 3% after a switch in systems, then in the most extreme case, a party that gets, say, 18% in my model would get perhaps 23% or so if turnout were accounted for (supposing a small majority of the eligible population votes). Note that this turnout differential is the only effect that my model is missing: I do explicitly model people who vote under SMDP but would vote differently under PR, so we are only missing the people who abstain under SMDP but would vote under PR. So a reader who believes that the model is missing people under PR who would vote a certain way under SMDP is encouraged to add or subtract a few percentage points from the party vote shares reported in this chapter and the next one.

The next section concludes Chapter III with an empirical check of the model, applying it to the real electoral system change that happened in New Zealand in 1996. In the interest of ensuring that the structure of the model is completely clear, though, the reader may first wish to consult the Appendix, which includes a step-by-step walkthrough of the model.

Empirical model test

Before drawing substantive conclusions, New Zealand's electoral system change provides an opportunity to test the model. Can the model, using a national election study conducted between New Zealand's last SMDP election and its first Mixed-Member Proportional election, accurately estimate the result of the first election under the new system?

In New Zealand, since the election reform in 1996, electors cast both a local constituency vote and a nationwide party vote. The constituency vote is used to elect one local representative, by plurality. Then seats are distributed to each party in a way that ensures that their seat share in parliament resembles their share of the nationwide party vote. To apply the sincere model in this situation, I first take the preference orderings reported in the 1996 New Zealand Pre-Election Study (Vowles et al., 1996). For each person who voted for a given party in a given district in the 1993 election, the model generates an elector with one of the orderings reported in the 1996 pre-election study, subject to the constraint that they do not hold that party in last place. Then the electors report their sincere top choices. The nationwide top preferences form the sincere party vote, and top preferences by district form the sincere constituency vote.

To implement a strategic model, voters use equation 1 to compute their best response vote choice within their local constituency, iteratively updating until they reach an equilibrium as described in Figure 5. I take the standard assumption that strategic voting in the nationwide proportional party vote is negligible, so the party votes are sincere (Karp et al., 2002).

The severity of the data constraints make this test more challenging. Only a few thousand complete preference orderings were reported in the election study, so in some electoral districts there are only about half a dozen preferences to sample from. There was also substantial redistricting between 1993 and 1996. To obtain the results of the previous election in each new district, I categorized which old districts were adjusted or abolished to make room for a new district, and then I averaged the results in those old districts to seed the election results in the new district. The model's estimates are shown Figure 6.

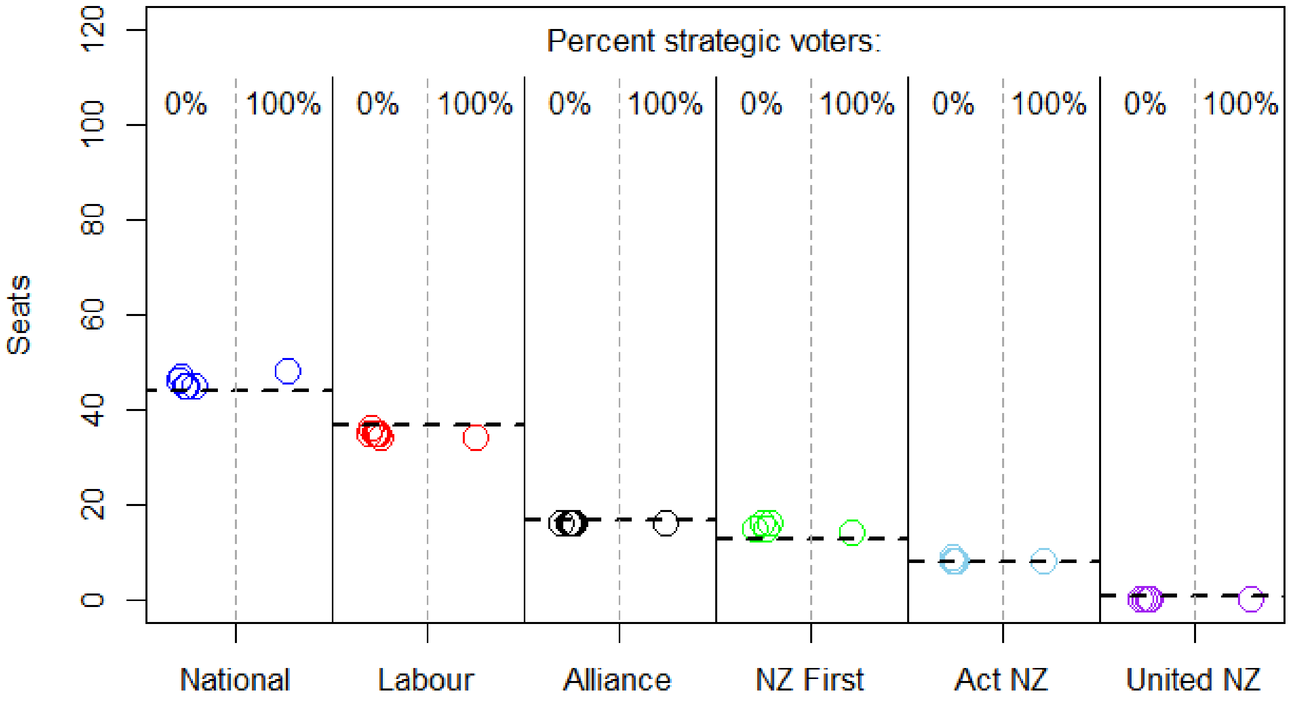

Figure 6: The results of the sincere and strategic model in New Zealand's first MMP election. The results are based on public opinion data for that election, together with the last SMDP election results. The real results of the first Mixed-Member Proportional election are marked with horizontal lines.

Figure 6 demonstrates that preference data from before the 1996 election, combined with the results of the last SMDP election, are enough information for the model to almost exactly anticipate the results of the first Mixed-Member Proportional election.[31]

In Chapter IV, there will be one more validation test performed in cases where the electoral system did not change. When applying the model to an SMDP country in the following chapter, I will first check that the model replicates the results of the real SMDP election, before speculating about the effects of electoral system changes.

---

[31] Two parties, ACT New Zealand and United New Zealand, won seats in 1996 but did not yet exist in 1993. Preference data is available for ACT, but in the simulation no electors are assigned to have voted for it in 1993. No preference data is available for United New Zealand, so it is simply assigned 0 seats in the simulation.

Conclusion

In this chapter, I constructed a method for estimating how a specific political party in a specific election would have been affected by a switch to a different electoral system. There is deep consensus that electoral system changes happen when the people in charge of powerful political parties believe that the change would be in their party's interest. But there is no tool for understanding precisely what those interests are: exactly how many more or fewer seats would a certain party receive under a particular electoral system? This makes it extremely plausible that party leaders have taken positions on electoral system reform that they *think* are in their party's interests, but really aren't. So, this chapter outlined a framework for constructing computer simulations of elections in which it is possible to change just the electoral system without changing anything else. This makes it possible to estimate what would have happened to a specific party in a specific election if the electoral system had been different. I showed that this model closely matches the real example of New Zealand's 1996 electoral system change. In the following chapter I will further validate it, and then apply it to the cases of Canada, Britain, and Alaska.

CHAPTER IV

Simulating Election Reforms in Canada, Britain, and the US

In this chapter I apply the model that was developed in Chapter III, based on ideas from Chapter II.

When applied to a recent Canadian election, the model will suggest that Canada's smaller leftist parties could have benefited greatly from joining the Liberal Party in their support for Ranked-Choice Voting. Whether people adjust their vote to the new electoral system, or whether they simply vote for the party they rank highest on surveys, Ranked-Choice Voting may benefit many more parties than the conventional wisdom suggests. At the same time, the model largely vindicates the other parties' positions on alternative systems.

In Britain, the results of the 2019 election appear much more stable than expected when the system is changed. That was a wave election, and some observers attributed the substantial Conservative majority to distortions by the SMDP system. The model disagrees: it suggests the preferences that electors reported in that election combined with the results of the previous election would have produced a similar outcome under even more proportional systems.

I then model the extremely elabourate new system for electing federal representatives from Alaska. The model produces results in line with the conventional wisdom that the new electoral system is a bespoke political institution tailored to the interests of one specific incumbent. At the same time, it suggests that other systems would have been even more favourable to that incumbent.

Comparing model results across countries suggests some tendencies for how different systems affect different types of parties. For example, the model hints that Ranked-Choice Voting might boost broadly palatable issue parties more than expected.

Electoral system change in Canada

When Canada's left-of-centre parties could not agree on which system to replace SMDP with, election reform was doomed. The majority Liberal party judged that Ranked-Choice Voting would favour it more than SMDP, but that SMDP was better than any form of PR. The smaller leftist parties, the New Democratic Party (NDP) and the Green Party, favoured PR over SMDP, and were unwilling to consider Ranked-Choice Voting. The unified party of the right, the Conservatives, did not support any change, as did the regional Bloc Québécois. But how do parties make such fine judgments about the expected effects of specific alternative electoral systems, when even the academic literature is full of dire warnings like that of Taagepera (2007, p. 36): to predict the specific workings of complex electoral systems is to "kid yourself"?

In judging the difference between SMDP and PR, at least, the simple approach seems to explain why the parties behaved as they did. Figure 7 shows, for every election in the current party system, the difference between how many seats each party would have won under a purely proportional system and the number of seats that they actually won under the existing SMDP system, if electors did not adjust their votes. So positive numbers are a better result under PR than SMDP, and negative numbers are a worse result under PR.

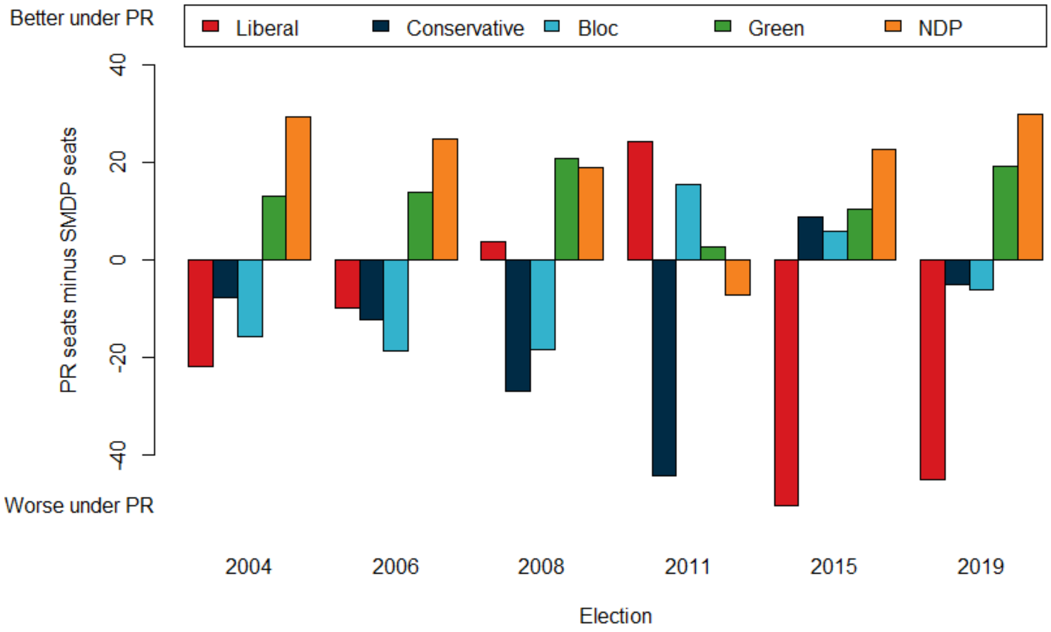

Figure 7: Estimating SMDP versus proportional results in Canada without any strategic voting. The values are the seats that each party would have won under a purely proportional ruleset minus the seats it actually won, for each election of the current party system.

Figure 7 clarifies the party leaders' reactions to the prospect of PR. Looking just at the proportion of votes that each party usually wins, the New Democratic Party and the Green Party should expect to benefit from PR, while the Conservatives, Bloc, and nearly always the Liberals would expect that change to harm them. But of course, this simple comparison suffers from *both* problems discussed in the last chapter under the heading "Modeling party fortunes: two problems and a solution". First, there is no similarly obvious way to capture the parties' guesses about how Ranked-Choice Voting would affect them. And second, voters would not have voted the same way under PR as they did under SMDP.

Thankfully, the conventional wisdom on Ranked-Choice Voting in Canada is quite simple. As Bryden (2015) wrote, commentators have been unwaveringly certain that the Liberals "would have secured an even more lopsided majority under ranked balloting" in the 2015 election. Dutil (2017) goes further, writing that "[i]t took little time for experts to predict, using past results and some imagination, that under [Ranked-Choice Voting] the Liberals would be guaranteed a place in government forever." And this is why Canada does not use Ranked-Choice voting today: it was viewed as such a gift to the Liberals that it "was a non-starter for the majority of non-Liberals on the [parliamentary election reform] committee" (Dutil, 2017).

Abstractly, the idea that Ranked-Choice Voting would likely help the centrist Liberal party stands on firm ground. The original proposal by de Borda (1781) to rank candidates on ballots was motivated by the observation that single-vote elections might not select the person who is most broadly acceptable across the electorate, and that same framing spurred the modern revival of the ranked ballot idea (Black, 1958, p. 157). The promise of Ranked-Choice Voting is that it will often "favor politicians that are more centrist" (Tolbert and Kuznetsova, 2021, p. 266). So it is thoroughly reasonable to expect that Ranked-Choice Voting would likely help the Liberal Party, a party so close to the median voter that they are commonly known as the country's "natural governing party". However, it requires a tremendous leap to start with the broad tendency that Ranked-Choice Voting usually helps centrists and conclude that the Liberals would be *guaranteed a place in government forever*.

And suppose we knew confidently that Ranked-Choice Voting were the best option for the Liberals and the worst option for everybody else. Even if we granted that, it glosses over crucial details. How much exactly would Ranked-Choice Voting help the Liberals? The New Democratic Party and Greens flatly opposed Ranked-Choice Voting because it would help the Liberals, but could it not at the same time help them, at the expense of the Conservatives or the Bloc?

In the following section I make the first effort to answer these questions empirically. First, I check the performance of the model with varying levels of strategic voting against the real results of Canada's 2019 election (which was its first election after the election reform discussions, and produced almost identical seat counts to the subsequent 2021 election). The model closely replicates the election results. Then, I vary the electoral system in two ways. First, I improve on the simple PR analysis of Figure 7 by simulating a Canadian election under a common type of PR, also with varying levels of strategic voting. These models suggest that PR would harm major parties and help minor parties as expected, but that the level of strategic voting is extremely important in determining how large this effect is. Then, I model an illustrative type of Ranked-Choice Voting. Here, the model departs strongly from conventional wisdom: the Liberals perform much worse under Ranked-Choice Voting than expected, and the New Democratic Party and Greens perform much better, whereas the other two parties are hardly affected. This is true in the purely sincere model, meaning that it is a pattern that exists in the survey data alone, and it is even more true when strategic voters are included. I then compare these substantive findings about Canadian politics to findings about other cases, and suggest a few general patterns in electoral system changes.

Empirical model test: Simulating SMDP in Canada

Before addressing any substantive questions about other electoral systems, there is an opportunity to check that the model agrees with reality: using the 2019 Canadian National Election Study (Stephenson et al., 2020), combined with Canada's 2015 election results, does the model closely replicate the 2019 Canadian election results *without changing the electoral system*? For every party, and every alternative system, I will present three results: one from a run of the sincere model, one using a point estimate of the number of strategic voters in Canada, and one in which every elector is strategic.

I report all three of these results because there is no one consensus proportion of how many people vote strategically, and in fact, measured proportions vary quite widely. By examining both the system with no, middling, and complete strategic voting, we get three pieces of information, each of them interesting. The sincere model is the result of an election in which the only ingredient is public opinion data. Comparing those results to the system with entirely strategic voting shows the direction of strategic behaviour. Together, these provide a range of results, depending on how much strategic voting a reader expects to occur. In the middle is a point estimate that approximately matches the level of strategic voting (about 30%) identified in a recent Canadian election by Eggers et al. (2021).[32]

---

[32] I use this estimate because it is from a recent Canadian election, it uses a definition of insincere voting that is in the

So, the core property to be validated is that the number of seats that each party won in the real 2019 election should lie in between the sincere and strategic models. We should also expect both the sincere and strategic models to be close to the real results, and the model with 30% of electors voting strategically should be closer than either of them. The results of this validation check are shown in Figure 8.

