

NBER WORKING PAPER SERIES

KEEP YOUR ENEMIES CLOSER:  
STRATEGIC PLATFORM ADJUSTMENTS DURING U.S. AND FRENCH ELECTIONS

Rafael Di Tella  
Randy Kotti  
Caroline Le Pennec  
Vincent Pons

Working Paper 31503  
<http://www.nber.org/papers/w31503>

NATIONAL BUREAU OF ECONOMIC RESEARCH  
1050 Massachusetts Avenue  
Cambridge, MA 02138  
July 2023

For suggestions that have improved this article, we are grateful to Daniel Auer, Gustavo Bobonis, Laurent Bouton, Pierre Boyer, Micael Castanheira, Christian Cox, Ruth Dassonneville, Gemma Dipoppa, Allan Drazen, Jean-Francois Godbout, Joanne Haddad, Nicolas de Laubier Longuet-Marx, Ro'ee Levy, Pierre-Guillaume Meon, Juan Morales, Carlo Prato as well as seminar and conference participants at CREST, the Universite de Montreal, HEC Montreal, the London School of Economics, the University of Toronto, Wilfrid Laurier University, Sao Paulo School of Economics FGV, the 2022 American Political Science Association Meeting, the 2022 Montreal Applied Microeconomics Workshop, and the 2023 Barcelona Graduate School Summer Forum. We thank Emiliano Grossman for sharing data from the French Agenda Project. The views expressed herein are those of the authors and do not necessarily reflect the views of the National Bureau of Economic Research.

NBER working papers are circulated for discussion and comment purposes. They have not been peer-reviewed or been subject to the review by the NBER Board of Directors that accompanies official NBER publications.

© 2023 by Rafael Di Tella, Randy Kotti, Caroline Le Pennec, and Vincent Pons. All rights reserved. Short sections of text, not to exceed two paragraphs, may be quoted without explicit permission provided that full credit, including © notice, is given to the source.

Keep your Enemies Closer: Strategic Platform Adjustments during U.S. and French Elections  
Rafael Di Tella, Randy Kotti, Caroline Le Pennec, and Vincent Pons  
NBER Working Paper No. 31503  
July 2023  
JEL No. D72,P0

### **ABSTRACT**

A key tenet of representative democracy is that politicians' discourse and policies should follow voters' preferences. In the median voter theorem, this outcome emerges as candidates strategically adjust their platform to get closer to their opponent. Despite its importance in political economy, we lack direct tests of this mechanism. In this paper, we show that candidates converge to each other both in ideology and rhetorical complexity. We build a novel dataset including the content of 9,000 primary and general election websites of candidates for the U.S. House of Representatives, 2002-2016, as well as 57,000 campaign manifestos issued by candidates running in the first and second round of French parliamentary and local elections, 1958-2022. We first show that candidates tend to converge to the center of the ideology and complexity scales and to diversify the set of topics they cover, between the first and second round, reflecting the broadening of their electorate. Second, we exploit cases in which the identity of candidates qualified for the second round is quasi-random, by focusing on elections in which they narrowly win their primary (in the U.S.) or narrowly qualify for the runoff (in France). Using a regression discontinuity design, we find that second-round candidates converge to the platform of their actual opponent, as compared to the platform of the runner-up who did not qualify for the last round. We conclude that politicians behave strategically and that the convergence mechanism underlying the median voter theorem is powerful.

Rafael Di Tella  
Harvard Business School  
Soldiers Field Rd  
Boston, MA 02163  
and NBER  
rditella@hbs.edu

Caroline Le Pennec  
HEC Montreal  
3000 Chemin De La Cote-Sainte-Catherine  
Montreal, Quebec H3T 2A7  
Canada  
caroline.le-pennec@hec.ca

Randy Kotti  
CREST Ecole Polytechnique  
Boulevard des Marechaux  
Palaiseau 91120  
France  
kotti.randy@gmail.com

Vincent Pons  
Harvard Business School  
Morgan Hall 289  
Soldiers Field  
Boston, MA 02163  
and CEPR  
and also NBER  
vpons@hbs.edu

An online appendix is available at <http://www.nber.org/data-appendix/w31503>

# 1 Introduction

A key tenet of representative democracy is that politicians' campaign platforms and the policies they implement should follow voters' preferences. A cornerstone of modern political economy, the median voter theorem predicts that electoral incentives will generate such alignment even if candidates are only interested in their own success ([Hotelling, 1929](#); [Downs, 1957](#)). In a two-candidate race, assuming that candidates maximize votes generates two predictions: conditional on one of the candidates' proposal, the other would like to choose one that is as close as possible to the first; and both candidates will propose the policy position preferred by the median voter in equilibrium.<sup>1</sup> In [Palfrey \(1984\)](#)'s apt summary, because candidates "follow" each other, they will end up close to the "center."

Contrasting with this view of electoral politics is parties and candidates' frequent claim that they stand for ideas. As Margaret Thatcher famously declared during her 1979 campaign: "I am not a consensus politician. I am a conviction politician." Candidates' reluctance to defend propositions they do not favor may limit convergence or prevent it altogether ([Wittman, 1973, 1977](#)). In an ideal typical citizen-candidate world, voters can only elect policies because candidates are unable to propose and implement any other than their preferred one ([Osborne and Slivinski, 1996](#); [Besley and Coate, 1997](#)). Additional forces may induce even strategic candidates to keep their distance from their rivals and prevent full convergence to the median ([Grofman, 2004](#)), including party discipline, the threat of third candidate entry, and the possibility that voters penalize flip-flopping from one election to the next or that they abstain when the choices available to them are too similar ([Adams and Merrill, 2003](#)).

A large empirical literature has investigated the extent to which politicians' positions and policies correspond to voter preferences, generally reporting positive but small correlations (e.g., [Ansolabehere, Snyder and Stewart, 2001](#)). However, these correlations do not necessarily reflect strategic behavior on the part of politicians: they may instead be driven by the influence of media coverage and other common factors on the ideology of both politicians and voters, or by the self-selection of politicians into constituencies aligned with their views. By contrast, we provide direct evidence on candidates' strategies by studying how they change their discourse during the electoral campaign. We show that candidates tend to adjust their platform to the platform of their opponent, which leads them to converge to the center. Our results provide the first test of the mechanism underlying the median

---

<sup>1</sup>The original version of the median voter theorem considers a one-dimensional policy space, in which voters choose the platform closest to their preferred policy position. Follow-up models have identified the conditions under which the prediction of convergence to the median holds when policy preferences are multi-dimensional ([Davis and Hinich, 1968](#); [Davis, Hinich and Ordeshook, 1970](#); [Calvert, 1985](#)). In the separate class of probabilistic voting models, where voter preferences are uncertain, office-motivated candidates competing on a multi-dimensional policy platform target the mean voter rather than the median ([Lindbeck and Weibull, 1987](#)). Despite important differences, these models share the prediction of strategic candidate convergence to the center.

voter theorem and they shed new light on the nature of electoral competition.

Our investigation focuses on elections with two rounds in the U.S. (the primary and the general election) and in France (the first and second rounds of the general election). Because two-round elections enable to observe the same candidate in the same year, but targeting two different sets of voters, they open a window into the heart of politics. We make five distinct contributions. First, we build a novel dataset including the content of thousands of French and U.S. candidates' platforms. Second, we provide systematic evidence on changes in electoral platforms during the campaign, and show that candidates tend to converge to the center between the first and second rounds. Third, we study the key force responsible for this convergence. Exploiting quasi-experimental variation in the identity of candidates qualified for the second round, we show that candidates strategically move their platform toward the platform of opponents who qualified marginally. Fourth, we ask which candidates are most likely to adjust their platform to the platform of their competitor and find that candidates move more when they have stronger incentives to do so – for instance, when their opponent is closer to the center – and when they are more successful and experienced. Fifth, while the literature has focused on candidates' ideology and, occasionally, the topics they discuss, we consider a novel dimension of political discourse: its level of complexity. We compare the amount of convergence in ideology and complexity and show that these dimensions are complements rather than substitutes.