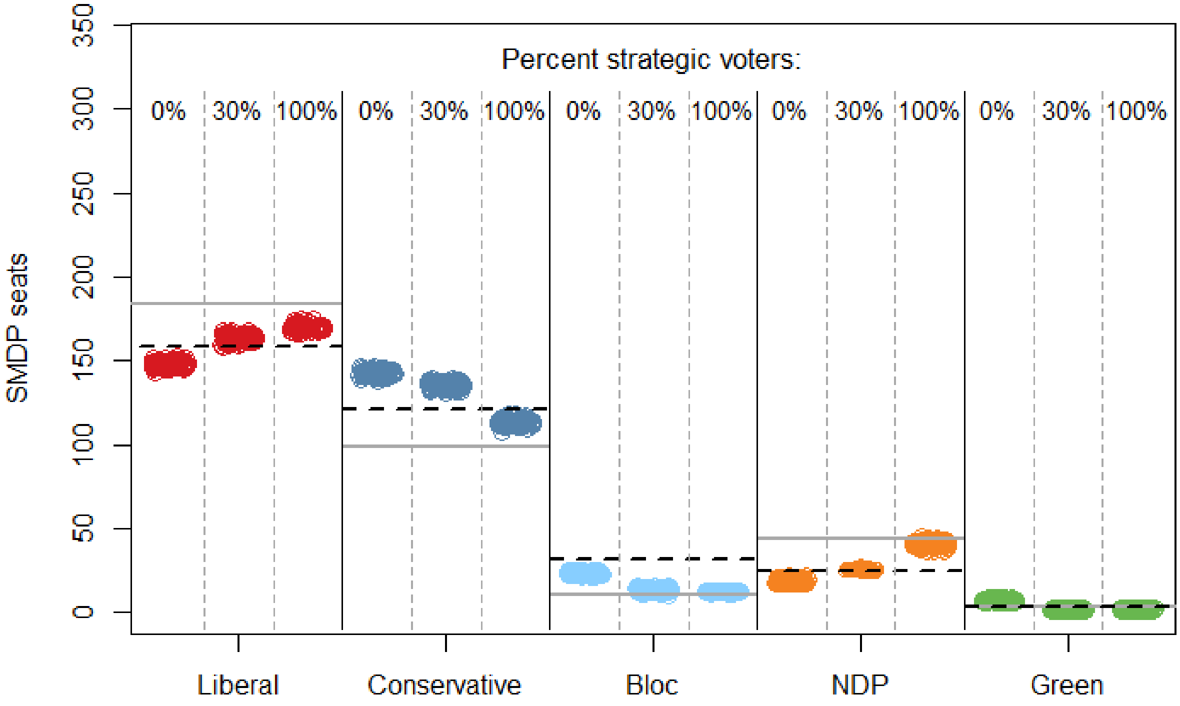

Figure 8: The results of the sincere and strategic simulations under SMDP, Canada 2019. Results presented with a small random horizontal jitter so that the density of points can be distinguished. The model is run 100 times with each strategic proportion. The real 2019 election results are represented with a black horizontal line, and the previous election's results with a grey line.

For three of the five parties, the model with realistic levels of strategic voting almost exactly reproduces the real seat total. The model slightly overestimates the Conservatives, and slightly underestimates the Bloc. For every party except the Bloc, the core validation check is satisfied: the real results lie between the sincere and strategic models.[33] Neither the entirely sincere model nor

---

same spirit as the strategic voting idea studied here, and is neither extremely high nor extremely low compared to other estimates of strategic voting in Canada, so that a reader will find it easier to make a rough guess about what the model would produce using those estimates instead.

[33] The problem is that the district-by-district public opinion data underestimates the Bloc to the tune of about 10 seats,

the entirely strategic model obviously performs better than the other.

Notice that the fully sincere and fully strategic models are both nearly correct for the Liberals and the Greens. In the real event, the Liberals won a narrow plurality of 157 seats. Across 100 runs[34] of the sincere model, the Liberals won a narrow plurality 96 times, twice tying with the Conservatives and losing to them another two times, and the median Liberal seat count was 147. The fully strategic model overestimates them by about the same amount as the sincere model underestimates them, boosting their median seat count to 170 and giving them a victory in every model run. So the model with 30% strategic voters is a successful compromise: it assigns them 162 seats, only 5 away from their actual 157.

The model also closely approximates the results of the smaller leftist parties. Because the New Democratic Party fortunes fell so much from 2015 to 2019, the sincere model does a better job than the fully strategic model of capturing their 2019 result: the sincere model gives them 19 seats on average compared to the strategic model's 41 seats, when in reality they won 24. But again the partially strategic model falls quite close to the truth, with 29 seats. And the sincere model assigns the Green Party 7 seats, but both the partially and fully strategic models correct this number down to 2, which is 1 less than their real total.

The biggest miss in any of the models is the sincere model's estimation of the Conservative Party. The public opinion data alone reports that substantially more Canadians place the Conservatives first in their preference ordering than election results would suggest, and the sincere model overestimates them accordingly, assigning them 142 seats to their actual 121. But the fully strategic model is highly successful in correcting for that, pulling them down to a median 113 seats, while in the partially strategic model they win 134. So although the partially strategic model is further off for the Conservatives than for the Liberals, New Democratic Party, or Greens, still it comes within 10% of the truth.

---

and because they did worse in 2015 than in 2019, strategic voting only exacerbates the problem. Since parliamentary seats are zero sum, the Conservatives are commensurately boosted, but in their case strategic voting does correct the issue.

[34] At the end of the model test I discuss why the number of runs can be arbitrary and does not to be particularly high. The reason is described at the end of this section: the variation in results is extremely low between model runs.

But while the Conservatives are overestimated in the survey's preference orderings, the Bloc are systematically underestimated. This leads to the strategic voting model's one systematic miss: because the Bloc performed so much worse in the 2015 election (when they only won 10 seats) than in the 2019 election, strategic voters in about a dozen Québec ridings flock to other parties. So while the Bloc won 32 seats in the 2019 election, the sincere model comes closest with an average of 23 seats, but the partially strategic model assigns them 11 seats while the fully strategic model gives them 10. But it is important to note that this miss is not enough to change any substantive findings from the model: in the real event, the Bloc placed third, performing slightly better than the New Democratic Party; in the average model run, they place fourth, performing slightly worse than the New Democratic Party. So even the model's biggest error is a small one.

I motivated the decision to examine three sets of results — no strategic voting, some strategic voting, and universal strategic voting — by arguing that, while some strategic voting is most empirically reasonable, it is substantively interesting to observe the direction of change between the results with no strategic voting and high strategic voting. In SMDP strategic voting should usually benefit larger parties and punish smaller ones, and that is nearly always true in Figure 8. One exception, the New Democratic Party, is because the previous election was their most successful ever. In most other elections, the arrow of strategic voting would push against them. But the other exception, the Conservatives, is consistent with the longstanding belief that strategic coordination by supporters of leftist parties works to the Conservatives' detriment.[35]

One more important detail is how often the model is correct about specific ridings. In a country with such notoriously regional politics, it is important that the model roughly captures the geographic distribution of seats, and not just the number of seats.[36]

Figure 9 maps the results of the 2019 election alongside the results of both the sincere and strategic models. The models are seen to broadly replicate the geographic distribution of results.

---

[35] CBC, 2008. "'Anything but Conservative' campaign hurting N.L. Tories." *Canadian Broadcasting Corporation.*

[36] I am grateful to Lori Thorlakson for emphasizing this point.

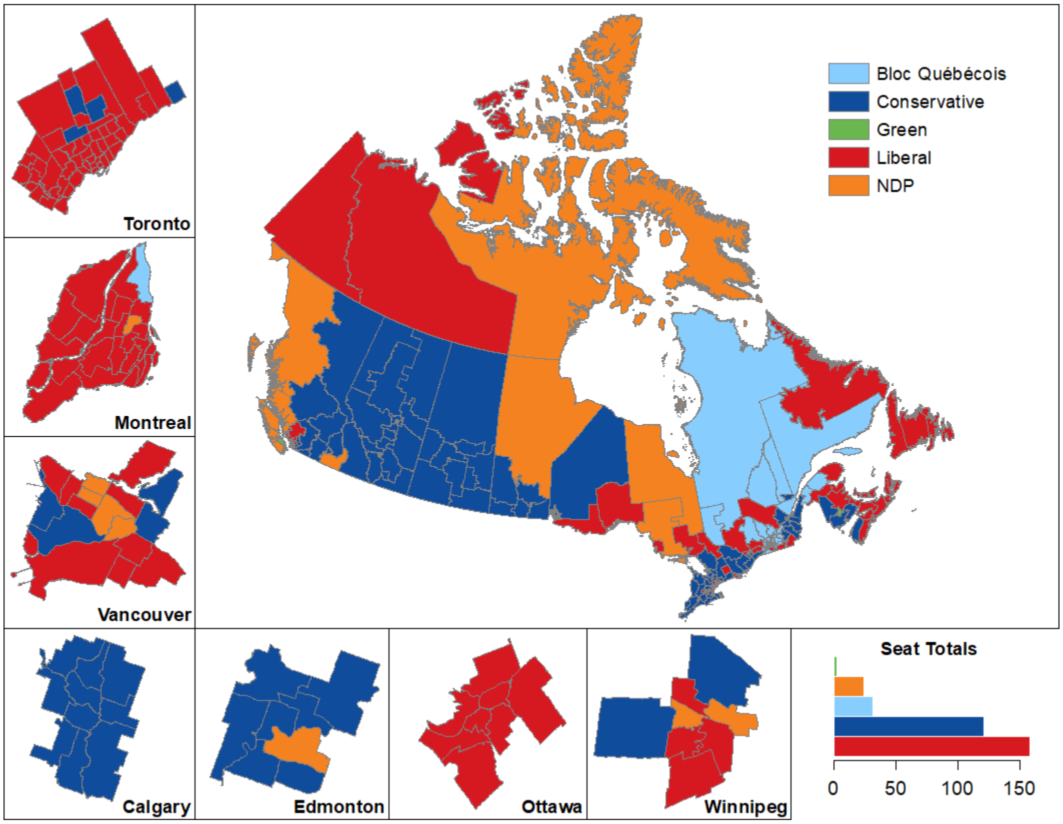

(a) Canada's real 2019 election results.

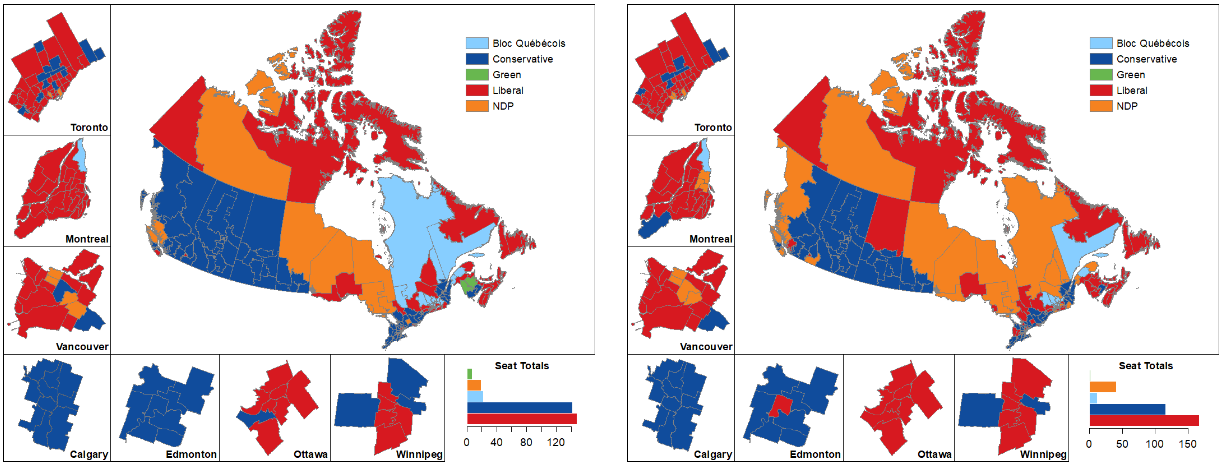

(b) Simulated SMDP, all sincere.

(c) Simulated SMDP, all strategic.

Figure 9: Canada's real 2019 election results, and sincere and strategic simulation results.

While the proportion of seats mis-allocated and maps of the results give a sense of the riding-level accuracy of the various models, this of course is also a calculable number: how many ridings did each model miss? In the Appendix, Table A.1 displays that result for every model under consideration, and for threshold pivotality with $\tau = 5\%$, in both Canada and Britain.

One noteworthy feature of the model's results is how close together the results of the different runs are. This might be unexpected since voting games are famously characterized by a multitude of best response equilibria. The simple answer is large sample statistics. For the sincere model, a list of dozens or hundreds of preferences is sampled from with replacement and with uniform probability, thousands of times per district. This is done independently district-by-district. Then the strategic model proceeds deterministically based on the initial vote distribution (this is true because there is one unique correct pivotal probability for any initial vote choice, and electors have full information), although it also bears mentioning that more complicated reasoning will exaggerate initial conditions more noticeably, so that for example the strategic model of RCV gives more varied results than the strategic model of SMDP. But all of this means that the probability of a systematic departure in the party-by-party seat counts is quite low: to randomly tip the scales towards one particular party would require a coincidentally high sampling of supporters of one party many thousands of times across districts. This is possible, but rare. So the model results are quite similar across runs.

Evidently the model has, in the language of Whicker and Sigelman (1991, p.66), substantial "event validity": in Canada and New Zealand both, it produced results that were similar to real events.

Of course, it does so even though real electors do not calculate their pivotal probability according to Equation 1, and since pivotal voting is only a stylization of real behaviour, we might expect that the model would be similarly accurate with a different model of strategic voting. In the following section, I follow those researchers who argue that people actually gauge the competitiveness of parties using simple heuristics, including some who have employed these types of rules in other computer models of elections (Bendor et al., 2011, Laver and Sergenti, 2012). I endow electors with a much simpler decision rule, and I find that the model produces nearly the same results.

A heuristic voting analysis

Now that we have a baseline model performance to compare it to, it is worth revisiting how much the precise strategic voting rule matters. Recall that in Chapter II we performed an event validity check comparing the results of the model using the minimalistic heuristic of threshold pivotality to the real results of the election. Throughout this chapter I will, for simplicity, use the election day stopping rule with a default number of days set to 0, so that the model ends the moment that iterative equilibrium is identified. This is an arbitrary choice since in Figure 3 it was demonstrated that the choice of stopping rule does not strongly affect the model results, at least in Canada's 2019 election. Here I reproduce the plot of the election day stopping rule for every possible value of τ:

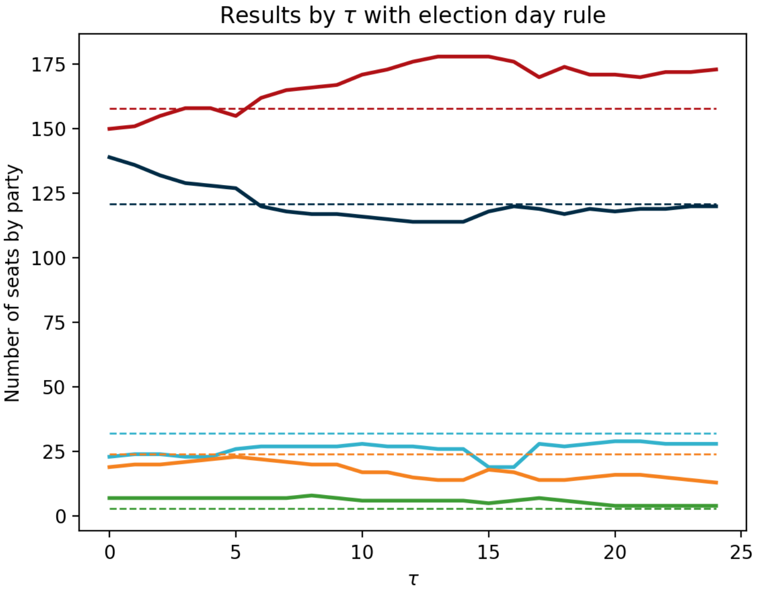

Figure 10: Results of the election day rule, which stops immediately at iterative equilibrium. This plot was originally shown alongside the results of other stopping rules in Figure 3 of Chapter II. As in that figure, dotted lines represent the real election results for each party, and solid lines show the simulated results, for each of the five parties in parliament: Liberals (red), Conservatives (dark blue), Bloc (light blue), NDP (orange), and Green (green).

Table 3 shows examples of typical runs of the sincere and strategic models, along with the heuristic model results with the election day rule, with the number of iterations set to 0, and with

$\tau = 5\%$.[37]

	Liberals	Conservatives	Bloc	Greens	NDP	**Errors**
Real	158	121	32	3	24	
Sincere model	149	142	24	6	17	
	(9)	(21)	(8)	(3)	(7)	**24**
Strategic model	169	115	12	2	40	
	(11)	(6)	(20)	(1)	(16)	**27**
Threshold pivotality, $\tau = 5\%$	155	127	26	7	23	
	(3)	(6)	(6)	(4)	(1)	**10**

Table 3: Results from sincere, strategic, and heuristic model versions, and mis-allocated seats. The numbers in parentheses show the party-by-party difference, and the Errors column shows the total number of mis-allocated seats (half of the sum of the party-by-party differences). The heuristic version uses the election day stopping rule with $\tau = 5\%$.

As we found in Chapter II, what stands out about the heuristic model is how exceptionally close to the actual results it is. This further underscores why it is useful to present both the sincere and strategic model results, since the heuristic method is something like an average of these two. In light of the strategic model's success in matching the real election results, producing less than half the error rate of the sincere model and closer to a third the error rate of the strategic model, it is worth addressing why we do not simply use the heuristic model throughout. In fact, of the three models, the strategic model that required so much effort to produce in Chapter III actually performs worse than either the heuristic model or the sincere model.

One critique of simply using the heuristic model would be that the heuristic model which is represented here is actually very nearly the single best model out of 500 models run in Figure 4. However, studying that figure gives a ready response to this objection. First, it is worth noting that this value and stopping rule — the election day rule with 0 iterations, and $\tau = 5\%$ — were actually chosen as substantively reasonable choices before the quality of the different stopping rules and

---

[37] I chose the value $\tau = 5\%$ before running the model for all values of τ, so at the time that I chose it I was not aware that it was extremely close to the model that gets the closest results to the real election, which in the models that generated Figure 4 was $\tau = 6\%$.

values of τ was known. But more importantly, together Figure 3 and Figure 4 showed that the choice of stopping rule does not matter very much, and that every stopping rule with a value of τ between about 1% and 10% (which is probably also about the substantively reasonable range of what most electors might consider to be a "near-tie") yields a number of mis-allocated seats between about 10 and 20. Indeed, up to about $\tau = 30\%$, a value of τ that has exited the territory of near-ties and has entered the territory of true electoral landslides (a district where one party leads by 30% would surely be among the least competitive districts in almost any multiparty democracy), no stopping rule ever produces a number of mis-allocated seats above 25. So any stopping rule and nearly any reasonable choice of τ would actually beat the strategic model, and nearly beat the sincere model. So it is not hard to accept that the heuristic model is truly the most accurate of the three.

But a more serious question is why, if it tends to be more accurate, would it not be better to simply use this heuristic throughout? What was the point of the involved derivations in Chapter III if an extremely simple heuristic does a better job of matching real vote totals? The answer lies in those very derivations. After adopting the Poisson voting games framework, there is one unique correct way to arrive at Equation 1. With Equation 1 in hand, there is just one correct way to derive Equation 2. Extra assumptions (like that the population of electors is Poisson distributed, and not known exactly to all electors) had to be made to construct a model for which expected vote shares have a unique correct expected value. These assumptions were not directly motivated by any substantive reality; they are reasonable enough representations of the world, but their main motivation is that they produce much more usable mathematical representations of elections (for example, once the population of electors is Poisson distributed, so too are the expected vote shares of the candidates, which makes the expected difference between those vote shares calculable using known distributions).

It is only to be expected that these assumptions might drag the model's results slightly away from reality. This is one manifestation of a classic trade-off: in exchange for a few more mis-allocated seats, the model can be confidently extended to cover other voting systems. In the language of Page

(2018, p. 6) the Poisson voting model is further abstracted away from reality than the heuristic model is, but in exchange for that, the Poisson voting model is also more formalized and more precise.

With less structured descriptions of reality, like the threshold pivotality story, it is not at all obvious how we could broaden threshold pivotality logic to cover voting in a Borda count election. We could probably quickly make up half a dozen similarly plausible analogues to threshold pivotality that describe how people might fill out Borda count ballots. But the model will not help us to discern which is correct among them. The strategic voting models of Chapter III are quite different: Equation 2 is the unique correct calculation of the utility of Borda count ballot that corresponds to Equation 1. Of course all the same reasoning applies to why the strategic model run in Table 3 performs very slightly worse than the sincere model, and why that does not threaten its value (although it will soon be shown that they make very different mistakes, and the strategic model actually does correct some of the sincere model's problems).

But it also is important not to overstate the differences between the models. Of the 338 seats in question, the strategic model only mis-allocated about two dozen more than the heuristic model, and about the same number as the sincere model. Really, Table 3 shows close agreement between the heuristic model and the strategic model. To facilitate comparison with the analyses of threshold pivotality voting in Chapter II, the same random number seed from those analyses was used for the heuristic model run shown in Table 3 (so this is the same model run that appears in the figures in that chapter with $\tau = 5\%$), but this was different from the seed used to generate the strategic model run in the table. When the seeds were instead synchronized (running the heuristic model on the same seed as the sincere and strategic model), the heuristic decided 28 (8%) of the 338 ridings differently than the sincere model, and 50 (15%) differently than the strategic model. If 1 in every 6 races were *systematically* different from the strategic model (say, if 50 Liberal wins in the strategic model results were given to the Conservatives in the heuristic results) then this could produce a fairly large difference in seat counts. But the heuristic does not point to the strategic model strongly over- or under-estimating any party. The heuristic model decides about a dozen seats more in favour of the Bloc and about a dozen fewer in favour of the New Democratic Party than the strategic model, but it hews very closely to the strategic model's expectations about the other parties. This is an example of a pattern demonstrated in a broader context by Vasselai (2021a), that seemingly quite different ways of calculating pivotal probabilities in these sorts of games tends to produce very similar results, which this example further supports.

In light of this agreement, a reader who is skeptical of the SMDP model's applicability to real elections, and who believes that simple heuristics are a better description of voting decisions, should at most add or subtract a few percentage points to the model's anticipated party vote shares. Of course the Ranked-Choice Vote model is more complicated, and may depart more from whatever heuristic one might imagine voters using in an Ranked-Choice Vote election. But the Ranked-Choice Vote model is just an extension of the SMDP model which involves the same exact pivotality calculations repeated more times, so the fact that simple rules replicate the results of those pivotality calculations should lend credence to the Ranked-Choice Vote model too. So one does not need to believe that pivotal voting is a literal description of reality in order to find the results of these models informative.

With that result in hand, the model can now be used to anticipate election results under counterfactual systems.

Simulating subnational PR in Canada

With no further amendments to the model, we can directly simulate a type of more proportional system. Two observations make this possible. First, many more proportional systems are constructed by simply increasing the size of districts to create large multi-member districts. In the limit as district size approaches population size, the system becomes increasingly proportional. Second, Equation 1 can also be used to simulate elections for more than one seat so long as votes are non-transferable (Cox, 1994).

The most natural multi-member districts to draw are to amalgamate the single-member districts up to the size of Canada's provinces or territories, and to endow those hypothetical province- or territory-wide multi-member districts with magnitude equal to the number of single-member districts that they contain in reality. I choose this structure because multi-member districts nearly always conform to subnational administrative boundaries, and this very idea has been used by LeDuc (1999, p. 72 – 74) to estimate a more proportional Canadian Senate, as well as being reminiscent of previous reform proposals like that of the Macdonald Commission. I model a PR system with seats allocated by the D'Hondt rule, but I assume that electors do not strategize around the negligible chances that their vote affects a D'Hondt allocation; that is, electors strategize as though they are casting a non-transferable vote. Then we can simply run the strategic model as discussed so far, seeding each elector with the results that their hypothetical districts would have seen in 2015 by also aggregating the real district results up to province- or territory-wide results in 2015. The results of that contest are shown in Figure 11, while the following figure, Figure 12, maps the results.

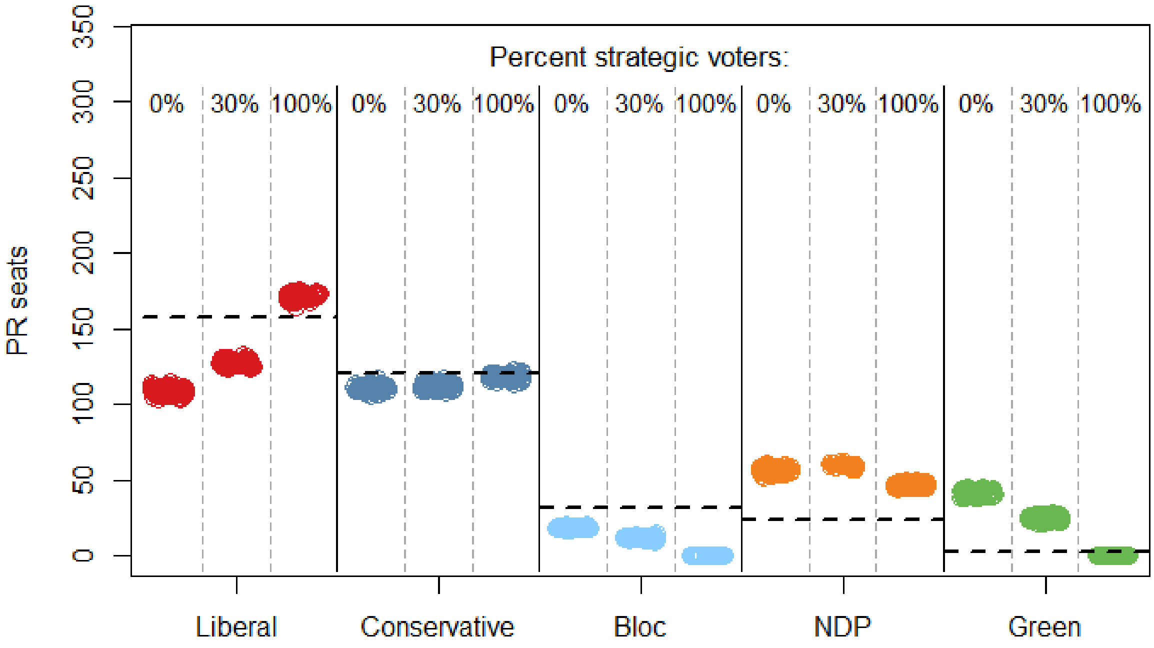

Figure 11: Results from a hypothetical Canadian PR election, varying levels of strategic voting. The model is run 100 times with each strategic proportion. The real 2019 election results under SMDP are marked with horizontal lines.

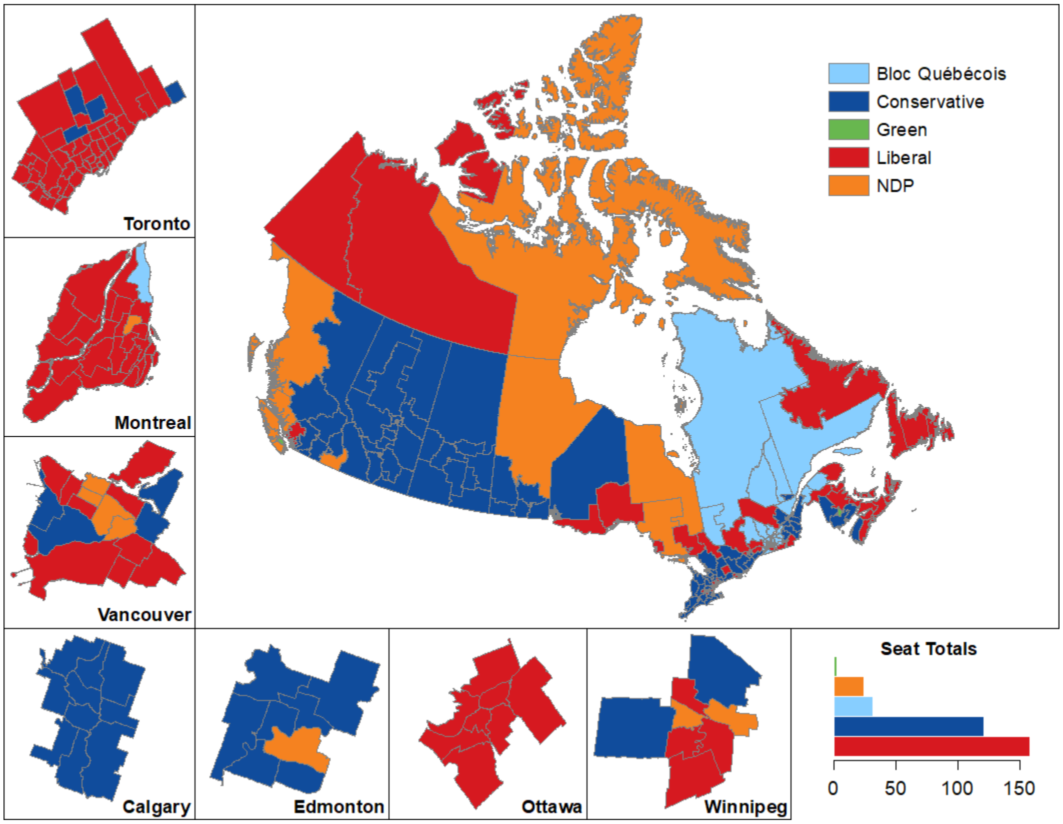

(a) Canada's real 2019 election results.

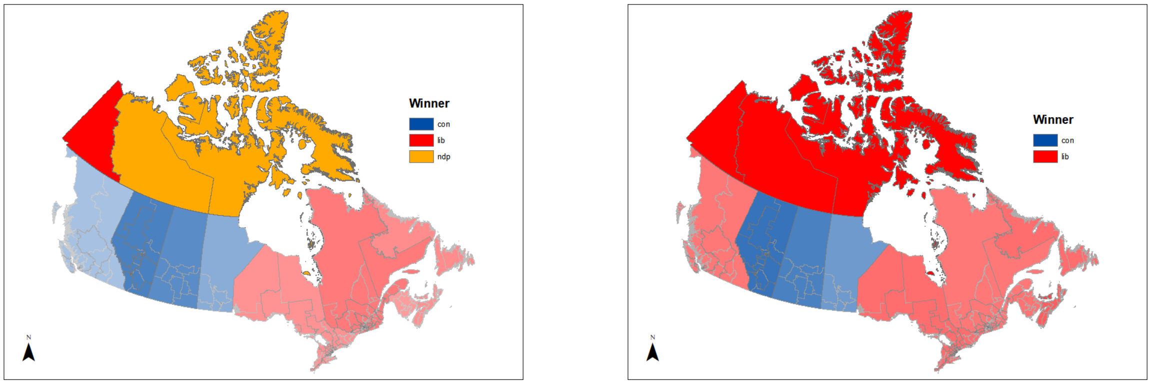

(b) Simulated PR, all sincere.

(c) Simulated PR, all strategic.

Figure 12: Sincere and strategic PR simulation results, alongside Canada's real 2019 election results.

The results of the sincere model in Figure 11 show another view of the conventional wisdom: the Liberals, Conservatives, and Bloc expect to lose seats from a change to a more proportional system, while the New Democratic Party and Greens expect to benefit.

However, the model departs in two ways from the conventional wisdom. Although the Conservative Party does indeed win fewer seats under this more proportional system, in runs of the sincere PR simulation they actually frequently win a very narrow plurality of seats. While the Conservatives are correct to expect to lose seats after a switch to PR, they might actually benefit from the relative distribution of seats. A second noteworthy pattern is that the results heavily depend on how much strategic voting the parties expect to encounter in a more proportional Canada. When only 30% of the population votes strategically, the results are hardly different from the sincere model. But in the extreme event where every elector votes strategically, the results show that strategic pressures acting at the provincial and territorial level can severely reduce the gains of parties that are not competitive at that level.

The reason is the one identified by Cox (1994): in a multi-member race it is often more desirable to help elect one more member of an already competitive party than to take a chance on a non-competitive one.[38] Of course, it might be hard to see how this could matter in a province like Ontario, where the volume of available seats suggests that there is minimal pressure to vote strategically. But imagine an elector in Saskatchewan who continues to see the Conservatives hold all of the province's 14 seats in every election, with just a few wins here and there by a Liberal or New Democrat, despite a switch to province-wide PR. It is easy to see why they might continue to think that a Green Party vote is a waste. So the strategic pressures under certain types of PR might be more to the detriment of the smaller leftist parties than is currently appreciated.

Simulating Ranked Choice Voting in Canada

There are many types of Ranked-Choice Voting, and nobody can know which type the Government of Canada would have adopted if the other parties had been willing to support the Liberals'

---

[38] Continuing to ignore vote allocation rules; that result applies to Single Non-Transferable Vote elections.

preference. But at issue is whether or not Ranked-Choice Voting is so centrist-favouring that it will ensure even greater Liberal dominance over the Canadian political scene than the dominance they have enjoyed for the last 150 years. To be as generous as possible to the conventional wisdom, the appropriate choice is to use a ruleset that heavily favours centrists. The main ranked choice alternative to Borda count, Instant Runoff Voting, has a well-known counterintuitive tendency to sometimes eliminate centrist candidates (Brams and Herschbach, 2001). Borda count, on the other hand, was designed specifically to boost centrists.[39]

Figure 13 shows the result of a simulated strategic Borda count election using length 3 complete ballots, retaining the same ridings as real Canadian elections. Figure 14 maps the results, comparing them the results to the real election.

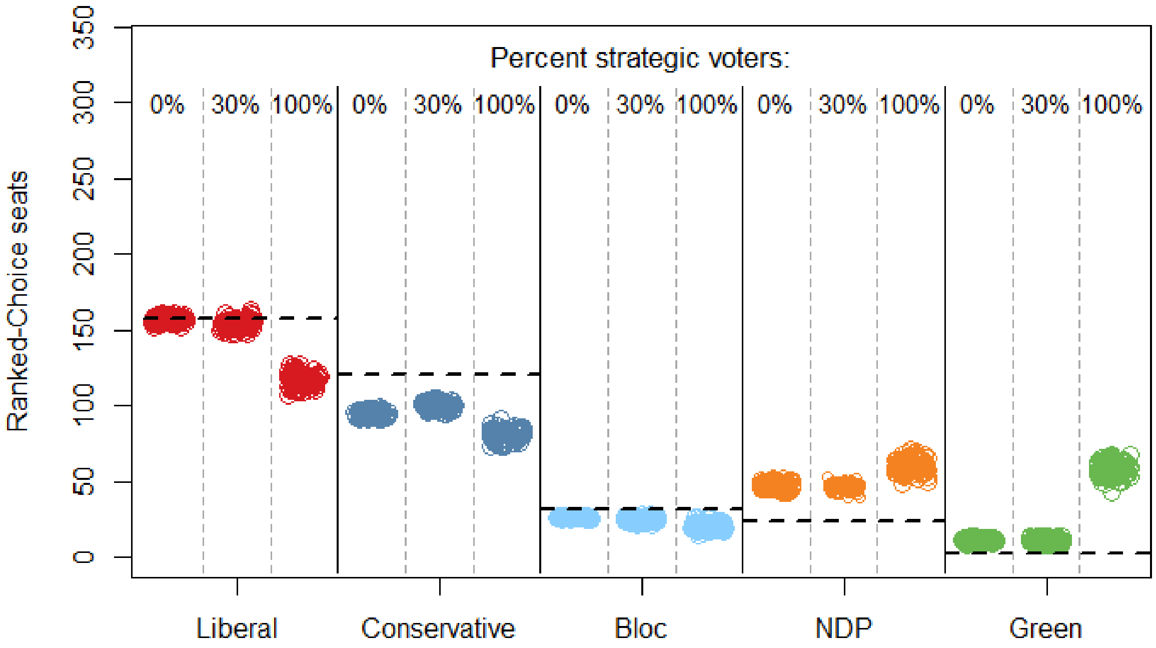

Figure 13: A hypothetical Canadian Borda count election with varying levels of strategic voting. The Borda count ballots have a length of 3 and must be completed. The model is run 100 times with each strategic proportion. The real 2019 election results under SMDP are marked with horizontal lines.

---

[39] One reason not to focus on Borda count is practical: it is extremely rare in national elections, used only in a few European and central Pacific countries. But so are all types of Ranked-Choice Voting; even Instant Runoff Voting, famous for its use in the Australian legislature, is used in almost no other national legislative elections.