Data on political language have long been limited to national party manifestos collected by the Manifesto Project ([Merz, Regel and Lewandowski, 2016](#)), and to presidential candidates' speeches ([Woolley and Peters, 2017](#)). [Gentzkow, Shapiro and Taddy \(2019\)](#) went one step further and collected speeches made by U.S. House representatives and senators. However, these data only cover election winners in the period following the election. Instead, our analysis requires measuring the campaign platforms of all candidates, including candidates who lost the general election or did not even qualify for it and thus fall under the radar of most data collection efforts. In addition, we need to observe the same candidate both in the general election and before, at a time where official records of candidates' activity are scarce. In the U.S., we rely on candidates' campaign websites and consider all House races for the period 2002-2016. While we were able to obtain the URLs of general election websites from the Library of Congress, this resource does not cover the websites of candidates that lost their primary. To overcome this challenge, we hired a team of research assistants to systematically search for primary election websites in the Wayback Machine, an archive of the Internet including over 800 billion web pages. Using this method, we collected multiple snapshots of a total of 9,000 candidate websites, both before and, for candidates running in the general election, after the primary. In France, we use individual candidate manifestos: two-page documents written by all candidates before the first round and, if present in the runoff, before the second round, and sent by the state to all voters. We observe manifestos for 57,000 candidates in more than 8,000 local and parliamentary elections from 1958 to 2022.

The inclusion of French and U.S. elections, made possible by the availability of data on candidate-level platforms in both countries, enables us to study politicians' strategies both in a two-party and a multi-party setting. On the one hand, the French multi-party setting is characterized by frequent changes in the party system, weaker party affiliations, and fuzzier party lines, which may give candidates more flexibility and stronger incentives to adjust their platform than in the U.S. two-party setting. On the other hand, the threat of third candidate entry is more serious in the French two-round elections than in the U.S., where the one-round plurality rule creates a barrier to entry. French candidates may be reluctant to move away from their initial platform if this means losing some voters to a new competitor. While the inclusion of these two settings enriches the analysis, we interpret differences in results with caution, since they could also reflect differences in the type of data (online websites vs. printed manifestos) and in the amount of time separating the two rounds (one week in France vs. multiple months in the U.S.), among other factors.

We use text analysis to locate candidate platforms on the left-right axis before the first round (or primary election) and then again before the second round (or general election). Each word is given a score depending on how often it is used by Democrats vs. Republicans (or left- vs. right-wing candidates), and each text is then given a score depending on the words it contains. As a complementary measure of ideology, we also identify the topics covered by each candidate. We exploit data from the Manifesto Project indicating which topics are covered in each paragraph of U.S. and French party platforms in order to train a text classifier and determine which topics candidates talk about.

Beyond the ideology and content of candidates' platforms, we analyze their level of complexity, a dimension which has previously been underexplored but which is a key parameter of candidates' strategies. Indeed, politicians put a lot of effort to understand their audience and to adjust the way they communicate accordingly. Our measure of complexity summarizes three indicators: syntactic complexity, semantic complexity, and conceptual complexity. Candidates' complexity may reflect their own level of education and sophistication and it may in turn seduce some voters while antagonizing others. Speaking (or writing) in a simple way increases candidates' chances to be understood even by the least educated and sophisticated citizens, and it may signal their willingness to defend the interests of "the people." On the other hand, a more complex discourse may enable candidates to convey their views with more nuance and increase their appeal among more educated voters. We complete our effort to decipher campaign language by computing the dissimilarity between the platforms of rival candidates in an unsupervised way. This last metric accounts for dimensions of language which may not be already captured by our metrics of ideology, topics, and complexity.

Our first set of analyses investigate the extent to which candidate platforms converge toward the center between the first and second rounds. In the U.S., the same candidate competes for the votes of their party base against rivals from their party, in the primary election, before disputing centrist voters to a candidate of the opposite party, in the general election. To the extent that

Downsian forces are at play, one would expect the changes in the set of competitors and in the location of the median voter to push the candidate’s platform towards the center between the two rounds. We plot the distributions of candidate ideal points on the left-right ideological axis in both rounds and verify this prediction: on average, Democrat and Republican candidates’ platforms shift inward by 0.33 and 0.42 standard deviation after the primary. While political parties are organized around the ideological divide, we find that changes in candidates’ level of complexity follow the same pattern. Candidates with below- and above-median levels of complexity in the primary election both converge to a more intermediate level between the primary and the general election, a result which is not explained by reversion to the mean. Furthermore, changes in complexity and ideology are complements rather than substitutes: candidates who adjust more on one dimension also tend to adjust more on the second. We also find that candidates tend to diversify the topics they cover in the general election, by decreasing the prevalence of topics that were over-represented in their primary election website and by giving more space to topics that they barely mentioned initially.

In France, the same pool of voters is called to participate in the first and second rounds. Nonetheless, like in the U.S., candidates generally target different sets of voters in the two rounds. For instance, with an average six candidates in the first round, a left-wing candidate is likely to have to compete for left-wing voters against ideologically close opponents. Once that candidate reaches the second round, electoral competition should push them toward the center if other left-wing candidates are gone and their only remaining opponent is on the center or on the right. [Le Pennec \(2020\)](#) documents inward shifts of the distributions of the platforms of left-wing, right-wing, and far-right candidates. We reproduce this result in our larger dataset, and extend it to complexity. Like in the U.S., candidates with simple first-round platforms shift to a more complex discourse in the second round, those starting with a more complex discourse follow the opposite trajectory, and convergence in ideology and complexity go hand-in-hand. Finally, French candidates also broaden the set of policy topics they cover in their manifestos between election rounds.

In the second part of our analysis, we explore whether the overall convergence taking place between the first and second rounds can be explained by candidates’ effort to adjust their platform to their second-round opponent, thereby testing the key mechanism underlying the median voter theorem. Convergence to a more centrist platform on the ideological and complexity dimensions could plausibly result from other factors, such as learning, a weakening of party discipline (allowing moderate candidates to move back to their ideal point), or changes in the types of individuals willing to contribute time or money to the campaign. To isolate the changes in candidate platforms explained by their effort to adjust to their opponent, we exploit quasi-experimental variation in the identity of that opponent. In the U.S., this variation is provided by close primaries. We ask whether Republican candidates move their platform closer to the platform of the winner of a close Democratic primary than to the platform of their runner-up. Formally, we use a regression discontinuity design with two observations for the same Republican candidate. Our outcome measures the extent to which the

change in the Republican candidate’s platform between the primary and the general election closes the gap with the primary platform of the Democratic winner, on the right of the threshold, and with the primary platform of the Democratic runner-up, on the left. The running variable is the winning margin of the Democratic primary, so that observations near the threshold correspond to close primaries. The size of the jump at the threshold, if any, measures the magnitude of the adjustment by the Republican candidate to their Democratic opponent, relative to the Democratic runner-up. Importantly, while candidates may in principle choose to emulate the winner of the opposite party’s primary because they infer that this platform was the most appealing to voters, our RDD shuts down this channel. Since we focus on close primary elections, similar numbers of voters are located close to the positions of the two Democratic candidates, so closing the gap with the winner demonstrates candidates’ strategic intent to get closer to the opponent they are facing.

We use a similar design in the French elections. In that case, first-round races in which the candidates ranked second and third have almost exactly the same vote share provide quasi-experimental variation in the identity of the opponent that the candidate ranked first in the first round will face in the second round. We test whether the first candidate moves their platform closer to the platform of the second candidate, between the first and second rounds, than to the platform of the third candidate.

Focusing first on changes in ideology, our design is best adapted to measure strategic adjustment in the French setting. To fix ideas, consider a race in which the candidate ranked first in the first round is from a centrist party, one of the two following candidates is on the left, and the other is on the right. Focusing on races in which the second and third candidates obtained nearly exactly the same vote share in the first round, our RDD tests whether the centrist candidate moves to the left when the left-wing candidate finished a marginal second in the first round, and to the right when the right-wing candidate finished ahead. We find that this is indeed the case: overall, the change in similarity between the first-ranked candidate’s ideology and that of their qualified opponent, from first to second round, is 0.46 standard-deviation larger than the change in similarity with the third-ranked candidate who failed to qualify. In the U.S., the Republican candidate will generally remain to the right of both the winner and runner-up of the Democratic primary even if they converge to their opponent. Therefore, the change in the distance between their platform and the platform of the primary winner will in general be identical to the change in the distance with the platform of the runner-up. Unsurprisingly, then, we do not find any significant differential convergence to the primary winner. While we find some evidence of convergence to the topics covered by the qualified opponent in France, we do not find any significant topic convergence in the U.S. either.