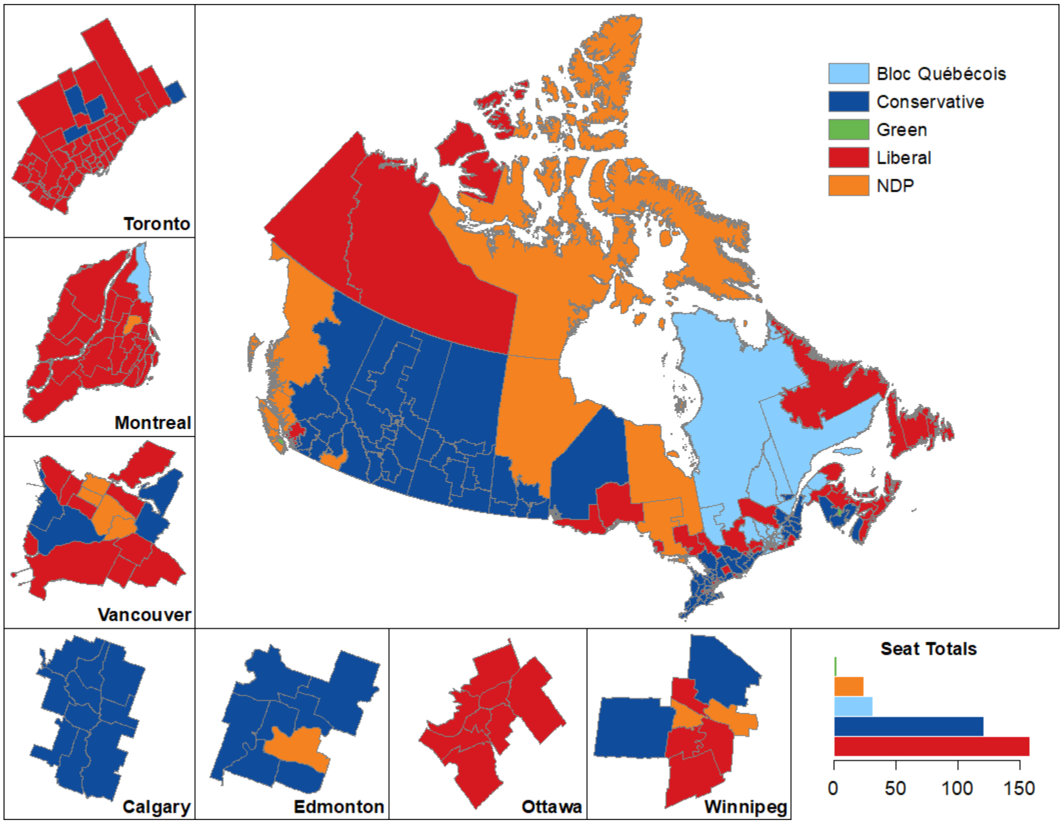

(a) Canada's real 2019 election results.

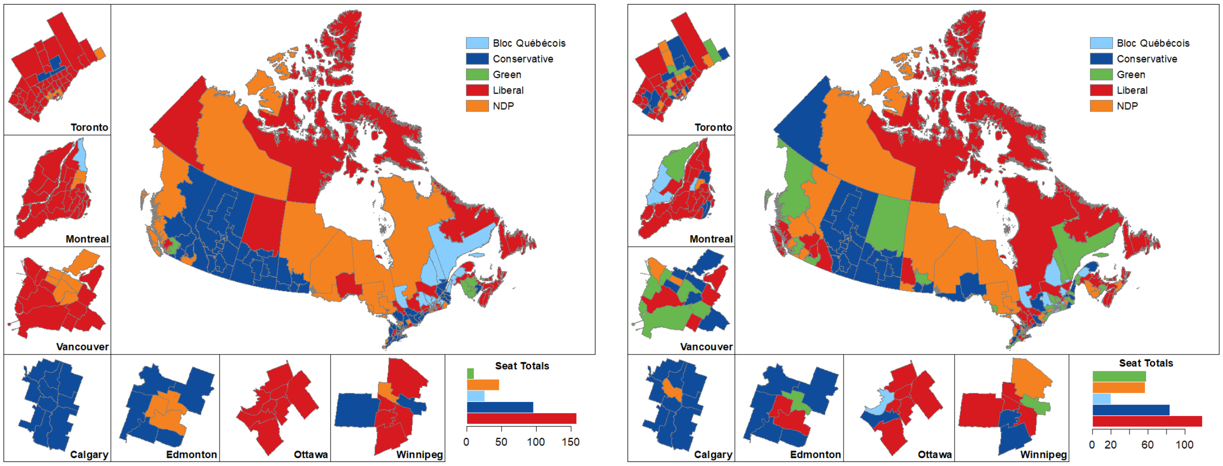

(b) Simulated Borda, all sincere.

(c) Simulated Borda, all strategic.

Figure 14: Sincere and strategic Ranked Choice simulation results, and real 2019 results.

The simulation departs from conventional wisdom in two major ways. Both Canadian parties and the political scientists who have studied possible electoral reforms in the country broadly agree that Ranked-Choice Voting would provide a colossal advantage to the Liberals. The simulation does not support this expectation.

In the sincere model as well as in the model with 30% of people voting strategically, the Liberals actually win almost exactly as many seats as they won in reality. The Conservatives are slightly harmed by the switch to Ranked-Choice Voting, but the Liberals' major leftist competitor, the New Democratic Party, are substantially strengthened. It is hardly obvious how rearranging the other parties' seats in this way would affect the Liberal Party.

The model with strategic voting holds greater surprises. With a conventional wisdom formed largely from treating historical vote counts as fixed and running them through different formulae, political scientists and party leadership alike have judged that it would not be in the interest of the New Democratic Party or the Green party to support the Liberals in switching from SMDP to Ranked-Choice Voting. But the Ranked-Choice Voting model with all voters acting strategically produces a surge for the Green Party. In Canada's real 2019 election, the Green Party won 3 of the 338 seats in parliament while receiving about 6.7% of the vote. In the sincere Ranked-Choice Voting model, they roughly quadruple this number, and they do similarly well with 30% of people vote strategically. But in a typical run of the fully strategic Ranked-Choice Voting model, they capture more than 15% of the seats in parliament.

The reason that this happens in the model is a familiar one in the study of Ranked-Choice Voting strategy. Consider a fairly typical Canadian district that is usually won by either the Liberals or the Conservatives, while seeing very few votes for the Green Party. Most survey respondents in this district either least-prefer the Liberals or least-prefer the Conservatives, while expressing vaguely positive feelings towards the Green Party. So if (say) a Liberal Party supporter in this district is asked to rank 3 parties, their highest expected utility ballot will almost always rank the Liberals first, the Green party second, and the New Democratic Party third, whereas a Conservative Party supporter would rank the Conservatives first, the Green Party second, and the New Democratic Party third. In many districts, the presence in second place on so many Borda count ballots is sufficient to boost the Green Party above either the Liberal Party or the Conservative Party. And because the model is iterative, there is a twist: in the next iteration, it is public information that the Green Party is leading. At this point, one of the formerly leading parties — Liberal or Conservative — has been knocked out of second place. The supporters of that party now expect a higher utility from backing the Green Party rather than sticking with their previous vote choice. In many districts they are able to consolidate behind the Green Party and boost it to first place. A phenomenon like this could play out in a real Canadian election as electors realize that the Green party enjoys a disproportionately high level of second- and third-place support.

The party that pushed a switch to Ranked-Choice Voting, the party that commentators believe would be the sole beneficiary of that system, only loses seats in both the sincere and strategic ranked choice models. Two of the parties that blocked a switch to Ranked-Choice Voting, the Green Party and the New Democratic Party, appear to unambiguously benefit from it. Even when the only ingredient in the model is public opinion data, a switch to Ranked-Choice Voting causes both parties' seat totals to surge. When a very large proportion of electors vote strategically, the fates of those parties only improves. In the next two sections I will build on these substantive findings by investigating whether or not similar patterns hold when the model is applied to Britain instead.

Electoral system change in Britain

In this section I apply the model to a case that resembles Canada, and then in the following section I will apply it to a case that is quite different. This provides an opportunity to check how many of the results in the previous section were specific to the Canadian case, and how many might also apply to other contexts.

Despite being the original Westminster system, the United Kingdom has seen major attempts to replace its electoral system in the past decade, including a failed 2011 referendum to convert the country's federal politics to the Alternative Vote. It is interesting as a prototype of the institutional challenges faced by former British colonies, but it is also a quintessential example of the disproportionalities of SMDP elections. As Shugart and Taagepera (2017, p. 81) note, "in the UK, no party has won a majority of the national vote since 1931, and yet in all but three elections since then, one party has had a majority of seats".

And the comparison to the Canadian case is tight. The party systems of the two countries are strikingly similar (Chhibber and Kollman, 2004, ch. 1). In both national party systems, one party stands alone on the right; in both cases, many parties make up a more fragmented left- and center-left, but only one of them consistently forms the government or opposition. And both countries have strong regional parties. As in Canada, commentators have speculated for many decades about what would happen if the United Kingdom replaced its SMDP system. And just as in Canada, the conventional diagnosis is that the left and the center-left parties split votes, allowing a conservative party to regularly turn pluralities of votes into majorities of seats (or even minorities of votes into pluralities of seats).[40]

So, in addition to addressing a major question about British politics, simulating a British election under some other electoral system gives us the opportunity to answer whether or not certain patterns in the Canadian results might be more general patterns. Figure 15 and Figure 16 show the results of a simulated 2019 UK election under the same alternative systems as were explored in the case of Canada, both of which have been proposed for British elections also. The following figures, Figure 17 and Figure 18, show the geographic distribution of results.

---

[40] British elections to the European Parliament have used more proportional systems (Hix, 2004, McLean and Johnston, 2009), but the question of how a different system would shape British federal elections remains largely open.

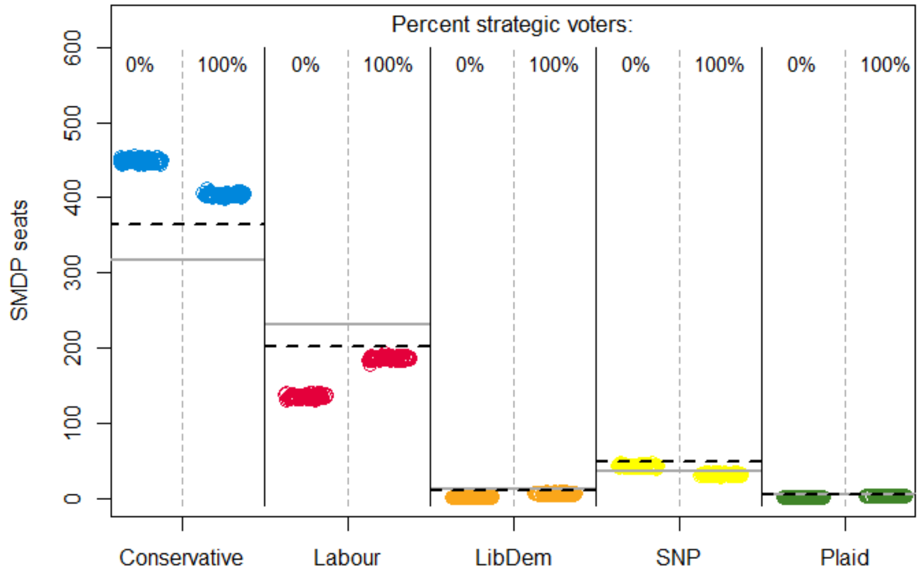

(a) Simulated SMDP.

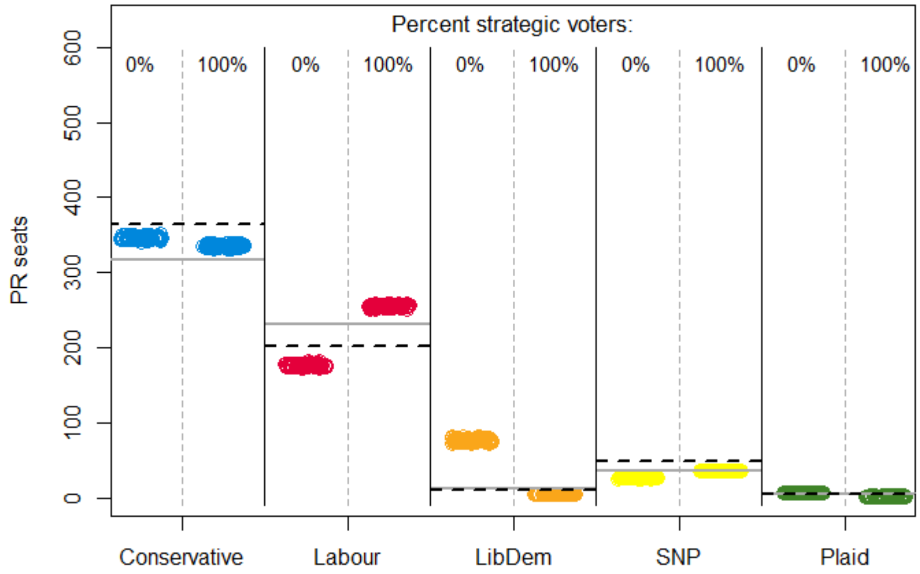

(b) Simulated PR.

Figure 15: SMDP and PR model results in Britain, 2019. Real 2019 election results represented with black horizontal lines, the previous election's results with grey lines.

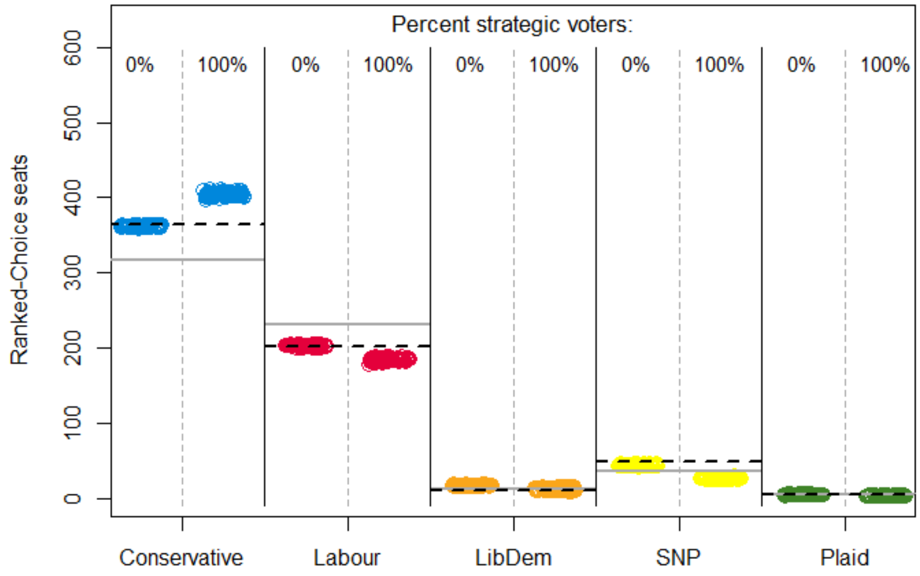

Figure 16: Simulated Ranked Choice Voting results in Britain, 2019. Real 2019 election results represented with black horizontal lines, the previous election's results with grey lines.

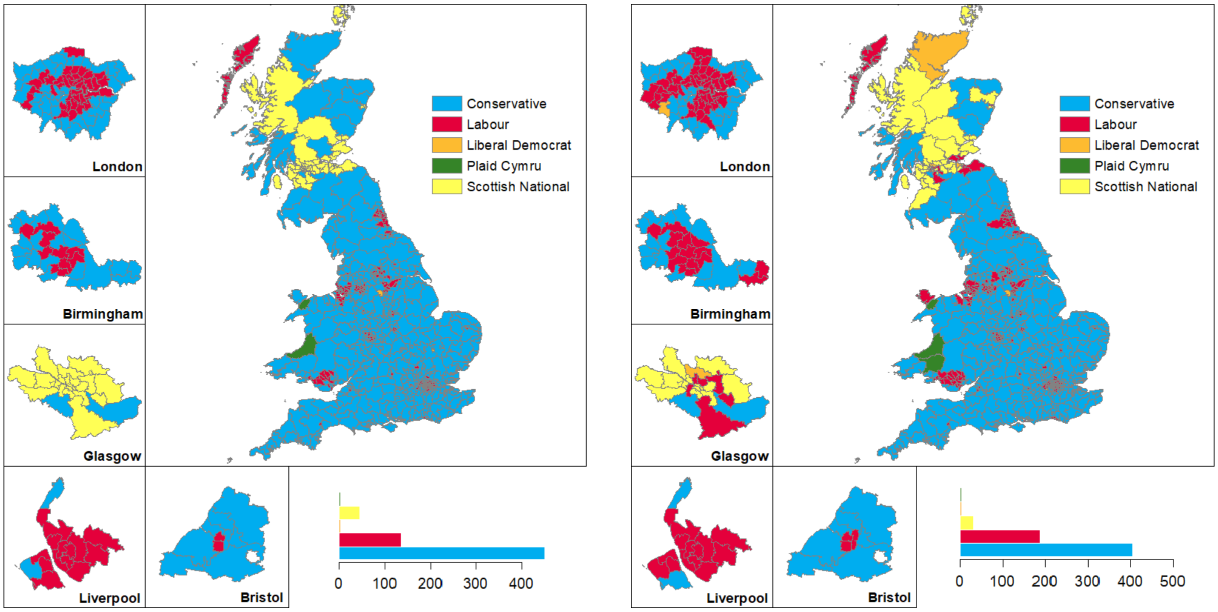

(a) Sincere SMDP.

(b) Strategic SMDP.

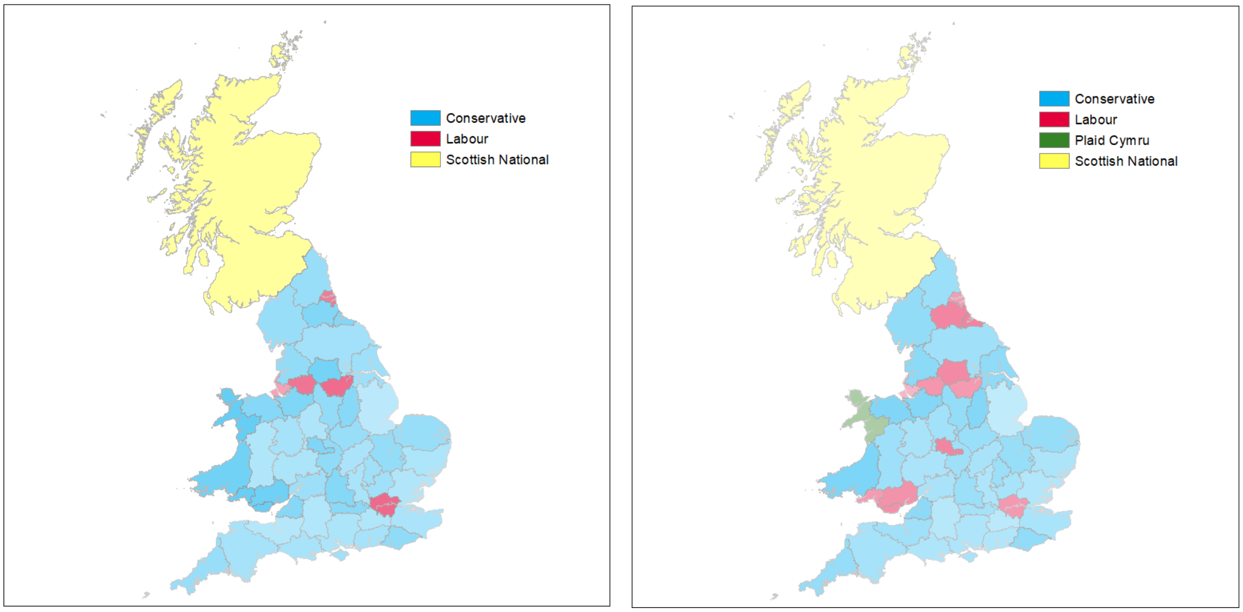

(c) Sincere PR.

(d) Strategic PR.

Figure 17: SMDP and PR simulation maps in Britain, 2019.

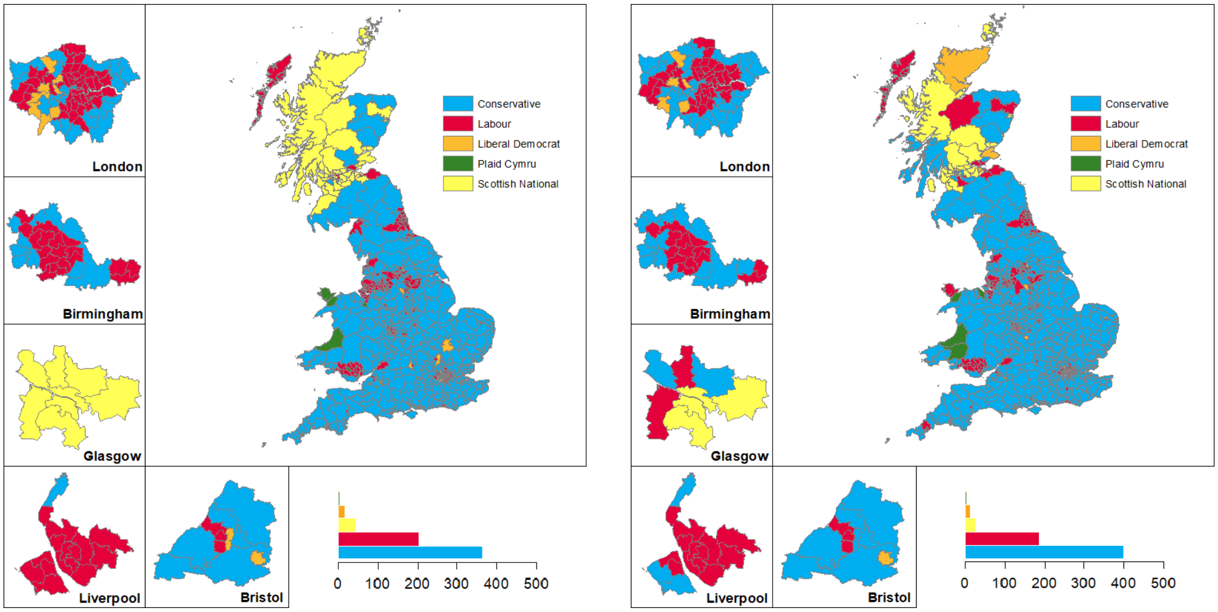

(a) Sincere Ranked Choice.

(b) Strategic Ranked Choice.

Figure 18: Ranked Choice Voting simulation maps in Britain, 2019.

Focusing first on the SMDP seat counts, the event validity of the British model is almost exactly the same as the event validity of the Canadian model. They share the direction of skew (towards the Conservative Party and away from its main competitor), the size of that skew (about 20%), and the fact that the strategic model is closer to reality than the sincere model alone.[41] How do the results compare?

All of the main patterns of the Canadian PR simulations appear in the British simulations too. In the sincere PR model both major parties are somewhat weakened, with the Conservative and Labour parties both receiving a few dozen seats less than they won in reality, but the decrease in their seat shares is small. And here too sincere PR helps smaller parties, as expected. The biggest beneficiaries are the Liberal Democrats, who have long supported switching to PR and who would indeed would multiply their seat share many times over under sincere PR. And yet, the model gives the same warning as in Canada: when all voters act strategically, Labour and the Scottish National Party completely erase the gains of the Liberal Democrats. Again in Britain, the details of the PR system and the level of strategic voting are hugely important for smaller parties.

Turning to the sincere Ranked-Choice Voting simulation, the most striking feature is actually their almost perfect correspondence to the real election results. Every single party receives almost exactly the same number of seats in the sincere Ranked-Choice Voting model as in reality.[42] But in the strategic Ranked-Choice Voting model, there is nothing like the strategic boosting of the Green Party of Canada. It is noteworthy that the only niche issue party in either of these two party systems is the only party that was substantially boosted; perhaps such parties benefit from not just a reduction in strategic behaviour, but also from the adoption of an electoral system in which strategy is dramatically more complicated. One final pattern of note is that all of the simulated electoral

---

[41] To draw district boundaries in the Multi-Member Districts simulations, I use the county boundaries identified in the United Kingdom parliament's official election results, which map counties one-to-many onto the SMDP ridings. This makes for multi-member districts that tend to contain a few dozen members. The preference data for these models are drawn from the 2019 British Election Study (Fieldhouse et al., 2019). Northern Ireland is not included in these simulations because of its distinct party system.

[42] The reason this happens is that the public opinion data slightly overstate support for the Conservatives, and preferences after the first preference are canceling out the top preferences' minor skew towards the Conservatives, just enough to pull the results down to the actual SMDP results. So we should interpret the model as suggesting that Ranked-Choice Voting would slightly reduce the Conservatives' seat share.

systems have broadly similar effects on the three regional parties.

Applying the model in this way to cases with a few more different institutional arrangements could flesh out a general theory of how vote and seat distributions might shift *in response to* an electoral system change, with and without voter strategy explicitly modeled. The least-differences approach of comparing the model in two extremely similar Westminster democracies in the same year has revealed that indeed, the model makes consistent predictions about the impact on the distributions of votes and seats, with only some differences that are due to the particular attitudes of the electorate towards niche parties. Perhaps the model has already furnished observations that could be tested through comparative observation.

In the next section, I illustrate the versatility of this modeling idea by applying it to an extremely different case, which requires some methodological adaptation and yields very different substantive insights.

The setting of electoral system change in Alaska

In 2008, Bowler and Donovan (2008, p. 102) could observe in the United States "an unchanging electoral system but constant change in electoral law", since the constant flux of American electoral procedures had not yet disrupted the nearly universal use of SMDP to elect national legislators. But in the last few years, this edifice has begun to seem much more fragile. Suddenly the conversation on electoral system reform is as active in the United States as anywhere else. Nearly every day, a major American newspaper publishes an opinion or analysis piece that advocates a switch away from SMDP.

Such proposals to change the country's electoral system might seem strictly theoretical, since it is hard to imagine that America's national election rules could be overhauled in the near future. But not all rule changes are federal. In recent national elections, states have begun to actually enact new systems for electing candidates to federal office. Six states have adopted modified two-round single-vote systems, while federal representatives in Maine are now chosen through Instant Runoff

Voting. But one rule change is particularly intriguing: as of November 2020, Alaskans will vote in a single-vote primary election to pick four winners, who then compete in a four-way Instant Runoff Voting contest.

Some conventional electoral systems, like SMDP or common types of proportional representation, have been the focus of decades of research across dozens of democracies. Their pros, their cons, and the kinds of governments they produce are well understood. But what can we expect from Alaska's two-stage hybrid, where a single-vote first round is followed by a four-way ranked choice face-off? As American states are beginning to practice electoral system husbandry, merging the canonical systems to create surprising mutations, there has been almost no scholarly attention to what might effects they might have.

Two parts of Alaska's new system are particularly unusual in the comparative study of electoral systems. First, there is a switch in vote-counting systems between the two stages of the election: in the language of the 2020 referendum, "[t]he four candidates with the most votes in the primary election would have their names placed on the general election ballot", while at the same time the referendum "would establish ranked-choice voting for the general election".[43]

Second, there is no obvious explanation for why four candidates advance to the ranked choice stage. Two-round contests arose in Western Europe in the late 1800s as a means to ensure that the eventual winner of a single-seat election would receive a majority of votes, and that is still a common use of two-round contests in national presidential systems (Colomer, 2004, p. 37 – 39). But there is no such guarantee if four candidates are included in the second round. And why do those four candidates proceed to an Instant-Runoff Voting election? Why not two single-vote rounds, or just one Instant Runoff election?

Crucially, these are not just questions about the politics of Alaska: Alaska has changed its method of electing *federal* representatives, and these representatives play an important role in American politics. The 2020 election, Alaska's last under SMDP, produced a Senate evenly split between the Democratic Party and the Republican Party, and Alaskan senator Lisa Murkowski in

---

[43] Quotations from the official Alaska Division of Elections 19AKBE ballot summary, 2019.

particular has distinguished herself as a crucial pivotal voter.

Initially appointed to her Senate seat in 2002 by the Governor of Alaska (who was also her father), Murkowski lost the Republican primary in 2010, but managed to become the second senator ever to win a general election as a write-in candidate. As a legislator she has cast deciding votes against the Republican party line on major legislation, and an open rift developed between Murkowski and her party's leadership during the Trump administration. Waves of media speculation that Murkowski would face a serious challenge from a Trump-aligned candidate were realized in summer 2021 when Trump endorsed Murkowski challenger Kelly Tshibaka. In the many peculiarities of the new electoral system, commentators have discerned the influence of Senator Murkowski.

Why do *four* candidates advance to the second round? Journalists reason that Murkowski, who had lost a primary from the right in 2010, has much less reason to worry about losing a top-four primary than a partisan single-winner primary.[44] Why Instant Runoff Voting in the second round? For Alaskan pollsters, Instant Runoff Voting is a mechanism for translating Murkowski's centrism into political support, and leaders of the state Republican party worry that it makes it harder for their candidates to win elections.[45] The fact that such a convoluted electoral system should so perfectly align with an incumbent's interests has prompted speculation that "Murkowski was behind it", and her "fingerprints are all over" the new system.[46] Add to these speculations the fact that Murkowski campaign staffers and volunteers heavily overlapped with the organization that promoted the election reform initiative, to the point that the main author of the ballot measure was a former Murkowski campaign worker, and this seems like a tidy example of how powerful people select electoral systems to protect their own interests (Ahmed, 2012).

But is the question really settled? There is almost no evidence about how exactly vote counts under the new system are expected to differ from vote counts under the previous system.[47] There has been at least one investigation of some of the system's properties, namely the study of its potential

---

[44] Kelsey Piper, "Alaska voters adopt ranked-choice voting in ballot initiative". *Vice News*, 2020.

[45] Liz Ruskin, "Murkowski, Who Voted Against Trump, Will Likely Survive Primary Challenge". *NPR*, 2021.

[46] James Brooks, "Local Alaska Republicans censure Sen. Lisa Murkowski, citing impeachment vote and other issues". *Anchorage Daily News*, 2020.

[47] But the appetite for speculation can hardly be overstated; as I was writing this paragraph, the *New Yorker* published a piece by Ed Kilgore called "Will Ranked-Choice Voting Help or Hurt Lisa Murkowski?"

effects on voters of color by Benjamin and Burden (2021), but given that the very existence of centrists in the United States Senate might hinge on the properties of this system, investigation of it has been quite limited. So it is worth examining its concrete vote allocation properties more closely.

In concrete terms of expected vote counts, the crucial questions are why the vote-counting method switches between rounds, and why four candidates compete in the second round. If the conventional wisdom that the system was specifically designed to boost Murkowski's electoral prospects is correct, then the answer should be "yes" to the following three questions:

1. Does having four candidates in the second round help Murkowski?
2. Does the second round's Instant Runoff Voting help Murkowski?
3. Would simple plurality or pure Instant Runoff Voting have been better for Murkowski?

In addition to these substantive questions, addressing a case like this also has methodological interest, and in the following section I impose a series of constraints that serve two purposes: they make it possible to apply the model to the extremely low information environment of Alaska state politics, and they also test the limits of the model and verify that it can work with much less rich information than we fed it in the cases of Canada and Britain.

Testing a much more constrained model

So far, the election simulation idea has performed well under reasonably favourable circumstances. In this section, I deliberately apply the model to a case about which there is much less data. I have claimed that this model needs only two pieces of information: The results of a recent SMDP election, and data on voters' sincere preference orderings. In this section, I do without one of these ingredients entirely, and use a much rougher form of the other. Rather than respondent-by-respondent preference rankings, I check how the model performs when I reconstruct preference orderings probabilistically using only crosstabulations that were publicly reported by media outlets and private pollsters. In applying the model to Alaska, I construct preference orderings from crosstabulations with support by party ID reported in three surveys. From a survey by *The New York Times*/Siena College Research Institute, October 9-14, 2020, I take the ratings of Donald Trump, the archetype of a competitive Democrat (for which I use ratings of Al Gross, a Democrat-aligned independent), and the archetype of a more conventional Republican (for which I use Dan Sullivan, Alaska's other incumbent Republican senator). From *The Alaska Survey* in the third quarter of 2016 I take ratings of a Libertarian candidate (for which I use Alaskan politician Joe Miller), and of Lisa Murkowski herself. In the actual model application, to understand the effects of the electoral system change, I will run the model on hypothetical election setups, so there is no past vote count available to seed the strategic models; instead, I will use the results of the sincere model as a seed for the strategic model.

I first use official state data to estimate the breakdown by party ID. Then I use crosstabulations from private polling data to reconstruct the probability of assigning a given Likert score to each of these candidates conditional on belonging to the Republican Party, the Democratic Party, some other party, or no party at all. I construct preference orderings for the simulated electorate by independently drawing Likert scores for each candidate according to the proportions in the survey data, propagating the survey's reported margins of error.[48] So instead of using actual reported preference orderings, I fix the number of Republicans, Democrats, Independents, and those with no affiliation, and then I draw opinions about the candidates in the proportion that they were reported by members of those groups in recent surveys.[49] Then I run the model as before. For simplicity I will not develop a new model of strategic Instant Runoff Voting, and will only discuss sincere behaviour under that system.

Before proceeding to a hypothetical election, I run two types of validation checks. It is natural to run the same sort of validation check that we performed in both previous cases: that is, to run this modeling idea on previous statewide races in Alaska and compare the simulated vote totals to the

---

[48] This is a crucial step. If the margin of error is not simulated, then there is almost no variation in preferences between model runs, and no variation in results.

[49] This resembles the approach devised by Ndegwa (1997, p. 28) for a similar problem in the case of Kenya's 1992 election, except that in Ndegwa used election results rather than summaries of surveys.

real vote totals, before changing the electoral system. But those previous two cases also present an opportunity to perform an event validity check on the idea itself. In simulating the Canadian election, we used the 2019 Canadian National Election Study, which featured individual-level thermometer scores of the different political parties, and each respondent's electoral district was identified. In the Alaskan example, we wish to challenge the model by doing without individual-level preference orderings: instead, the idea is to use publicly available polling data which only includes the frequency with which respondents of each party identification ranked the parties in each possible position. So, for example, in the Alaskan case we might have the probability that a Democrat placed Murkowski second in their preference ordering; analogously, we can use the individual-level responses in a Canadian electoral district to calculate how often, say, a Liberal placed the Conservatives second in their preference ordering. If this aggregation does not lose much information, then this pared down version of the model passes an event validity check in the example of that Canadian electoral district.

We can check how close this method gets to both the real results and to the simulation with full individual-level preference orderings by simply picking an electoral district in the Canadian election (since, for the purpose of US Senate races, Alaska is also just one electoral district), aggregating the individual-level preference orderings in that district up to rankings of the parties, and then running the model with electors whose preferences are generated according to those probabilities rather than being sample from real individual-level preference orderings in a survey.

I suggest two criteria for choosing an electoral district to perform an event validity check on. First, the district should not be completely dominated by one party over another. If there is one party that is completely dominant among local electors' preference orderings, even very different versions of the model will tend to produce the same answer, because a model would have to depart very strongly from reality in order to generate sincere preferences that cause another party to win. Second, however, there should be a clear right answer: the contest in 2019 should not have been particularly close.

Among the ridings that satisfy both conditions, I arbitrarily chose Calgary Centre. Even though the Conservative Party has won six of the seven elections in that riding since the formation of Canada's modern party system, the Liberal Party actually won the most recent election before 2019, holding the district from 2015 until 2019. This suggests that the preference structure in the electorate is not so lopsided that the model is guaranteed to be right. Second, however, one party had a solid victory in 2019, so there is a clear right answer to compare the model to: the Conservative candidate won by more than 20 percentage points.

The results of applying both ways of imputing preferences in the case of Calgary Centre are shown in Table 4. The table contains 15 model runs with each method for imputing preferences. This number of model runs is chosen arbitrarily; it does not need to be particularly large because the variation in vote totals and winners between runs is again very low.