By contrast, the level of complexity of U.S. House candidates is frequently located between the complexity of the top-two candidates in the primary of the opposite party. Focusing on these middle point races, we find that candidates adjust their level of complexity to their opponent in the general

election. By contrast, in France, we do not find any evidence of strategic adjustment on that dimension. Overall, an index summarizing convergence in ideology, complexity, topics, and unsupervised text distance shows an effect of 0.41 standard deviation in U.S. elections (significant at the 1 percent level), and 0.30 standard deviation in French elections (significant at the 5 percent level).

Our third set of results investigates which candidates are most likely to adjust their platform to the platform of the opponent they face in the final round of voting, and under which conditions. We show that candidates adjust their platform more when their incentives to do so are stronger, namely when their opponent is a greater threat. Using an alternative RDD, we find that facing an opponent who is likely to rally voters in the second round (because she is more moderate, because she is the incumbent, or because her party won more votes in the past) leads a candidate to converge more. The level of strategic adjustment is not just determined by incentives but also by politicians' type: experienced politicians, who were successful in the past (e.g., they won the previous election or received more votes in the first round of the election), adjust to their opponent the most even though they least need it.

Our paper builds on a large empirical literature studying the strategies of candidates and parties. Economists and political scientists have investigated the determinants of candidates' decision to run (instead of staying out of the race) as well as the factors facilitating dropout agreements between parties (e.g., [Anagol and Fujiwara, 2016](#); [Pons and Tricaud, 2018](#); [Granzier, Pons and Tricaud, 2023](#); [Dano et al., 2023](#)); the methods candidates use to contact voters (e.g., [Gerber and Green, 2000](#); [Pons, 2018](#); [Spenkuch and Toniatti, 2018](#)); and their efforts to increase their campaign resources, from fundraising to the mobilization of volunteers (e.g., [Bouton et al., 2022](#); [Adena and Hager, 2022](#)). We focus on a different but crucial aspect of candidate strategies: which platform they choose and which language they use to convey their propositions.

A core assumption of spatial competition models such as the Downs-Hotelling model ([Hotelling, 1929](#); [Downs, 1957](#)) and the probabilistic voting model is that voters are more inclined to vote for politicians that are ideologically closer to them. Consistent with this assumption, there is evidence that moderate candidates tend to win more votes than those who only appeal to a small number of extreme voters ([Erikson and Wright, 1980](#); [Canes-Wrone, Brady and Cogan, 2002](#)). For instance, using a regression discontinuity design, [Hall \(2015\)](#) finds that a party gets fewer votes in the general election when an extremist won the primary. Such results imply that candidates have an incentive to adjust their discourse to voters' preferences but they do not tell us whether and to what extent candidates actually respond to this incentive.

The vast majority of the papers addressing the latter question have studied the strength of the relationship between politicians' positions, as measured through surveys or based on their roll-call votes, and voters' preferences, as proxied for instance by the presidential vote shares in a certain district. Unsurprisingly, the prediction of full convergence has been repeatedly rejected. For instance, [Poole](#)



and Rosenthal (1985) show that Democratic and Republican senators from the same state generally hold different positions even though their constituents are, by construction, identical. Other studies have attempted to estimate the weight parliamentarians assign to their constituents' preferences in their congressional votes by correlating these two variables while controlling for other factors such as party ideology, and they have generally uncovered weak positive correlations (Levitt, 1996; Ansolabehere, Snyder and Stewart, 2001). However, the magnitude of these correlations could also reflect other mechanisms, such as selection: potential candidates are drawn from the population of voters, which may reduce the gap between their preferences, and those with unusual positions may reason that they are unlikely to win and refrain from running.

Data permitting, a more direct approach to study candidates' propensity to adjust their discourse to voters' preferences is to track changes in their position over time. Using national manifestos, existing studies have provided descriptive evidence that parties adjust their policy platform to changes in public opinion (see Adams (2012) for a review) and to shifts in other parties' success. For instance, mainstream parties tend to adapt their policy agenda in response to the growing success of the radical right (Abou-Chadi, 2016; Abou-Chadi and Krause, 2020). However, studies focusing on national parties' positions are constrained by their limited number of observations. By contrast, Lee, Moretti and Butler (2004) exploit the roll call votes of individual politicians. They find that a close victory by the Democratic party increases its electoral strength (due to the incumbency advantage) but fails to shift the roll call votes of politicians elected in the district to the left, suggesting that changes in voter preferences do not shape politicians' positions. Another possible interpretation is that close victories change vote shares while leaving voters' policy preferences unaffected. Strategic politicians making that interpretation would have no reason to change their position.

To understand how politicians choose and adjust their platform, there is no reason to stop at changes across elections: a lot can be learnt from changes by the same candidate across different rounds of the same election. In the U.S., for instance, the popular press is replete with anecdotes about primary candidates winning the nomination with an extreme platform that caters to their base, only to converge to the center as they advance to the general election. However, there is surprisingly little systematic evidence about the prevalence of this behavior and the magnitude of the corresponding shifts. Acree et al. (2020) focus on three U.S. presidential candidates in 2008 and 2012 and show that these candidates moderated their language between the primary and the general election. Burden (2001) shows that the roll-call votes of representatives in the 102nd legislature were more moderate after the 1992 primary election than beforehand, but this evidence is restricted to incumbents and may be unrepresentative of other elections. By contrast, we study the behavior of both incumbents and candidates of the opposing party, and we document systematic convergence in a sample including eight distinct U.S. election cycles from 2002 to 2016. We also establish systematic convergence to the ideological center between the first and second rounds of French elections by replicating and extending Le Pennec (2020)'s results. Our large sample considerably increases

the external validity of our results. Furthermore, the change in scale relative to previous work ensures that we have sufficiently many observations close to the threshold, in our RDD, enabling us to provide causal evidence on the extent to which changes in candidate platforms between rounds are driven by exogenous variation in the identity of their opponent in the second round.

Unlike previous studies, we also expand our analysis of convergence between rounds in France and the U.S. to a second dimension, complexity. The complexity of political discourse has received little attention heretofore. Exceptions include [Foarta and Morelli \(2021\)](#), who discuss the costs and benefits of producing complex legislative and regulatory texts (not politicians' discourse) and model the determinants of complexity. Closer to our paper, [Spirling \(2016\)](#) finds that the extension of the franchise in 19th century Britain decreased the complexity of speeches made by members of Parliament and interprets this result as an attempt to appeal to poorer and less educated voters. Building on this study spanning more than eight decades, we ask whether candidates will go all the way to adjust the complexity of their discourse within the same election.

Finally, our paper speaks to a broad set of studies which show that dramatic policy changes were brought about by the inclusion of new voters due to the expansion of the franchise to poorer households ([Meltzer and Richard, 1983](#); [Husted and Kenny, 1997](#)) and women ([Miller, 2008](#)), the enfranchisement of minorities ([Cascio and Washington, 2014](#)), the rollout of technologies facilitating the participation of uneducated voters ([Fujiwara, 2015](#)), or the adoption of compulsory voting ([Fowler, 2013](#)). The effects on policies and on social and economic outcomes documented by these papers are generally consistent with the median voter theorem prediction but they could also be explained by alternative mechanisms, such as the entry of new candidates, a change in the party in power, or increased salience of issues mattering to the new voters. While disentangling these different mechanisms is difficult, our results suggest that politicians' efforts to compete for newly enfranchised voters and court them by moving towards their preferences may have played an important role.

The remainder of the paper proceeds as follows. Sections [2](#) and [3](#) describe our data and outcomes. Section [4](#) provides evidence that candidates converge to the center between election rounds. Section [5](#) shows that this convergence is at least partly driven by candidates' adjustments to their opponent, and Section [6](#) concludes.