Method	Outcome	Vote type	Liberal	Conservative	NDP	Green
Real	Vote share		30%	51%	16%	2%
Individual	Mean vote share	Sincere	23%	61%	10%	6%
Individual	Mean vote share	Strategic	40%	60%	0%	0%
Aggregated	Mean vote share	Sincere	25%	49%	15%	11%
Aggregated	Mean vote share	Strategic	49%	49%	1%	0%
Individual	Win proportion	Sincere	$\frac{0}{15}$	$\frac{15}{15}$	$\frac{0}{15}$	$\frac{0}{15}$
Individual	Win proportion	Strategic	$\frac{0}{15}$	$\frac{15}{15}$	$\frac{0}{15}$	$\frac{0}{15}$
Aggregated	Win proportion	Sincere	$\frac{0}{15}$	$\frac{15}{15}$	$\frac{0}{15}$	$\frac{0}{15}$
Aggregated	Win proportion	Strategic	$\frac{5}{15}$	$\frac{10}{15}$	$\frac{0}{15}$	$\frac{0}{15}$

Table 4: A comparison of the fully sincere and fully strategic models in Calgary Centre, 2019. The "Real" row contains the parties' actual vote shares from the real election. The "Individual" data includes the full individual-level preference orderings, as imputed in Chapter II and Chapter III. The "Aggregated" data uses the probabilistic method proposed in this section for application to statewide elections in Alaska. Both methods are used across 15 model runs. The table includes the average party-by-party vote share across those 15 model runs, and the proportion of runs that each party won, for both fully sincere and fully strategic models. The fall in the Individual share for the Conservative Party between sincere and strategic models is an artifact of taking party-by-party averages across runs, not a reflection of any strategic voting behaviour.

In Table 4, the "Individual" data includes the full individual-level preference orderings, as imputed in Chapter II and Chapter III, while the "Aggregated" data uses the probabilistic method proposed in this section for application to statewide elections in Alaska. I will continue to use these words to refer to the two types of preference imputation ideas. The Individual and Aggregated methods produce nearly identical vote shares for the Liberal Party. The Individual method expects that the Conservative Party receives about a 10 percentage point higher sincere vote share than the Aggregated method, which comes equally from the shares of the NDP and the Green Party.

Like in previous validations, the Individual method overstates the strength of the Conservative

Party and the Green Party among sincere preferences, while slightly underestimating the Liberal Party and the NDP. The real Liberal Party and Green Party vote totals are again between the sincere and strategic simulations for the Individual method in this riding, though the same is not true for the other two parties. The Aggregated method is actually similarly successful, compared to the real vote totals: the real Conservative Party vote total is actually only 2 percentage points different from the sincere and strategic vote total with the Aggregated method, and the NDP vote total is only 1% off from the sincere estimate, whereas the Aggregated method underestimates the Liberal Party and overestimates the Green Party. However, its misses are quite comparable in magnitude and type to the misses of the Individual method. In all 15 runs of the sincere model with each of the Individual and Aggregated methods, the Conservative Party won every time.

Importantly, the Individual method and the Aggregated method tell exactly the same strategic voting story: nearly every voter flees from the NDP and the Green Party (the vote shares never fall to exactly 0, but they round down to 0%), and a large majority of them instead choose to support the Liberal Party. The difference is how much support the NDP and Green Party began with. Since the sincere preferences of the Individual model represent the strength of the NDP and Green Party more closely, there is a smaller pool of electors to boost the Liberal Party when they flee from those smaller parties. This wave of support is never enough for the Liberal Party to win, so under the Individual method, the Conservative Party wins every strategic model run. In contrast, because the Aggregated method has a larger pool of sincere supporters of the NDP and Green Party, at the expense of the Conservative Party, when these voters flee to the Liberals it is enough to boost them into first place in one third of the runs.

The Aggregated method is therefore not quite as successful at matching the real results as the Individual method was, but the differences in vote share between Aggregated method and the Individual method in this example tended to be within the range of about 5% and 10%, and it produced a different winner only one third of the time. The model has similar event validity to the Individual model.

Next, I check the results of the Aggregated model in a case that does not have the rich individual-level data to compare it against: a series of recent Alaskan elections. The last three senate races had varied numbers and types of serious competitors. In 2010, Murkowski lost to Joe Miller in the Republican primary, but narrowly defeated him in the general election, and a Democratic candidate also received a substantial number of votes. In 2014, Republican Dan Sullivan narrowly beat the incumbent Democrat (in 2020 Sullivan would again narrowly beat a Democrat-aligned candidate, by a slightly wider margin). In 2016, Murkowski faced Miller again, with a Republican-leaning independent and a Democrat also winning votes. The model's expectations for the mean vote count of each competitor in these matchups is shown in Table 5.

I also report the percentage of times that each candidate would have won if instead the election had been held as just one single-stage Instant Runoff Voting election.

2010

Outcome	Ruleset	Vote type	Murkowski	Democrat	Libertarian
Real results	Single vote		39.5%	23.5%	35.5%
Mean vote share	Single vote	Sincere	40%	34%	26%
Mean vote share	Single vote	Strategic	60%	33%	6%
Win proportion	IRV	Sincere	$\frac{100}{100}$	$\frac{0}{100}$	$\frac{0}{100}$

2014

Outcome	Ruleset	Vote type	Republican	Democrat
Real results	Single vote		48%	45.8%
Mean vote share	Single vote	Sincere	56%	44%
Win proportion	IRV	Sincere	$\frac{100}{100}$	$\frac{0}{100}$

2016

Outcome	Ruleset	Vote type	Murkowski	Republican	Democrat	Libertarian
Real results	Single vote		44.4%	13.2%	11.6%	29.2%
Mean vote share	Single vote	Sincere	26%	31%	28%	15%
Mean vote share	Single vote	Strategic	20%	37%	29%	4%
Win proportion	IRV	Sincere	$\frac{45}{100}$	$\frac{55}{100}$	$\frac{0}{100}$	$\frac{0}{100}$

Table 5: The model's expectations for recent Alaskan senate elections, from 2010 through 2016. For the sincere and strategic single-vote models, the mean vote shares are shown. Then an Instant Runoff Voting contest is held using the sincere model results. The percentage shown in this row is not a vote share, but rather the percentage of model runs that each candidate won. Note that in the two-way contest, a strategic model would be the same as the sincere model, since there are only two candidates.

Table 5 shows broad success in replicating the results of recent elections using only opinion surveys, with the major exception of one type of candidate. In the 2010 election the sincere single-vote model is almost exactly correct about Murkowski's eventual vote share. The model also very closely captures the dynamic of the 2014 race, with Sullivan leading his opponent by a relatively narrow margin. Importantly, in both cases the model is easily correct about the winner of the election, and it always declares that the Instant Runoff Voting winner would have been the same as the single-vote winner (in the case of the two-candidate race, trivially so).

The pattern observed in the simulated 2016 race is more subtle. Just looking at an average of vote proportions, the model appears to badly underestimate Murkowski. But the reason that the average is so low is that the results are cleanly divided into two equilibria. In the more common one, the conventional Republican and the Democrat randomly begin with a slightly edge over Murkowski in the sincere vote totals. Through vote redistribution, the Republican always defeats the Democrat. In the less common one, Murkowski edges out the Democrat in the sincere vote distribution. When she does so, she invariably faces the Republican, and invariably wins through vote redistributions.

The one consistent problem is that the model seriously underestimates the Libertarian candidate. In the 2010 race this does not change the result, though the model expects the Democrat to be Murkowski's main challenger when in reality it was Libertarian-aligned Republican Joe Miller who placed second. In the 2016 election, it may affect the prominence of the more common equilibrium.

This provides an initial answer to a few of the substantive questions. Even in a model that underestimates Murkowski enough to believe that she would almost never have won a single-vote election that she in fact did win, an Instant Runoff Voting model — specifically, a simple simulation of the Instant Runoff Voting ruleset using publicly available data of the preference orderings reported by Alaskans — assigns her a victory almost half the time. This strongly bolsters the idea that Instant Runoff Voting is advantageous to Murkowski.

Also, it might help to explain why the number of competitors in the second round is four. Despite the fact that the model was slanted towards the Democrat and away from Murkowski, it never produces a win for either the Democrat or the Libertarian candidate under Instant Runoff Voting, even though it asserts that the Democrat out-competes Murkowski in a single-vote contest. The model suggests that Murkowski sits at a special place in the Alaskan political spectrum where the major types of competitors in Alaskan elections would mostly lose to Murkowski in an Instant Runoff Voting contest, but might defeat her in a single-vote contest.

Understanding the effects of Alaska's new electoral system

We can now apply the model to a stylized version of Alaska's 2022 race. In addressing the three substantive questions about the new electoral system, there is one structural limitation on the model: the restriction to a four-way race is only relevant in a field of more than four candidates. If four or fewer candidates compete, then every candidate will win the first round, and Instant Runoff Voting converges to a single-vote election as the number of candidates decreases to two. But typical United States Senate races in Alaska in recent years have featured between two and four major competitors. So instead I consider the smallest race — a five-candidate contest — in which the new electoral system could make an important difference in the results.[50]

There is also a scientific reason to construct the most challenging realistic field for Murkowski: to do so plays against the claim that the new electoral system makes her more likely to be elected. For the archetype of a Trump-aligned challenger, I set the very high bar that the challenger manages to become perfectly aligned with Trump, and I give them Trump's own approval ratings among Alaskans. I label this candidate Tshibaka, since Kelly Tshibaka has been endorsed by Donald Trump, though this a stylization of reality — here Tshibaka is standing in for a candidate who receives *exactly* the same Likert ratings as Trump. For the other three preference orderings I use the reported preferences about a highly successful Alaskan Republican, Democrat, and Libertarian. Table 6 shows the proportion of model runs in which each competitor wins under the single-vote *status quo* (both sincere and strategic), Alaska's new electoral system, and a Borda race (both sincere and strategic).

---

[50] Such a large race appears plausible. In polling conducted in the very early stages of the 2022 election, Lisa Murkowski already appeared to face a Trump-aligned challenger (Kelly Tshibaka), and a strong Democratic-aligned independent (Al Gross). At the same time, a series of Libertarian and more conventional Republican politicians were actively signaling that they might enter the race.

Candidate	Single-vote (sincere)	Single-vote (strategic)	New system	Borda (sincere)	Borda (strategic)
Tshibaka	$\frac{919}{1000}$	$\frac{1000}{1000}$	$\frac{880}{1000}$	$\frac{0}{100}$	$\frac{100}{100}$
Murkowski	$\frac{0}{1000}$	$\frac{0}{1000}$	$\frac{115}{1000}$	$\frac{100}{100}$	$\frac{0}{100}$
Democrat	$\frac{81}{1000}$	$\frac{0}{1000}$	$\frac{2}{1000}$	$\frac{0}{100}$	$\frac{0}{100}$
Republican	$\frac{0}{1000}$	$\frac{0}{1000}$	$\frac{3}{1000}$	$\frac{0}{100}$	$\frac{0}{100}$
Libertarian	$\frac{0}{1000}$	$\frac{0}{1000}$	$\frac{0}{1000}$	$\frac{0}{100}$	$\frac{0}{100}$

Table 6: The results of an imagined Alaska Senate election under Alaska's new electoral system. Each fraction is a number of model runs in which the candidate won, divided by the total number of model runs of that type. The single-vote and IRV results are based on 1000 model runs, and the Borda results are based on 100 runs.

Table 6 shows that the electoral system change indeed favours Murkowski. A challenger who somehow manages to perfectly align themselves with Donald Trump wins an overwhelming proportion of the single-vote races, only being overtaken by the Democrat (due to right-of-center vote-splitting) less than one in every ten times. High levels of strategic voting make that candidate essentially unbeatable. And yet, under the new electoral system, the Democrat's successes flow suddenly to Murkowski: even against a challenger who is perfectly identified with Trump, Murkowski wins more than one in ten contests under the new system, up from no contests under the old system.

In answer to question 1, Murkowski wins a large number of Instant Runoff Voting contests among four candidates, so that number appears relatively beneficial for her. It is clear why a two-way runoff would be harmful for Murkowski: even if she made such a runoff, she would then have to directly defeat a candidate aligned with Trump. But at the same time, it is hard to see how she could be harmed by facing a fifth or sixth candidate in the second round. So from Murkowski's perspective, a three-way runoff would probably be worse, but a runoff between five or more candidates would perhaps be just as good.

The answer to question 2 is comparatively simple: Instant Runoff Voting in the second round is unambiguously to Murkowski's advantage. And turning to question 3, the results are similarly unambiguous that the new two-round system is better for Murkowski than plurality would have been. However, it is not at all clear what distinguishes the current system from a pure Instant

Runoff Voting election. In fact, in sincere simulations, these two systems have identical results: a single-round Instant Runoff Voting election simply begins with the weakest of the five candidates being eliminated through Instant Runoff Voting rather than through a single-vote round, and that candidate is always the same. It remains unknown why the designers of the electoral system would construct a brand new merger of plurality and ranked choice rules when they appear to produce the same results as a common ranked choice ruleset alone.

And the model offers another major challenge: Instant Runoff Voting does not appear to be the most favourable system for Murkowski. If peoples' first-, second-, and third-place choices are added up in a Borda ballot, Murkowski has such a commanding lead that she always leads a Borda election where every elector casts a sincere ballot. If those who engineered this electoral system indeed believed that Instant Runoff Voting was beneficial to the incumbent, they may have erred in the ranked choice allocation method they selected.[51] But of course, it is exactly these sorts of detailed comparisons that would not have been possible without a model of what might happen under different systems.

By studying the effects of a new electoral system on the politics of Alaska, we can draw clear connections to the nation's politics. Alaska is represented by one of the senate's pivotal voters, and federal legislation could be meaningfully changed by her defeat. But in a climate where new systems for electing federal representatives are being slowly but steadily rolled out across the nation, the main use of this model should be to discern the possible effects of different electoral systems in the specific politics of different states *before* they are instituted.

This exercise has also shown that even when the available data are tremendously sparse, the results resemble reality. This was just about the worst-case scenario, since large polls in Alaska are notoriously infrequent. The model used only a few hundred respondents to a handful of private surveys, in which the real preference orderings of individual respondents are not known and can only be imputed from crosstabulations, and there are no real results of a previous election available to inform the electors' judgments of the competitiveness of each candidate. And yet, despite

---

[51] Supposing that sufficiently few people vote strategically. If electors perform strategic calculations and select their best-response ballot, then the victory shifts uniformly to the Trump-aligned candidate.

performing worse than it did when much higher-quality data was available, the model still generates results under the previous electoral system that roughly resemble the real results, and these results still provide insights into the effects of a new electoral system.

CHAPTER V

Conclusion

Chapter II of this dissertation studied a model of elections where electors iteratively signal their intended vote choice, simultaneously update their intended vote choice based on that signal, signal their updated choices, and so on. To obtain an election result from that type of model, I introduced the notion of an iterative equilibrium, which consists of all of the vote totals that will arise as long as the model runs. If one vote total leads to another vote total with certainty, then iterative equilibrium is inevitable. I also proved several constraints on how complicated electors' behaviour can be at iterative equilibrium. If electors respond deterministically to vote totals, then their vote choices cannot change in ways that are much more complicated than the changes in the aggregate vote totals.

In Chapter II I also introduced a type of voting that I call threshold pivotality. In this extremely minimalistic voting heuristic, electors check if multiple parties are within a threshold τ of the party that is expected to place first. If so, they vote for whichever party they most-prefer in that near-tie. If not, they vote for their sincerely most-preferred party. I applied threshold pivotality to the 2019 Canadian election, and with reasonable values of τ, threshold pivotality correctly assigned about 95% of the seats in parliament to the party that actually won that district. I also showed that it often but not always produces two party competition in the districts of that SMDP election, and in that example I demonstrated that very different ways of resolving the districts with iterative equilibria produced very similar results.

Chapters III and IV of the dissertation built on the setting of Chapter II, and in these chapters I developed the first method to estimate how a change to a *specific* alternative electoral system might affect the elections of a *specific* country, accounting for how electors might respond to that institutional change. The method closely replicates how the results of real elections shifted after the electoral system change in New Zealand, as well as the results of actual SMDP elections in Canada, Britain, and Alaska. And while it relies on and extends the explicit computations of pivotal probabilities, the results are very similar to the results of applying a tremendously simple heuristic.

In the case of Canada, the SMDP model reinforced the conventional idea that strategic voting is bad for the Conservative Party. The PR results supported the expectation that Canada's smaller leftist parties would benefit from a switch to a more proportional system, but contained a harsh warning about the dangers of proportional systems that have similar strategic incentives to the ones that already depress those parties' vote shares under SMDP. It also hinted that PR might not be as bad for the Conservative Party as is commonly thought. The model sharply disagreed with the almost-universally accepted claim that Ranked-Choice Voting would entrench Liberal Party dominance, instead showing a much more complicated picture in which a Liberal plurality would face a weakened Conservative Party but substantially stronger leftist parties. The model called into question the behaviour of Canada's left-of-centre parties during the election reform debate, when by banding together they might have been able to institute either PR or Ranked-Choice Voting; Ranked-Choice Voting especially appears beneficial to all of them.

Applying the same model to the highly similar case of Britain showed that some of these findings might be generalizable, and suggest areas for future empirical or observational testing. The model is largely at odds with the widespread belief that the two parties that tend to dominate Westminster politics are animals mostly of SMDP. In particular, if the wave election of 2019 had been held under another system, it seems the Conservatives would still have dominated parliament. And the fact that no minor parties in Britain benefited from strategic Ranked-Choice Voting the way that Canada's Green party did suggests that there is something special about that party — perhaps the fact that it is the only competitive niche issue party in either country — that benefits from strategic voting under Ranked-Choice Voting.

Finally, applying the model to the particularly strange new electoral system in Alaska showcased how a model of the specific effects of electoral systems can help political scientists make sense of why systems are designed and instituted. In this case, the model was able to identify potential upsides for the incumbent in each component of the convoluted system, which is consistent with the widespread belief that supporters of this incumbent were involved in engineering and implementing the new system.