## 2 Setting and Data

### 2.1 U.S. elections and candidate websites

**U.S. elections** Elections for the U.S. House of Representatives occur every two years and elect 435 representatives. A plurality of votes is enough to win the election, except in Georgia and Louisiana, where runoff general elections are held when no candidate reaches a majority of votes. We rely on the general election results made available by [MIT Election Data and Science Lab \(2017a\)](#). Although candidates from third parties run in 54% of the general elections held between 2002 and 2016, 99% of these elections were won by either a Republican or a Democrat candidate and the combined vote share of these two parties was 95% on average. In other words, these elections are largely bipartisan.

Primary elections are held in many congressional districts among candidates from the Democrat and Republican parties prior to the general election. Between 2002 and 2016, 50% of general elections were preceded by at least one competitive primary, with two candidates or more competing against each other. When both parties hold a primary in a district, these two primaries are held on the same day. In ten states, a majority of votes is required to win the primary, and a runoff takes place if no candidate reaches 50% of votes ([Underhill, 2017](#)).<sup>2</sup> In all other states, a plurality of votes is sufficient to win and qualify for the general election. We collected the primary election dates from the Federal Election Commission’s website<sup>3</sup> and use the primary election results gathered by [Pettigrew, Owen and Wanless \(2014\)](#) for the 1956 to 2010 elections, and by [Miller and Camberg \(2020\)](#) for the 2012 to 2018 elections.<sup>4</sup>

**Candidate websites** Campaign websites are an ideal source of data to study political messaging. Their content is directly elaborated by the candidate and their team, they target a broad audience of voters, and they paint a comprehensive picture of a candidate’s program ([Druckman, Kifer and Parkin, 2010](#)). According to campaign insiders surveyed by [Druckman, Kifer and Parkin \(2018\)](#) between 2008 and 2016, websites remain an important element of candidates’ campaigns, despite the increased use of other forms of communication such as social media. Websites are used to present candidates’ background and political positions, on top of signing up volunteers and raising funds.

We collect the content of candidate websites through the Wayback Machine, a non-profit initiative that aims to keep track of the Internet’s most important content.<sup>5</sup> The Wayback Machine enables

<sup>2</sup>In North Carolina, a runoff is not required unless the second-ranked candidate calls for one, which they have the possibility to do if the first-ranked candidate gets less than 30% of votes.

<sup>3</sup><https://www.fec.gov/introduction-campaign-finance/election-and-voting-information/2002-us-congressional-primary-election-dates-and-candidate-filing-deadlines-ballot-access/>.

<sup>4</sup>We note that 62 primary runoffs were missing from [Pettigrew, Owen and Wanless \(2014\)](#). We added results from these runoffs manually.

<sup>5</sup>The Wayback Machine can be accessed at <https://archive.org/web/>.

us to view campaign websites as they appeared online at multiple points in time throughout the campaign and to browse through them. As an example, Appendix Figure A.1 shows the main page of a website capture.

In order to find the captures of a website on the Wayback Machine, we first need to identify the website's URL. Website URLs of candidates present in a *general* election are curated and archived by the Library of Congress (LoC).<sup>6</sup> By the time our data analysis was completed, in spring 2023, this database ran between 2002 and 2016. More recent years had been curated but were still under embargo. We linked these URLs to the corresponding candidates in our results' database using fuzzy string matching. Since the LoC only archives URLs of general election candidates, we adopt an alternative approach to find the URLs of candidates present in a *primary* election. We use the Wayback Machine's search function, which matches keywords with URLs and titles of over 820 billion archived web pages. To collect as many primary election websites as possible, we first used brute search before turning to manual search and a series of systematic manual checks. Importantly, primary election winners' websites are referenced by the LoC database – since they qualified for the general election – but primary losers' websites are not. To avoid gathering artificially more websites from primary election winners than losers, we separated the search procedure for primary and general election candidates. Therefore, websites of primary election winners are gathered twice: once using URLs from the LoC database, and once using URLs found on the Wayback Machine's search interface.

After retrieving the URLs pointing to House of Representatives' campaign websites, we scraped all the textual content displayed on the main page and all sub-pages accessible with one click from the main page, for each time capture of each website.<sup>7</sup>

In total, we were able to find and verify 35,427 website captures taken between the beginning of the election year and the day of the primary election for 3,185 candidates across 2,022 competitive primary elections (an average of 11 captures per candidate, and 8 sub-pages per capture). In the remaining of the paper, we refer to these websites as *primary election websites*.<sup>8</sup>

We were also able to collect 156,943 website captures (an average of 27 time captures per candidate and 12 sub-pages per capture) for 5,792 candidates present in a general election, across 3,036 general elections. Specifically, 44,871 captures were taken between the beginning of the election

---

<sup>6</sup><https://www.loc.gov/collections/united-states-elections-web-archive/>.

<sup>7</sup>For more details on both the URL search and the scraping process, see Appendix A.

<sup>8</sup>In districts where a competitive primary is followed by a runoff (3.2% of the districts in our sample), we followed a specific set of rules to determine the date at which we stop collecting a primary election website's captures. If both the Democrat and Republican parties hold a runoff, or if one of the two parties holds a runoff while the other party holds no primary at all, we collect all the captures taken before the date of the runoff. If one party holds a primary with a runoff while the other has a primary but no runoff, the distinction between the pre-primary and post-primary periods is more ambiguous. In these cases (accounting for only 1.1% of all districts), we stopped collecting primary elections websites' captures after the date of the primary's first round.

year and the date of the primary election, and 112,072 captures were taken between the primary election and the general election.

Appendix Table A.1 provides the number of primary and general election races for which we have collected at least one website, the total number of candidates competing in these races, and the number of candidates for which a website is available.

## 2.2 French elections and candidate manifestos

**Electoral rule** We combine our U.S. dataset with data from French local and parliamentary elections held between 1958 and 2022. Parliamentary elections are held every five years to elect the 577 representatives seating in the *Assemblée Nationale*, the lower house of the French Parliament. Local elections are held every six years to elect the representatives seating in departmental councils, which hold the legislative power in each of the 101 *départements* – the territorial and administrative units in charge of education, transportation, and social assistance, among other prerogatives.

Parliamentary and local elections follow the same electoral rule: each constituency (*circonscription* for the former election type, *canton* for the latter) elects one representative under a two-round plurality rule. A candidate needs to obtain the absolute majority of the votes cast in their constituency to win in the first round, and these votes need to account for at least 25% of all registered voters. If no candidate is elected in the first round, the two candidates who received the most votes and any other candidate who obtained the votes of at least 12.5% of the registered voters qualify for the second round. The runoff takes place among all qualified candidates who choose to stay in the race, and the candidate who receives a plurality of votes gets elected.<sup>9</sup>

While this electoral rule has remained relatively stable over the period that we study, a few historical changes must be noted. First, the first round vote share required to qualify for the second round has changed. In parliamentary elections, it went from 5% of the expressed votes in 1958 to 10% of the registered voters in 1966 and to the current threshold of 12.5% of the registered voters in 1975. In local elections, it went from 10% of the registered voters to 12.5% in 2010. Second, the 1986 parliamentary election used a proportional list system at the *département* level instead of the two-round plurality rule in single-member constituencies described above. We exclude this election year from our sample. Third, local elections were substantially transformed through a large reform prior to the 2015 elections. Until 2015, each *canton* elected one council member for a six-year mandate and, every three years, half of the constituencies voted to renew their representatives. Since 2015, local elections have been taking place every six years in the entire country, and they elect tickets composed of a male and a female candidate and campaigning on a common platform.

---

<sup>9</sup>Over our sample period, a runoff was held in 86% of parliamentary elections and 69% of local elections.

Electoral results at the candidate level for each of these elections, along with candidate characteristics such as their gender, their incumbency status, and whether they have run in the past or not, were obtained from [Granzier, Pons and Tricaud \(2023\)](#) and [Dano et al. \(2023\)](#).