Across the three cases, a few particularly dramatic changes stand out. On the whole, changing the electoral system matters more in Canadian election results than in British election results. The Proportional Representation model suggests that, if people vote sincerely, Canada's two largest parties would be substantially hurt, with the Liberal Party losing about 15% of the seats in parliament and the New Democratic Party and Green Party commensurately gaining. In contrast, British results would change fairly modestly, and the main difference is a jump in Liberal Democrat holdings to the tune of about 10% of the seats in parliament. If voters are instead strategic, the results in Canada converge to the real election results, with the exception that about 10% of the seats in parliament flow from the Bloc Québécois to the New Democratic Party. In Britain, the strategic voting model under Proportional Representation allocates the gains that were made by the Liberal Democrats in the sincere Proportional Representation model to the Labour Party instead.

The most striking finding of the sincere Ranked Choice Voting model across both Canada and Britain is how exceptionally close both are to the real election results under Single-Member District. The British results in particular under Ranked-Choice Voting with sincere electors are almost identical to the real election results. This is somewhat less true of the Canadian results, with several parties receiving almost their real vote totals, and the Conservatives giving up a modest number of seats to the New Democratic Party. Taken together with the fact that respondents in both countries were more likely to report preferring the Conservative parties more than voters were to actually cast votes for them — much moreso in Britain than in Canada — this suggests that Ranked Choice Voting with mostly sincere voters would reduce the seat count of both countries' Conservative parties. The Alaskan results, however, heavily depend on which type of sincere Ranked Choice Voting is being simulated. The sincere Borda count results showed a dramatic swing towards Murkowski, far more than the sincere Final-4 simulations.

If instead voters are expected to be highly strategic under Borda count, then the potential effects of this system were much larger in Canada than in either Britain or Alaska. Canada's Green Party received about 15% of seats in parliament in a typical strategic Borda count model, whereas in the real election it received less than 1% of the seats (though it also received nearly 7% of the votes), and the Liberal Party saw a similar drop of about 10% of the seats in parliament. In contrast, in Alaska the strategic Borda count model saw Tshibaka winning every run, and even the switch to Final-4 only featured Murkowski wins in about 10% of model runs.

The substantive results in these chapters suggest tendencies that can be further tested using observational data. But the ideal substantive application is not to apply the model after electoral systems have already been implemented, or after electoral system change has already failed. The history of attempts to change electoral systems away from SMDP is a history of non-action: one filled with attempts to adopt a more proportional system, broadly perceived (whether correctly or incorrectly) as being in the public interest, but easily quashed by the self-interest of powerful politicians. The ironic core of this dissertation is that, evidently, incumbents may often fail to enact an electoral system change that would actually be in their interest. Lacking any obvious tool to judge what would help and what would hurt them, these party leaders choose to stick with the devil they know. The method introduced in this dissertation provides an empirical baseline for future election reform debates. In doing so, it also illustrates how computer models can be used to better understand elections.

APPENDIX A

Riding-level Successes

Table A.1 shows the riding-level success rate of each of the models in both Canada and Britain.

Country	Model	Seat allocation successes	Riding-level successes
Canada	Sincere	92.5%	81.1%
	Strategic	91.4%	78.9%
	30% strategic	95.3%	82.6%
	$\tau = 5\%$	94.7%	79.9%
Britain	Sincere	86.6%	84.9%
	Strategic	93.9%	89.0%
	30% strategic	88.9%	86.2%
	$\tau = 5\%$	87.3%	84.7%

Table A.1: Mean percentage of seats and ridings each model assigns correctly across 25 runs.

APPENDIX B

A Step-By-Step Example of the Model

In this section I use prose, examples, and pseudocode to walk, step-by-step, through a simple concrete example of how the model produces results. I base the example in two adjacent electoral districts in Michigan, shown in Figure B.1a: the 11th Congressional District, a reasonably competitive seat that typically leans Republican, and the 12th Congressional District, an extremely un-competitive district that tends to vote Democratic by about a ratio of about $2:1$.

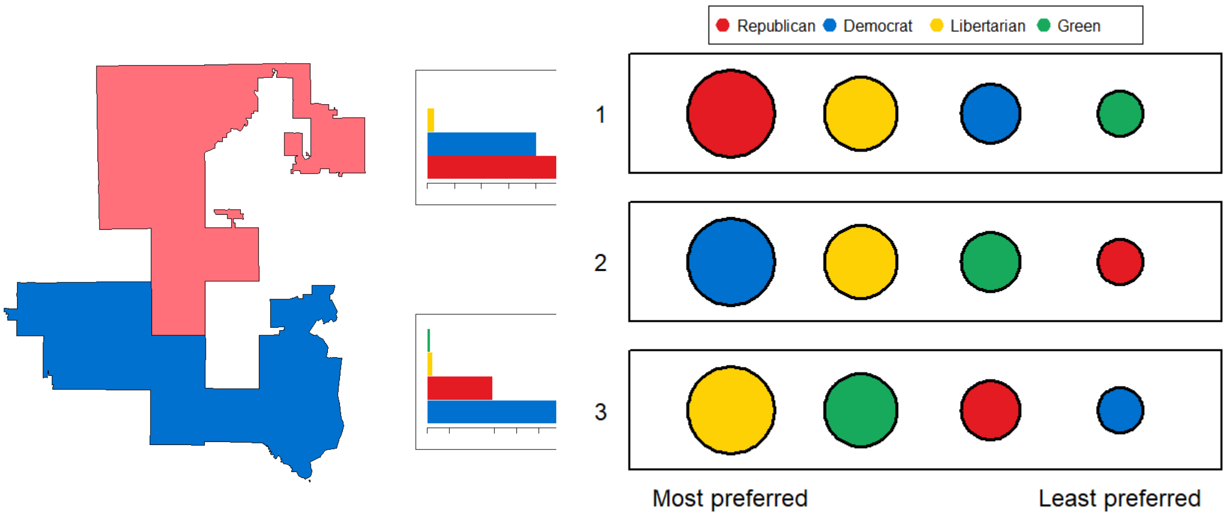

(a)

(b)

Figure B.1: An example of how the model works using real districts with hypothetical preferences. Subfigure B.1a shows the two districts in which the model example is set. Subfigure B.1b shows example preference orderings for electors in a district.

To keep the numbers manageable and concrete, imagine that the districts contain only 50 voters. In the model, the preference orderings of these voters would be sampled with replacement from survey data until the population size was reached. But for simplicity, imagine that every elector in the district has one of the preference orderings represented in Figure B.1b. These rankings include the Democratic Party, the Republican Party, the Libertarian Party, and the Green Party. Of the 50 people in each of our imagined districts, suppose that 20 people have the preference ordering in the top row of that figure (so, Republicans $\succ$ Libertarians $\succ$ Democrats $\succ$ Greens), 20 have the preference ordering represented in the second row, and only 10 have the preference ordering represented in the third row.

Already we have specified enough to state the results of the sincere model. In such a district, the Republicans would receive 20 votes, the Democrats would receive 20 votes, and the Libertarians would receive 10 votes.

To understand how the strategic model works, first take up the case of Congressional District 11 in the year 2016. In the 2014 election, the republican candidate received a vote share of 0.559, the Democrat 0.405, and the Libertarian 0.03. There was no Green Party candidate, so they are assigned a previous vote share of 0. These vote shares are then scaled to the population size we arbitrarily selected (in the model runs for this dissertation, the population size was one tenth of the real district's population size, which was usually in the thousands, though the scaling factor was smaller for PR runs in exceptionally large districts). So in this case, electors begin with the expectation that 28 voters will begin supporting the Republicans, 20 will support the Democrats, and 2 will support the Libertarians.

This completes the setup of the strategic model. Next, electors iteratively compute their best response to the expected vote choice. As explained in the text, the expected utility of casting a vote for party i is given by

$$u(i) = \sum_{\substack{j=1 \\ i \neq j}}^{M} \mathbb{P}(v_i = v_j \wedge v_j > v_k) \cdot (u_i - u_j) + c \cdot \mathbb{P}(v_i = v_{j+1} \wedge v_j > v_k) \cdot (u_i - u_j)$$

for all other parties k, where u_i and u_j are the utility obtained from the victory of i or j respectively and v_i and v_j are their expected vote totals, M is the number of parties in the election, and c is the probability that a party in a first-place tie will win the election, taken here to be $\frac{1}{2}$.

First let us compute the pivotal probabilities in this expression. Take the example of an elector whose preference ordering is of the first type, so they most prefer the Republicans. This means that they count themselves among the 28 expected Republican voters. But they are in the act of computing their own best vote choice, so it does not make sense to include their own vote in the expected vote count. So the first thing they do is to subtract out their own vote, leaving 27 expected Republican votes, 20 expected Democratic votes, and 2 expected Libertarian votes.[1] The Skellam distribution is only defined when its second and third parameters are positive integers, so even though the Green Party is expected to receive 0 votes, we force it to have a single expected voter (we can be quite certain that this will not change any results).

Focus first on the voter's expected utility of voting for the Republican Party. The probability

---

[1] Actually, in a Poisson voting game, the conceit is that the population size itself should be drawn from a Poisson distribution where the paremeter, in this case, would be 50. But in large electorates this does not affect the substantive model results, so in this simple example I assume a fixed electorate of 50.

p_{TIE} that the vote total v for the Republican Party will equal the vote total for one of the other parties is the following (ignoring, as mentioned in the body of the text, the possibility of ties between more than two parties, which I take to be negligible *for elections that have realistic levels of voter turnout*, so that the following is true by the assumption of mutual exclusivity):

$$p_{TIE} = \mathbb{P}(v_{Republican} = v_{Democrat}) + \mathbb{P}(v_{Republican} = v_{Libertarian}) + \mathbb{P}(v_{Republican} = v_{Green})$$

Each term is the probability that there is difference zero between two Poisson random variables. As discussed in Mebane et al. (2019), that difference is Skellam distributed, in our example as follows:

$$p_{TIE} = \mathcal{S}(0; 27; 20) + \mathcal{S}(0; 27; 2) + \mathcal{S}(0; 27; 1)$$

This expression yields a probability that the elector has an opportunity to break a tie between the Republican Party and any other party. But of course, breaking a tie that is not a first-place tie is not useful to the elector; the only situation in which this provides the opportunity to cast a pivotal vote is if the two parties under consideration have more expected votes than any other party. This too can be explicitly calculated using the Skellam distribution. Consider for a moment just the leftmost term in the sum above: the probability that the elector is able to break a tie between the Republican and Democratic parties. We are concerned with the situation where their vote counts are equal, so if one has more expected votes than all other parties, then so does the other. So the probability that these two are leading is just the CDF of the Skellam distribution down to a vote difference of 1. That is, it is the probability that the Republican Party has one more vote than (say) the Libertarian Party, plus the probability that it has two more votes than the Republican Party, and so on, up to the largest possible vote difference. So, keeping with the standard (simplifying) assumption of the independence of vote counts, the probability that the Republican Party leads both parties other than the Democratic party is

$$p_{FIRST} = \sum_i \mathcal{S}(i; 27; 2) \cdot \sum_i \mathcal{S}(i; 27; 1)$$

where i indexes the other electors. The probability that a Republican-Libertarian tie or a Republican-Green tie is in first place is computed similarly. Finally, a voter is also concerned with the probability p_{MAKE} that they have the opportunity to create a tie between two parties:

$$p_{MAKE} = \mathbb{P}(v_{Republican} = v_{Democrat} - 1) + \mathbb{P}(v_{Republican} = v_{Libertarian} - 1) + \mathbb{P}(v_{Republican} = v_{Green} - 1)$$

Putting all of these probabilities together, according to mutual exclusivity and independence, the pivotal probability p of our voter actually casting a vote for the Republican Party in this SMDP election is:

$$p = (p_{TIE} + p_{MAKE}) \cdot p_{FIRST}$$

and the pivotal probability for each other party is computed in the same way. Performing those computations in our example, the pivotal probabilities (denoting the pivotal probability of casting a vote for each party as p with the first letter of that party's name as the subscript) are as follows, rounding to the nearest two decimal places:[2]

$$\begin{bmatrix} p_R & p_D & p_L & p_G \end{bmatrix} \approx \begin{bmatrix} 0.03 & 0.03 & 3.6 \cdot 10^{-13} & 5.0 \cdot 10^{-16} \end{bmatrix}$$

---

[2] Crucially, I do not round pivotal probabilities in the model. Even when they are exceptionally small — and they can have exponents in the negative hundreds — it is essential that they not be rounded to zero, in which case electors might find equal strategic utilities for casting two votes which in reality have unequal pivotal chances.

Note that the Republican and Democratic pivotal probabilities are not actually exactly equal, though rounding to two decimal places makes them appear so; really, the Republican pivotal probability is very slightly higher, because the probability that the Republican Party lands in a first place tie with either the Libertarian or Green party is higher than the probability that the Democratic Party reaches a first place tie with those parties. It is also worth noting that in a real electorate these numbers would be tremendously smaller than they are in a 50 person electorate. But all that matters is whether or not their ordering is variant under multiplication by small natural numbers, which does not change in realistic applications whether a population has hundreds or thousands of electors.

Carrying through the example of this elector's vote decision, how do they use the probability of having a pivotal vote to calculate the expected utility of every vote? Consider the evaluation of the probability that the Republican candidate will be either exactly tied with the Democratic candidate, or be one vote behind, and that both parties have more votes than all other parties. This is the dominant term in the full pivotal probability of a vote for the Republicans: our elector has about a 3% chance of casting a pivotal vote between those alternatives. So the utility of that vote is the difference between their sincere utility for the Republicans and their sincere utility for the Democrats, times the probability that it is decisive, or

$$(3 - 1) \cdot 0.03 = 0.06$$

Indeed, because this term is so much larger than both other terms in the sum of the full strategic utility for voting Republican (that is, the utility of creating or breaking a first-place tie between Republicans and the Libertarians, or between the Republicans and the Greens), this is also the rounded strategic utility of voting Republican. Conversely, the strategic utility of this voter casting a vote for the Democrats is about -0.06. Votes for the Libertarians and Green candidates are very small negative numbers. So our elector will choose, in this iteration, to vote Republican.

All electors perform these calculations simultaneously, in one iteration. At the end of the iteration, in our example we find that voters with the first preference ordering will all vote Republican, voters with the second preference ordering will all vote Democratic, and voters with the third preference ordering find that their highest expected strategic utility is to vote for the Republicans. This completes one iteration.

At the end of an iteration, the electors announce their vote choices: there are now expected to be 30 votes for the Republican candidate, 20 votes for the Democratic candidate, and no votes for either of the other two parties. Now the electors repeat all of the calculations above, but with these new numbers instead of the original distribution. The Republican party, now both preferred and supported by a majority, will hold on to all of its voters, since all 30 will find their highest strategic utility is to cast a vote for the Republicans. The Democratic party likewise retains its 20 voters. No elector has a unilateral incentive to change their vote, so this is a mutual best response equilibrium. In all future iterations, the computations will be identical, so the model has hit a fixed equilibrium and this is the result of our imagined election in Michigan's 11th congressional district.

Do the two largest parties always drag electors inexorably towards themselves, as represented in Figure B.2a, or can the situation in B.2b also occur, with third parties also receiving votes?

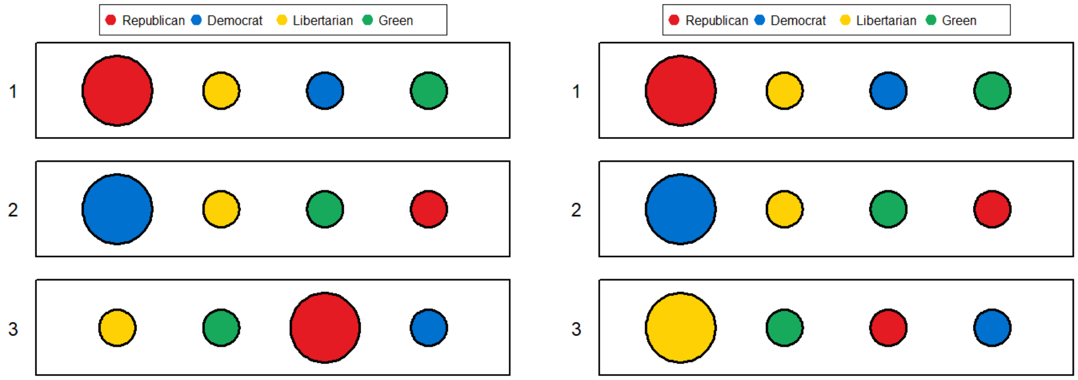

(a) (b)

Figure B.2: Example of strategic voting in two real districts with hypothetical preferences. Voters in this district vote (as symbolized by a larger circle) for their favourite candidate that has a meaningful chance of winning. In Figure B.2a the tight race between Democrats and Republicans drives voters to one of those options. In the less close election, voters have less incentive to vote for a party other than their most preferred option, resulting in the situation shown in B.2b.

To see one reason why third party votes do occur, consider the less competitive district, Michigan's 12th. Imagine that we have exactly the same distribution of preferences in the electorate: 20 with the first preference ordering, 20 with the second, and 10 with the third. And let us stylize the elections in this district by supposing that the Libertarian Party is nearly as popular here as the Republican Party (this exaggerates the situation in Michigan's 12th district, but in some years third parties do give the less popular of the two biggest parties a close race for second in places like Alaska and Brooklyn, and of course this is extremely common across parliamentary SMDP countries). So imagine that the initial vote count is an expected 25 votes for the Democrats, 13 votes for the Republicans, and 12 votes for the Libertarians.

With these initial expectations, the voters with the first preference ordering will support the Republicans, and voters with the second preference type will support the Democrats. But voters with the third preference ordering will find that the pivotal probability of voting Republican is not sufficiently bigger than the pivotal probability of voting Libertarian to motivate compromising and voting Republican. Rather, their highest expected strategic utility comes from a Libertarian Party vote. So the expected vote total after one round is 20 for the Republicans, 20 for the Democrats, and 10 for the Libertarians. However, we have already seen what happens when this population of electors encounters this vote total: in the next iteration, the voters consolidate to 30 votes for the Republican and 20 votes for the Democrats, which is the same Nash equilibrium seen in the last congressional district.