**Party system** Unlike the two-party U.S. setting, the French political system is multipartisan. French politics have been historically dominated by a left-wing block organized around the Socialist party and a right-wing block organized around the conservative Gaullist party (currently called *Les Républicains*), but other parties have also been present on the ballot in many constituencies throughout our sample period. Important examples include the Communist party, the Green party, the centrist party MODEM, and the far-right party RN. Beyond these large national parties, elections often feature candidates affiliated with smaller issue-specific parties. Candidates may also run as independents, without the endorsement of any national party.<sup>10</sup>

Following [Granzier, Pons and Tricaud \(2023\)](#), we allocate candidates to six political orientations (far-left, left, center, right, far-right, and non-classified) based on political labels assigned to each candidate by the Ministry of the Interior – the official publisher of the election results. This classification applies to independent candidates as well, since candidates who are not affiliated with any of the main parties might nonetheless have a clear political orientation, indicated by labels such as "diverse left" or "diverse right."

**Candidate manifestos** During the official campaign period preceding both local and parliamentary elections, candidates are invited to issue a campaign manifesto (or *profession de foi*), in which they may advertise their policy platform as well as their personal attributes. These two-page documents are mailed to all registered voters by the state a few days before the first round.<sup>11</sup> Importantly, candidates who qualify for the runoff are invited to issue a new manifesto prior to the second round, that is also sent to voters. Additional details are provided in Appendix A and an example of a candidate manifesto is shown in Appendix Figure A.2.

Our dataset of candidate manifestos builds on a corpus assembled by [Le Pennec \(2020\)](#) for the parliamentary elections held between 1958 and 1993. [Le Pennec \(2020\)](#)'s data only cover a single year post 1993: the 2017 parliamentary elections.<sup>12</sup>

We first complete these data with the manifestos for the 1997 parliamentary elections, which were collected by [Cagé, Le Pennec and Mougin \(Forthcoming\)](#).

<sup>10</sup>Independent candidates account for 31% of all candidates running in our sample period and for 18% of the candidates who are ranked first, second, or third in the first round.

<sup>11</sup>The expenditures associated with the preparation and printing of manifestos are fully reimbursed by the state, provided that the candidate obtains at least 5% of the votes in the first round of the election.

<sup>12</sup>The 1958-1993 manifestos were digitized by the Archelec project ([Gaultier-Voituriez, 2016](#)), and the 2017 manifestos were published online (at the candidate's discretion) on the Ministry of the Interior's website.

Second, we scraped the manifestos published online for the 2021 local election and the 2022 parliamentary election.<sup>13</sup>

Third, we made a big effort to find manifestos issued for the intermediate parliamentary elections, held in 2002, 2007, and 2012, and for the local elections held between 1979 and 2015. We visited departmental archives, municipal archives, and town halls throughout *Ile de France*, France’s most populated region (which includes Paris and its surroundings) and were able to find and digitize a subset of the manifestos for these elections.<sup>14</sup> We applied optical character recognition to convert their content into machine-readable text.

In total, our dataset contains first-round manifestos issued by 46,607 candidates across 8,156 races, and second-round manifestos issued by 10,310 runoff candidates across 5,209 races. Appendix Table A.2 indicates the number of first- and second-round races for which we have at least one manifesto, for each local and parliamentary election, the total number of candidates competing in these races, and the number of candidates for which a manifesto is available.

### 3 The Multiple Dimensions of Political Discourse

We consider three dimensions of political discourse: ideology, the complexity of candidates’ language, and the policy topics they talk about. We use text analysis to construct measures associated with each of these dimensions. We also construct an agnostic measure of text similarity between any two documents.

#### 3.1 Ideological score

Our first measure captures the ideology of candidates’ discourse. We use a supervised approach to scale candidate websites and manifestos from left to right. Intuitively, we use candidates’ party affiliation to identify all words associated with each ideological side and we then estimate candidates’ ideology based on their choice of words. Within each country, we treat each election year separately in order to account for changes in the ideological leaning of words over time. Our approach builds on the Wordscores method introduced by Laver, Benoit and Garry (2003).

The ideological score of document  $j$  is defined as:

$$S_j = \sum_w p_{wj} \cdot s_w,$$

<sup>13</sup>These manifestos were available at <https://programme-candidats.interieur.gouv.fr/>.

<sup>14</sup>See Appendix A for more details on this data collection effort.



where  $s_w$  is the score of word  $w$ , as defined below, and  $p_{wj} = \frac{c_{wj}}{m_j}$  is the frequency of word  $w$  in document  $j$ , with  $c_{wj}$  the number of occurrences of word  $w$  in document  $j$  and  $m_j$  the total number of words in document  $j$ .

In France, a document is either a first-round or second-round manifesto. In the U.S., we concatenate all the different captures of a campaign website prior to the primary election date to have one primary election data point per candidate. Similarly, we concatenate all the captures of a website between the day of the primary election and the day of the general election to have one general election data point per candidate. Therefore, in both countries, we have at most two data points (and two scores) per candidate: one for each election round. In the remainder of the paper, we use the term "first round" to refer both to primary elections in the U.S. and to first election rounds in France, and the term "second round" to refer both to general elections in the U.S. and to runoffs in France.

We construct the word scores  $s_w$  in the following way. In the U.S., we use the content of websites published by Democratic candidates (labeled as "left") on one side, and Republican candidates (labeled as "right") on the other, excluding websites from independent or third-party candidates. In France, we aggregate the content of manifestos issued by candidates labeled as left-wing or right-wing, excluding manifestos from centrist and non-classified candidates.<sup>15</sup> We only use the primary election websites and first-round manifestos as reference texts.<sup>16</sup> Word scores are then defined as:

$$s_w = \frac{p_w^R}{p_w^L + p_w^R} - \frac{p_w^L}{p_w^L + p_w^R},$$

where  $p_w^I = \frac{1}{|I|} \sum_{j \in I} p_{wj}$  is the average frequency of word  $w$  among documents from ideological side  $I$  (with  $I = L, R$ ). The score of each word ranges from  $-1$ , when it is only used by left-wing candidates and never by right-wing ones, to  $1$ , when it is only used by right-wing candidates and never by left-wing ones. We exclude words that are used too infrequently in order to limit the influence of rare words that appear extreme because they are, by chance, used only by one or few candidates from the same ideological side, even though they do not carry any partisan meaning. We also exclude the most commonly used words, which do not differentiate partisan sides, to improve computational efficiency.<sup>17</sup>

Intuitively, a website or a manifesto with a negative (positive) score contains mostly words that are primarily used by left (right)-wing candidates, while a document with a partisan score close

<sup>15</sup>Left-wing candidates include candidates from parties such as the Communist Party or the Socialist Party, and candidates labeled as "Other Left." Right-wing candidates include candidates from parties such as the National Front or Rally for the Republic, and candidates labeled as "Other Right."

<sup>16</sup>If candidates start using words from the opposite ideological side in the second round, using second-round documents as reference texts would result in more neutral word scores, and candidates who use these words in the first round would receive a more neutral ideological score as well. This could lead us to underestimate candidates' changes in ideological scores between rounds as well as their adjustments to the opponent.

<sup>17</sup>Specifically, we exclude words that are used by fewer than 0.5% and more than 80% of all left-wing and right-wing candidates.



to zero either contains words that are just as common among left-wing and right-wing candidates of words that are primarily used by the left and words primarily used by the right.<sup>18</sup> Appendix Tables B.1 and B.2 show the twenty words with the highest (most right-leaning) and lowest (most left-leaning) score for the U.S. and French elections in the sample. In both countries, we find words referring to economic policy and redistribution (e.g., "wealthiest," "richest," "equality," "trades") and the environment (e.g., "renewable," "solar," "pollution") on the left. In the U.S., many left-wing words also refer to minority rights (e.g., "minority," "discrimination," "gay") while many right-wing words refer to religion (e.g., "sanctity," "pray," "bible") and abortion (e.g., "unborn," "abortions"). In France, many right-wing words refer to immigration policy (e.g., "stay," "clandestine") and crime (e.g., "military police," "brigade," "terrorism").

### 3.2 Complexity score

Our second measure captures the complexity of the language used in candidates' websites or manifestos. Little attention has been heretofore given to discourse complexity in political economy (but see [Spirling \(2016\)](#)). This gap is surprising given that campaigns are, at their core, acts of communication, and politicians are often thought as experts at understanding their audience and finding the most efficient way to communicate with it. We may thus expect them to adjust the sophistication of their discourse depending on the voters they are trying to persuade. First, language complexity may send a signal of who the politician is and which group of voters they will represent once in office (e.g., "the people" vs. the elite). Second, voters may be more receptive to a campaign message that they can understand without being too simplistic, and different voters may prefer different levels of complexity. For instance, using a simple language may appeal to less educated voters as well as immigrants whose native language is not English, while using a complex rhetoric may appeal to highly-educated voters.

One of the most widely used proxies for textual complexity is the Flesch–Kincaid readability metric. However, [Benoit, Munger and Spirling \(2019\)](#) found that more recent metrics brought by advances in software tools perform closer to human ratings when it comes to measuring different dimensions of textual sophistication in the field of politics. We combine their findings with the approach of [Tolochko and Boomgaarden \(2018\)](#) and [Hurka and Haag \(2020\)](#) to define textual complexity along three main dimensions: the complexity of sentences, the complexity of words, and the overall conceptual complexity.

First, we measure the syntactic complexity, which refers to the complexity of sentences' structure, by computing the ratio of the number of subordinating conjunctions and relative pronouns to the

---

<sup>18</sup>The final score of each document is further normalized as recommended by [Martin and Vanberg \(2007\)](#). See Appendix B.2 for more details.

total number of words (Montoro and McIntyre, 2019; Benoit, Munger and Spirling, 2019).<sup>19</sup>

Second, we measure the semantic complexity, which refers to the complexity of words themselves. We approximate words' complexity by measuring their entropy, using Google Books as reference text (Michel and Orwant, 2011).<sup>20</sup> For any word appearing with frequency  $f$  in the Google Books corpus, we define its entropy as  $-f \log f$ . As a result, the entropy is mostly a decreasing function of words' frequency: the rarer a word, the higher its entropy. However, below a certain threshold, the entropy is an increasing function of words' frequency. This ensures that extremely rare words found in candidate websites and manifestos (such as typos or OCR errors) are not considered complex words.<sup>21</sup>

Third, we measure the conceptual complexity of texts, which refers to the difficulty to understand the ideas embedded within a text, regardless of its form. Levy, Razin and Young (2022) use the ratio of unique words to total words in speeches as a proxy for the complexity of ideas: more diverse words point to a more complex idea. This metric is also known as Type-Token Ratio (TTR), but it is strongly correlated with text length (e.g., a 10,000 word text will necessarily use the same words several times and will therefore have a lower ratio of unique words as compared to a very short and simpler text). To mitigate this problem, we use the Moving Average Type Token Ratio (MATTR), which computes the average TTR through a moving window of 200 words (Covington and McFall, 2010).<sup>22</sup>

We define the complexity score of a document as the standardized average of these three components: the share of subordinating words and relative pronouns, the average word entropy, and the MATTR.<sup>23</sup> We discuss the validity of this approach in Appendix B.3 and compare the complexity of our corpus against some benchmarks. We find that the average U.S. website is equally complex as the business/financial section of the New York Times, but less complex than its book review section and more complex than its sports section (Appendix Figure B.1). The 2022 French manifestos are

<sup>19</sup>Some manifestos were processed through optical character recognition (OCR), which may limit our ability to identify punctuation. Instead of relying on sentence length, we count subordinating conjunctions such as "when," "where," "whether" in English and "que," "quoi," "où" in French. In practice, we use the R implementation of the OpenNLP package to identify subordinating conjunctions (part-of-the-speech tagging) in the U.S. Since this package does not support French language, we use lists from "Le Robert" dictionary to detect subordinating conjunctions and relative pronouns in French manifestos: <https://dictionnaire.lerobert.com/guide/conjonctions-de-subordination> and <https://dictionnaire.lerobert.com/guide/pronoms-relatifs>.

<sup>20</sup>We use the latest Google Books corpus available: 2008 in English and 2009 in French.

<sup>21</sup>Historically, semantic complexity was measured by the average number of syllables per word (Flesch, 1948) or the proportion of "difficult" words, i.e., words that are not included in a list of most common words (Dale and Chall, 1948). Benoit, Munger and Spirling (2019) use the average word rarity as given by the Google books corpus, an approach that is not immune to extremely rare words (likely to be typos or OCR errors). The measure of word entropy which we rely on, like Hahn and Sivley (2011) and Katz (2013), is analogous to Shannon's entropy in information theory (Shannon, 1948), which measures the computational resources required to process a telecommunication.

<sup>22</sup>We use the Quanteda package (Benoit et al., 2018) to measure MATTR.

<sup>23</sup>For the complexity analysis, we cannot concatenate all the captures of a U.S. candidate's website together like we do for the ideological score, because some metrics such as the Type-Token Ratio are not linear. Instead, we define the complexity score as the average complexity of all website captures.

more complex than their "easy to read and understand" counterparts, which candidates had the possibility to publish online along with their original manifesto in that specific year (Appendix Figure B.2).

### 3.3 Topic distribution

Besides their ideological tone and the complexity of their rhetoric, candidates may strategically choose which topics to focus on to persuade voters. To quantify the relative importance of different topics, we implement a supervised machine learning model trained on the manifestos issued by national parties.

For the U.S., we rely on the Manifesto Project (Lehmann et al., 2021), which gathered manifestos for the 2004-2020 period and hand coded each sentence to fit into one of 31 topics such as "Human rights," "Protectionism," "Education," or "Agriculture and farmers." For France, the party manifestos hand coded by the Manifesto Project are only available for the 2012 and 2017 elections. Therefore, we use an alternative data source: the French Agenda Project (Grossman, 2019), which gathered party manifestos for the 1981-2017 period and hand coded each sentence to fit into 27 topics such as "Economic regulation," "Health," "Education," or "Immigration." Appendix Table B.3 provides the full list of topics in each country, together with the words that are most predictive of each topic.

We feed the sentences from the national manifestos into a TF-IDF vectorizer and train a Support Vector Machine (SVM) model to predict each sentence's topic.<sup>24</sup> Once trained on the national manifestos, we used the SVM model to predict the topics appearing in our corpus. For each document, the model outputs a vector indicating the likelihood that each topic is discussed.<sup>25</sup>

To check the quality of our topic predictions, we regress each measure of topic prevalence on our ideological score (Appendix Figure B.3) and complexity score (Appendix Figure B.4). On average, controlling for year fixed effects, candidates more to the left talk more about education and equality, while candidates more to the right focus on markets and administrative efficiency. More complex candidates speak more of international affairs, while less complex candidates refer more to labour groups.

---

<sup>24</sup>This method yields a higher out-of-sample cross-validated accuracy than the alternative classifiers that we tried. We explain at greater length why we chose this approach in Appendix B.4.

<sup>25</sup>Similarly as for the ideological score, in the U.S., we concatenate all the captures of a candidate's website prior to the day of the primary election on the one hand, and all the captures of their general election website between the primary and the general election on the other hand.

### 3.4 Text similarity

Finally, we construct a measure of textual similarity between two documents based on their overall content. This measure is based on the average cosine similarity between the vector representations of texts, across multiple choices of such representations – including TF-IDF, Word2Vec, and BERT. Regardless of the technique chosen to transform texts into vectors, textual similarity relies on an unsupervised approach and is therefore more "agnostic" than measuring the distance between two texts' ideological scores, complexity scores, or topic distributions. It captures whichever parts of the text make two documents similar or different from each other, including other important dimensions that we could miss by just focusing on ideology, complexity, and topics. See Appendix B.5 for more details on this measure.