Another important way to get votes for parties other than the two most competitive is to allow preference orderings to be non-strict. If both of the two most competitive parties are tied for last, the elector will not be able to cast a vote for either. Strategic payoffs for an elector's least-preferred party are coerced to $-\infty$, because it is never strategic to boost one's least-preferred party over a more preferred alternative.

This model illustration has not covered two features of the model. First, this only illustrates a run of the model under SMDP; I describe elsewhere how the model is modified for other electoral systems. Second, it is not always true that the model reaches one fixed equilibrium for all eternity. This situation is discussed at length in Chapter II of this dissertation, in which I showed that for the main case under consideration, the particular stopping rule is not important. So throughout I will use the election day rule with a default of 0 days: as soon as iterative equilibrium is identified, the result is reported.

BIBLIOGRAPHY

Acemoglu, D., Egorov, G., and Sonin, K. (2009). Equilibrium refinement in dynamic voting games. *Paper presented at the Massachusetts Institute of Technology Department of Economics Working Paper Series*.

Ahmed, A. (2012). *Democracy and the Politics of Electoral System Choice: Engineering Electoral Dominance*. Cambridge University Press.

Aldrich, J. H., Blais, A., and Stephenson, L. B. (2018). Strategic voting in political institutions. In Stephenson, L. B., Aldrich, J. H., and Blais, A., editors, *The many faces of strategic voting*, chapter 1, pages 1–27. University of Michigan Press.

Alvarez, R. M. and Nagler, J. (2000). A new approach for modelling strategic voting in multiparty elections. *British Journal of Political Science*, 30:57–75.

Baker, A. and Scheiner, E. (2007). Electoral system effects and ruling party dominance in japan: A counterfactual simulation based on adaptive parties. *Electoral Studies*, 26:477–491.

Bechtel, M. M., Hangartner, D., and Schmid, L. (2016). Does compulsory voting increase support for leftist policy? *American Journal of Political Science*, 60(3):752–767.

Bendor, J., Diermeier, D., Siegel, D. A., and Ting, M. M. (2011). *A behavioral theory of elections*. Princeton University Press.

Benjamin, A. and Burden, B. C. (2021). Consequences of final-five voting for communities of color.

Benoit, K. (2004). Models of electoral system change. *Electoral Studies*, 23:363–389.

Benoit, K. (2007). Electoral laws as political consequences: Explaining the origins and change of electoral institutions. *Annual Review of Political Science*, 10:363–390.

Benoit, K. and Schiemann, J. W. (2001). Institutional choice in new democracies: Bargaining over hungary's 1989 electoral law. *Journal of Theoretical Politics*, 13:158–182.

Binmore, K. (1990). *Essays on the foundations of game theory*. Basil Blackwell, 1 edition.

Black, D. (1958). *The theory of committees and elections*. Cambridge University Press.

Blais, A. and Carty, R. K. (1991). The psychological impact of electoral laws: Measuring duverger's elusive factor. *British Journal of Political Science*, 21:79–93.

Blais, A., Héroux-Legault, M., Stephenson, L., Cross, W., and Gidengil, E. (2012). Assessing the psychological and mechanical impact of electoral rules: A quasi-experiment. *Electoral Studies*, 31:829–837.

Blumenau, J., Eggers, A. C., Hangartner, D., and Hix, S. (2017). Open/closed list and party choice: Experimental evidence from the uk. *British Journal of Political Science*, 47(4):809–827.

Blumenau, J. and Hix, S. (2015). What would the election look like under pr? *LSE blog*.

Boix, C. (1999). Setting the rules of the game: The choice of electoral systems in advanced democracies. *American Political Science Review*, 93:609–624.

Bowler, S. and Donovan, T. (2008). Election reform and (the lack of) electoral system change in the us. In Blais, A., editor, *To keep or to change first past the post? The politics of electoral reform*, chapter 3, pages 90–111. Oxford University Press.

Brams, S. J. and Herschbach, D. R. (2001). The science of elections. *Science*, 292:1449.

Brandenburg, H. (2019). What would the british parliament look like under proportional representation? *The conversation*.

Bryden, J. (2015). Justin trudeau denies electoral reform favours liberal party. *The Canadian Press*.

Chhibber, P. K. and Kollman, K. (2004). *The formation of national party systems: Federalism and party competition in Canada, Great Britain, India, and the United States*. Princeton University Press.

Cochrane, C. (2017). The imaginary worlds of electoral system reform. In Potter, A., Weinstock, D., and Loewen, P., editors, *Should we change how we vote? Evaluating Canada's electoral system*, chapter 5, pages 42–48. McGill-Queen's University Press.

Colomer, J. M. (2004). Introduction: The strategy and history of electoral system choice. In Colomer, J. M., editor, *Handbook of Electoral System Choice*, chapter 1, pages 3–80. Palgrave Macmillan.

Coughlin, P. J. (1992). *Probabilistic voting theory*. Cambridge University Press.

Cox, G. W. (1994). Strategic voting equilibria under the single nontransferable vote. *American Political Science Review*, 88:608–621.

Cox, G. W. (2015). Electoral rules, mobilization, and turnout. *American Review of Political Science*, 18:49–68.

Cusack, T. R., Iversen, T., and Soskice, D. (2007). Economic interests and the origins of electoral systems. *American Political Science Review*, 101:373–391.

de Borda, J.-C. (1781). Mémoire sur les élections au scrutin. *Mémoires de l'Académie Royale*.

de Condorcet, M. (1785). *Essai sur l'application de l'analyse à la probabilité des décisions rendue à la pluralité des voix*. L'imprimerie Royale, Paris.

Demichelis, S. and Dhillon, A. (2010). Learning in elections and voter turnout. *Journal of Public Economic Theory*, 12(5):871–896.

der Straeten, K. V., Laslier, J.-F., Sauger, N., and Blais, A. (2010). Strategic, sincere, and heuristic voting under four election rules: an experimental study. *Social Choice and Welfare*, 35:435–472.

Downs, A. (1957). *An Economic Theory of Democracy*. Harper Collins.

Dutil, P. (2017). Why trudeau abandoned electoral reform. *Literary Review of Canada*.

Duverger, M. (1951). *Les Partis Politiques*. Librarie Armand Colin.

Eggers, A. C. (2015). Proportionality and turnout: Evidence from french municipalities. *Comparative Political Studies*.

Eggers, A. C., Rubenson, D., and Loewen, P. J. (2021). Who votes more strategically? evidence from canada. *Journal of Politics*, Forthcoming.

Evci, U. and Kaminski, M. M. (2021). Shot in the foot: unintended political consequences of electoral engineering in the turkish parliamentary elections in 2018. *Turkish Studies*, 22:481–494.

Feller, W. (1968). *An Introduction to Probability Theory and Its Applications*. Wiley, 3 edition.

Fey, M. (1997). Stability and coordination in duverger's law: A formal model of preelection polls and strategic voting. *American Political Science Review*, 91(1):135–148.

Fieldhouse, E., Green, J., Evans, G., Mellon, J., Prosse, C., de Geus, R., and Bailey, J. (2019). 2019 bes pre-election random probability survey.

Fiva, J. H. and Folke, O. (2014). Mechanical and psychological effects of electoral reform. *British Journal of Political Science*, 46:265–279.

Flis, J., Słomczyński, W., and Stolicki, D. (2018). Pot and ladle: Seat allocation and seat bias under the jefferson-d'hondt method. *ArXiV*.

Gelman, A. and King, G. (1994). A unified method of evaluating electoral systems and redistricting plans. *American Journal of Political Science*, 38(2):514–554.

Golder, M. (2003). An evolutionary approach to party system stability. *Paper presented at the Annual Meeting of the Midwest Political Science Association, Chicago*.

Gosnell, H. F. (1930). *Why Europe Votes*. University of Chicago Press.

Heller, A. L. (2021). Public support for electoral reform: The role of electoral system experience. *Electoral Studies*, 72:179–208.

Hicken, A. and Pundit, B. (2016). The effects of thailand's proposed electoral system. *Thai Data Points*.

Hix, S. (2004). Electoral institutions and legislative behavior: Explaining voting defection in the european parliament. *World Politics*, 56(2):194–223.

Jackson, W. K. (1993). The origins of the electoral referendums. In McRobie, A., editor, *Taking it to the people?*, chapter 1, pages 15–23. Hazard Press.

Karp, J. A., Vowles, J., Banducci, S. A., and Donovan, T. (2002). Strategic voting, party activity, and candidate effects: testing explanations for split voting in new zealand's new mixed system. *Electoral Studies*, 21:1–22.

Kavner, J. and Xia, L. (2021). Strategic behavior is bliss: Iterative voting improves social welfare. *ArXiV working paper*.

Kawai, K. and Watanabe, Y. (2013). Inferring strategic voting. *American Economic Review*, 103(2):624–662.

Kedar, O. (2014). *Voting for Policy, Not Parties: How Voters Compensate for Power Sharing*.

Kim, S., Taber, C. S., and Lodge, M. (2010). A computational model of the citizen as motivated reasoner: Modeling the dynamics of the 2000 presidential election. *Political Behavior*, 32(1):1–28.

Kloeckner, B. R. (2020). Bad cycles and chaos in iterative approval voting. *ArXiV working paper*.

Kollman, K., Miller, J. H., and Page, S. E. (1992). Adaptive parties in spatial elections. *American Political Science Review*, 86(4):929–937.

Kollman, K., Miller, J. H., and Page, S. E. (1997). Landscape formation in a spatial voting model. *Economics letters*, 55(1):121–130.

Kollman, K., Miller, J. H., and Page, S. E. (1998). Political parties and electoral landscapes. *British Journal of Political Science*, 28(1):139–158.

Konishi, H. and Ray, D. (2003). Coalition formation as a dynamic process. *Journal of Economic Theory*, 110:1–41.

Krook, M. L. (2018). Electoral systems and women's representation. In Herron, E. S., Pekkanen, R. J., and Shugart, M. S., editors, *The Oxford Handbook of Electoral Systems*, chapter 9, pages 175–192. Oxford University Press.

Laslier, J.-F. (2009). The leader rule: A model of strategic approval voting in a large electorate. *Journal of Theoretical Politics*, 21:113–136.

Latner, M. S. and Roach, K. (2011). Mapping the consequences of electoral reform. *California Journal of Politics and Policy*, 3:1–22.

Laver, M. and Sergenti, E. (2012). *Party competition: An agent-based model*. Princeton University Press, 1 edition.

LeDuc, L. (1999). New challenges demand new thinking about our antiquated electoral system. In Milner, H., editor, *Making every vote count: Reassessing Canada's electoral system*, chapter 4, pages 63–78. Broadview Press.

Lepore, J. (2020). *If then: How the Simulmatics Corporation invented the future*. Liveright Publishing Corporation.

Lev, O. and Rosenschein, J. S. (2012). Convergence of iterative voting. *Proceedings of the 11th International Conference on Autonomous Agents and Multiagent Systems*.

Lijphart, A. (1994). *Electoral systems and party systems: A study of twenty-seven democracies 1945-1990*. Oxford University Press.

Linzer, D. and Honaker, J. (2003). A theory of the evolution of voting: Turnout dynamics and the importance of groups. *Presented at the American Political Science Association Annual Meeting*.

Loewen, P. (2017). Democratic stability, representation, and accountability: A case for single-member plurality elections in canada. In Potter, A., Weinstock, D., and Loewen, P., editors, *Should we change how we vote? Evaluating Canada's electoral system*, chapter 3, pages 14–22. McGill-Queen's University Press.

Lublin, S. and Bowler, S. (2018). Electoral systems and ethnic minority representation. In Herron, E. S., Pekkanen, R. J., and Shugart, M. S., editors, *The Oxford Handbook of Electoral Systems*, chapter 8, pages 159–174. Oxford University Press.

McGhee, E. (2020). Partisan gerrymandering and political science. *Annual Review of Political Science*, 23:171–185.

McKelvey, R. (1976). Intransitivities in multidimensional voting models and some implications for agenda control. *Journal of Economic Theory*, 12(3):472–482.

McLean, I. and Johnston, R. (2009). When is proportional not proportional? great britain's 2009 elections to the european parliament. *Journal of Representative Democracy*, 45(4):349–355.

Mebane, Jr., W. R., Baltz, S., and Vasselai, F. (2019). Using agent-based models to simulate strategic behavior in elections. *Presented at the 2019 meeting of the Society for Political Methodology*.

Meir, R., Polukarov, M., Rosenschein, J. S., and Jennings, N. R. (2010). Convergence to equilibria in plurality voting. *Presented at the Twenty-Fourth AAAI Conference on Artificial Intelligence*.

Montoro-Pons, J. D. and Puchades-Navarro, M. (2001). An evolutionary model of voting. *European Journal of Economic and Social Systems*, 15(3):171–184.

Mun, N. T. (2020). Electoral system at a crossroads: Recalculation of the 2015 election results under the proportional representation system. *The Salween Institute for Public Policy*.

Myerson, R. B. (1998). Population uncertainty and poisson games. *International Journal of Game Theory*, 27:385–392.

Myerson, R. B. (2000). Large poisson games. *Journal of Economic Theory*, 94(1):7–45.

Nagel, J. H. (1994a). How many parties will new zealand have under mmp? *Political Science*, 46:139–160.

Nagel, J. H. (1994b). What political scientists can learn from the 1993 electoral reform in new zealand. *PS: Political Science and Politics*, 27:525–529.

Navia, P. D. (2003). *You select the rules of the game and lose? Advantages and constraints when choosing electoral rules: The case of Chile*. PhD thesis, New York University.

Ndegwa, S. N. (1997). The relevance of the electoral system: A simulation of the 1992 kenyan election. *African Journal of Political Science*, 2:12–40.

Norris, P. (1997). Choosing electoral systems: Proportional, majoritarian and mixed systems. *International Political Science Review*, 18(3).

Page, S. E. (2018). *The model thinker*. Basic Books.

Palfrey, T. R. and Rosenthal, H. (1985). Voter participation and strategic uncertainty. *American Political Science Review*, 79(1):62–78.

Planck, M. (1901). Ueber das gesetz der energieverteilung im normalspectrum. *Annalen der Physik*, 309:553–563.

Polga-Hecimovich, J. and Siavelis, P. M. (2015). Here's the bias! a (re-)reassessment of the chilean electoral system. *Electoral Studies*, 40:268–279.

Pollock, J. K. (1930). Why europe votes. *American Political Science Review*, 25:195.

Quinn, K. M. and Martin, A. D. (2002). An integrated computational model of multiparty electoral competition. *Statistical Science*, 17(4):405–419.

Riambau, G., Stillman, S., and Boe-Gibson, G. (2021). What determines preferences for an electoral system? evidence from a binding referendum. *Public Choice*, 186:179–208.

Riker, W. H. and Ordeshook, P. (1968). A theory of the calculus of voting. *American Political Science Review*, 62(1):25–42.

Sadiraj, V., Tuinstra, J., and van Winden, F. (2004). A computational electoral competition model with social clustering and endogenous interest groups as information brokers. *Working paper*.

Sadiraj, V., Tuinstra, J., and van Winden, F. (2006). Voting cycles in a computational electoral competition model with endogenous interest groups. *Working paper*.

Sanz, C. (2017). The effect of electoral systems on voter turnout: Evidence from a natural experiment. *Political Science Research and Methods*, 5:689–710.

Schofield, N. (1978). Instability of simple dynamic games. *The Review of Economic Studies*, 45(3):575–594.

Shugart, M. S. and Taagepera, R. (2017). *Votes from seats: Logical models of electoral systems*. Cambridge University Press.

Sieg, G. and Schulz, C. (1995). Evolutionary dynamics in the voting game. *Public Choice*, 85(1,2):157–172.

Siegel, D. A. (2018). Analyzing computational models. *American Journal of Political Science*, 62(3):745–759.

Simon, H. (1983). *Reason in human affairs*. Stanford University Press, 1 edition.

Stephenson, L. (2016). How does changing the voting system impact voter preferences? *Policy Options*.

Stephenson, L. B., Harell, A., Rubenson, D., and Loewen, P. J. (2020). 2019 canadian election study online survey.

Stiles, E. A., Swearingen, C. D., Seiter, L., and Foreman, B. (2020). Catch me if you can: Using a threshold model to simulate support for presidential candidates in the invisible primary. *Journal of Artificial Societies and Social Simulation*, 23.

Taagepera, R. (2007). *Predicting party sizes: The logic of simple electoral systems.*

Tal, M., Meir, R., and Gal, Y. (2015). A study of human behavior in online voting. *Proceedings of the 14th International Conference on Autonomous Agents and Multiagent Systems.*

Tolbert, C. J. and Kuznetsova, D. (2021). Editor's introduction: The promise and peril of ranked choice voting. *Politics and Governance*, 9:265–270.

Vasselai, F. (2021a). Pivotal probabilities in iterative voting under plurality. *Working paper.*

Vasselai, F. (2021b). Using computational game-theory and gis data to simulate election reforms. *Working paper.*

von Neumann, J. and Morgenstern, O. (1953). *Theory of games and economic behavior.* Princeton University Press, 3 edition.

Vowles, J., Aimer, P., Catt, H., Lamare, J., and Miller, R. (1995). *Towards consensus? The 1993 election in New Zealand and the transition to proportional representation.* Auckland University Press.

Vowles, J., Banducci, S., Karp, J., Aimer, P., Catt, H., and Miller, R. (1996). 1996 new zealand election study.

Whicker, M. L. and Sigelman, L. (1991). *Computer simulation applications: An introduction.* Sage Publishing.

Wuffle, A. (1984). Should you brush your teeth on november 6, 1984: A rational choice perspective. *PS*, 17(3):557–581.

Zucco, Jr., C. (2007). Where's the bias? a reassessment of the chilean electoral system. *Electoral Studies*, 26:303–314.