## 4 Convergence Between Rounds

We now investigate whether candidates adjust the content of their campaign communication to match their voters' preferences. We begin by computing correlations between voter characteristics and candidates' ideology and complexity scores in the second round.<sup>26</sup>

As shown in Appendix Tables C.1 and C.2, candidates tend to use more right-wing language in low-density districts (a correlation that is only significant in France) and in districts with stronger support for right-wing candidates in the previous presidential election (which is significant at the 1% level in both countries). Candidates running in more educated districts tend to use more left-wing language (significant at the 10% and 1% level in the U.S. and France, respectively) and more complex language (an estimate that is significant only in the U.S., but sizeable in France as well). These correlations suggest that candidates adapt the content of their communication to their electorate, but they could be confounded by reverse causality (e.g., if more educated voters sort where more complex politicians run) or omitted variable bias (e.g., if other factors explain both why more educated voters and more complex politicians live in the same area).

To overcome these issues, we exploit the two-round setting of the U.S. and French elections and study how the same candidate adjusts their discourse to different electorates. Focusing on within-candidate changes in discourse between rounds enables us to hold many factors constant, including the political climate and candidates' identity. The fact that the same candidate targets different electorates in the first and second rounds is obvious in U.S. elections. As candidates who win the primary election move from competing within their own party to competing against the other party, their electorate broadens and becomes more diverse. In France, the set of eligible voters remains

---

<sup>26</sup>See Appendix A.3 for more details on the voter characteristics we include in this analysis.

identical across rounds, but the composition of the actual electorate may change. For instance, extreme voters who participated in the first round may abstain from the second round if they feel too distant from any of the remaining candidates (Pons and Tricaud, 2018). Furthermore, the candidates who qualify for the runoff generally move from competing against many opponents in the first round, including other candidates from the same ideological side, to competing against a single candidate from the other side in the second round: 70% of runoff races in our sample oppose one left-wing and one right-wing candidate. Just like in the U.S., candidates have an incentive to target their base in the first round and to address a broader set of voters in the second round.

While candidates may be tempted to adjust their discourse to target a broader audience, flip-flopping may be costly, limiting the magnitude of the convergence. Indeed, since candidates anchored their platform in the first round, changing their discourse in the second round may hurt their reputation and cost them votes.

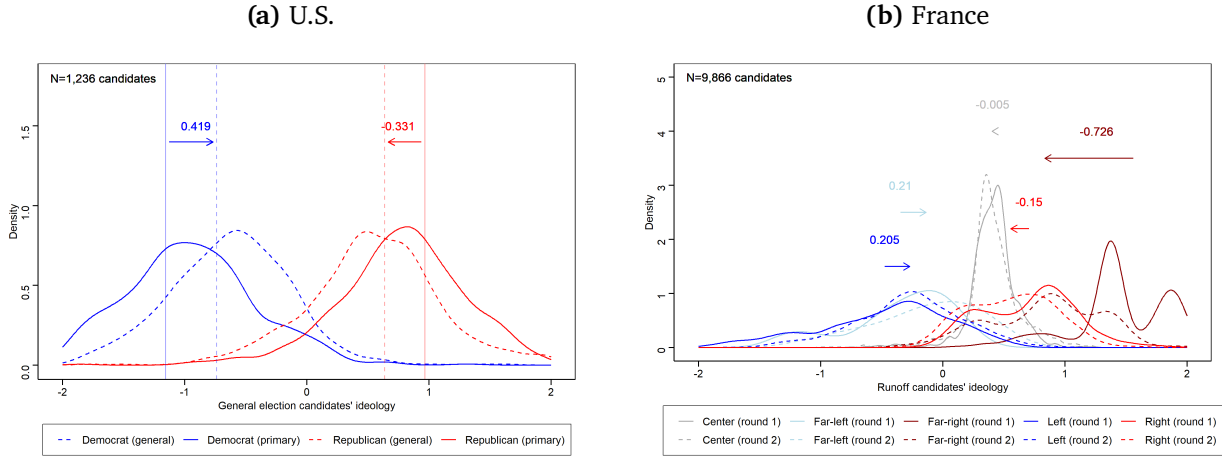
## 4.1 Ideological convergence

In Figure 1a, we plot the kernel density of ideological scores for Democratic and Republican candidates separately, pooling all elections year together. Our sample includes candidates who compete both in a competitive primary election (i.e., with two candidates or more) and in a competitive general election. The straight curves represent the distributions of ideological scores at the primary stage and the dashed curves represent the distributions of ideological scores among the same set of candidates at the general election stage.<sup>27</sup> Ideological scores are divided by their standard deviation at the primary stage.

We first observe that Democratic candidates tend to use left-wing language and Republican candidates right-wing language, both in the primary and the general election. This is somewhat mechanical since our method scores words primarily used by Democratic (Republican) candidates as left (right)-leaning. Second, both parties' distributions shift toward the center of the scale between the primary and general elections: the mean ideological score among Democrats shifts to the right by 0.42 standard deviation while the mean ideological score among Republicans shifts to the left by 0.33 standard deviation. Both estimates are significant at the 1% level, and they are not mechanical. These results indicate that the Democrats and Republicans that qualify for the general election tend to use more moderate language after the primary than before.

<sup>27</sup>As explained in Section 2.1, the website URLs of candidates present both in a competitive primary and a general election were collected through two different channels: the Wayback Machine search engine and the LoC, respectively. In this section, to ensure that we compare changes in ideological score from the same campaign websites (i.e., the same URLs) across election rounds, we focus on URLs taken from the LoC database and we use the captures of a candidates' general election website taken before the day of the primary election to calculate their primary ideological score. We use the captures of the same websites taken between the day of the primary and the general election to calculate the candidate's general ideological score.

**Figure 1: Ideology moderation**



Notes: We plot the kernel density of ideological scores for Democratic and Republican candidates in the U.S. (Figure 1a) and for the main political orientations in France (Figure 1b), pooling all election years together. The sample includes candidates who compete both in a competitive primary election and a competitive general election (Figure 1a), and candidates running both in a competitive first round and a competitive second round (Figure 1b). The solid curves represent the distributions of ideological scores in the first round and the dashed curves represent the distributions of ideological scores among the same set of candidates in the second round. In the U.S., candidates' ideological scores in the first round are calculated based on the captures of their general election website taken prior to the day of the primary election, while their scores in the second round are calculated based on the captures of their general election website between the primary and the general election.  $N=1,236$  candidates (Figure 1a) and 9,866 candidates (Figure 1b).

Similarly, Figure 1b plots the kernel distribution of ideological scores among French candidates running both in a competitive first round and a competitive second round (pooling all election years together), for each political orientation separately. We observe that candidates from left-wing parties use more left-wing language than candidates from centrist parties, who generally fall on the right side of the scale, but not as far to the right as candidates from right-wing parties. In addition, candidates from the far-right tend to use more extreme language than candidates from the mainstream right. Interestingly, candidates from far-left parties are located closer to the center than left-wing candidates, on average, revealing that the very few far-left candidates who qualify for the runoff tend to be relatively moderate.<sup>28</sup> Unlike in the U.S., these patterns are not mechanical, since centrist candidates were excluded from the construction of word scores and candidates from the right and the far-right (respectively left and far-left) were pooled together. As in the U.S., the distributions of scores shift toward the center in the second round. This shift is visible for all orientations, and particularly strong for far-right candidates, who shift to the left by 0.73 standard deviation on average.<sup>29</sup>

In both countries, these patterns are robust to restricting the sample to second rounds with only two competing candidates (Appendix Figure C.1),<sup>30</sup> which is the setting closest to the assumptions of

<sup>28</sup>By contrast, among all first-round candidates, including those who do not qualify for a runoff, the average ideological score of far-left candidates is -1.34 standard deviations, against -0.63 for left-wing candidates.

<sup>29</sup>We note that Figure 1b replicates results from Le Pennec (2020) on a larger sample, since hers includes only parliamentary elections between 1958 and 1993 while ours also includes candidates running in parliamentary elections between 1997 and 2022 and in local elections between 1979 and 2021.

<sup>30</sup>More precisely, we exclude general elections where smaller independent candidates are present and where a primary

the median voter theorem.

## 4.2 Complexity convergence

In addition to moderating their ideological tone, candidates may adjust the complexity of their language to appeal to a broader set of voters. For instance, politicians who catered to a highly-educated voter base in the first round may simplify their discourse in the second round to appeal to less-educated voters.

In Figures 2a and 2b, we plot the kernel density of complexity scores for the same samples of candidates as for ideology. We show two separate distributions, for candidates whose complexity score in the first round is below vs. above the median.<sup>31</sup> In both countries, we find that the distributions shift toward the center of the complexity scale. We observe the same pattern when we restrict the analysis to second rounds with only two candidates (Appendix Figure C.2).

A possible concern is that these patterns could be partly driven by a reversion to the mean. However, Appendix Figure C.3 shows similar results when we predict candidates' complexity score in the first round (based on a regression of their actual complexity on their observable characteristics) and use this *predicted* complexity score to separate candidates in two groups.<sup>32</sup>

Finally, we test whether candidates' propensity to moderate their ideology and their complexity are correlated. Appendix Figures C.4a and C.4b show a bin scatter plot of the mean complexity score against the mean ideology score in the first round as well as each bin's corresponding mean complexity and mean ideology in the second round. In the U.S., we observe that ideologically extreme candidates (either on the left or the right) tend to be more complex, whereas complexity does not vary as much in France. Despite these initial differences between countries, ideology and complexity tend to move toward the center in both settings, with larger adjustments for the most extremes candidates.

## 4.3 Topic convergence

A third way for candidates to appeal to a broader electorate is to expand the set of topics they discuss. Candidates who campaigned on very specific topics to appeal to certain voters in the first

---


election winner drops out before the general election (10 cases between 2002 and 2016), in the U.S.; and runoffs where more than two candidates are present in the second round as well as runoffs where two candidates qualify for the second round but one of them drops out of the race, in France.

<sup>31</sup>We compute the median complexity score in each election year separately.

<sup>32</sup>We predict complexity based on district fixed effects and candidate-specific variables: the candidate's party or political orientation, whether they are the incumbent, whether their party or political orientation won the previous election, their party or political orientation's vote share in the previous election, and the length of their website or manifesto.



**Figure 2: Complexity moderation**



Notes: We plot the kernel density of candidates' complexity score, pooling all election years together and splitting the sample between candidates whose complexity score in the first round is below the median score in a given election year, and those whose complexity score is above the median. Other notes as in Figure 1.

round may give more space to other topics that other voters care about in the second round.

We assess the prevalence of each topic in the candidate's website or manifesto and plot the kernel density of topic prevalence in each election round, for high- and low-prevalence candidates separately, and pooling across all topics.<sup>33</sup> Figure 3 shows convergence to the center in both countries. The shift is particularly striking in France, where candidates who insisted a lot on some topics in the first round reduce the prevalence of these topics by 0.44 standard deviation in the second round. We observe the same pattern when we restrict the analysis to general elections and second rounds with only two candidates (Appendix Figure C.5) and when we separate candidates based on their *predicted* topic propensities in the first round (Appendix Figure C.6).<sup>34</sup>

In sum, as their target electorate broadens and their number of competitors decreases, candidates move to the center of the ideology, complexity, and topics scales. In the next section, we show that candidates' convergence to the center is partly driven by their strategic convergence to their second-round opponent, which is the key mechanism underlying the median voter theorem.

## 5 Adjustment to Opponent


We now test whether candidates adjust their discourse to the opponent they are facing, exploiting races in which a candidate narrowly qualifies for the second round. We first provide two concrete

<sup>33</sup> Like for complexity, we distinguish high from low topic prevalence based on the median prevalence of that topic in the first round, in each election year separately.

<sup>34</sup> We use the same procedure and the same variables to predict topic propensities as we did for complexity.



**Figure 3: Topics moderation**



Notes: We plot the kernel density of candidates' topic prevalence, pooling all election years and topics together. For each topic, we split the sample between candidates whose topic prevalence in the first round is below the median topic prevalence in a given election year, and those whose topic prevalence is above the median. N=38,316 candidates×topics (Figure 3a) and 266,382 candidates×topics (Figure 3b). Other notes as in Figure 1.

examples of candidate adjustments.

In 2022, Franck Riester, a centrist candidate with an ideological score of 0.07, arrived first in the first round of parliamentary elections in the French *département* of Essonne. In the second round, he competed against François Lenormand, a far-right candidate with an ideological score of 0.93 who had narrowly qualified for the runoff after ranking second ahead of Cédric Colin, a left-wing candidate with an ideological score of -0.86. With Colin out of the race, Riester could reasonably expect left-wing voters to vote for him even if he did not make any effort to cater to them. Right-wing voters were more likely to hesitate between him and Lenormand. Moving his platform closer to Lenormand would help Riester persuade them and increase his vote share. Indeed, Riester shifted his manifesto to the right between the first and second round, reaching an ideological score of 0.30. Specifically, he started using right-wing words such as "community" and "defense" in the runoff, and stopped using left-wing words such as "minimum pension," "climate emergency," or "income tax." This closed the gap between his initial ideological score and that of his far-right competitor while increasing the gap with the left-wing candidate eliminated after the first round.

We now turn to a U.S. example related to candidates' complexity. In 2016, in the fourth congressional district of Tennessee, Steve Reynolds won the Republican primary with a complexity score of 0.27. In the general election, he faced Scott Desjarlais, who had narrowly won the Democratic primary with a complexity score of -0.08 against Grant Starrett, whose complexity score was 0.55. After the primary, Reynolds updated his website and decreased his complexity to 0.06, thus closing the gap with the complexity of his opponent as opposed to the complexity of the Democratic runner-up. He used less complex words (-0.41 standard deviation), fewer subordinates (-0.19 standard deviation), and slightly less diverse words (-0.04 standard deviation).

## 5.1 Regression discontinuity design

**Design** We now test more systematically whether candidates adjust their discourse to that of their competitor by estimating the impact for a candidate – the *leader* – to face a certain contender – the *opponent* – in the second round of the election, instead of another potential contender – the *runner-up*.

In France, remember from Section 2.2 that the set of candidates who qualify for the second round includes the two candidates who received the most votes in the first round and any other candidate who obtained the votes of at least 12.5% of the registered voters. We call the candidate ranked first in the first round the *leader*. We focus on races in which the vote shares of the second and third candidates were very close to each other and lower than the 12.5% qualification threshold, such that the second candidate qualifies for the second round but the third is eliminated. We call the second candidate the *opponent* and the third candidate the *runner-up*.

In the U.S., we exploit close primary elections in which the two top candidates obtained nearly exactly the same vote shares. We call the winner of the primary the *opponent* and the second candidate the *runner-up*. Each primary (Democratic or Republican) is linked to a *leader*, defined as the candidate of the opposing party in the general election.<sup>35</sup> For instance, the leader associated with a Democratic primary is either the Republican nominee or a contender from a third party, when no Republican candidate runs in the general election. Some primary races cannot be linked to a leader and are excluded from the sample, including primaries whose winner will run unopposed in the general election and elections in which several third party contenders run in the general election but no Republican candidate does.

Our outcome,  $Y_{i,l}$ , is a measure of discourse convergence between the leader  $l$  and opponent or runner-up  $i$  in the second round. If the leader strategically adjusts their platform to their opponent in order to attract undecided voters, we should expect the convergence between them to be stronger than between the leader and the runner-up. In general, this pattern could also emerge absent strategic adjustment. Indeed, the leader may decide to emulate their opponent's platform because they reason that this platform appealed to more voters than the runner-up's platform in the first round. Our RDD rules out this confounding mechanism by focusing on races in which the first round vote shares of the opponent and the runner-up were nearly identical and their platforms can thus be expected to be equally appealing to voters.

Formally, our design uses two observations per race, measuring the convergence  $Y_{i,l}$  between the leader and the opponent and the convergence between the leader and the runner-up, respectively. We define the running variable  $X$  as the difference between the vote shares obtained by the opponent

---

<sup>35</sup>Note that the leader is not necessarily the strongest candidate. We use this terminology by symmetry with the French context, in which the leader is the leading candidate in the first round.

and the runner-up in the first round. We set  $X$  as positive for the observation corresponding to the opponent, and negative for the observation corresponding to the runner-up. The treatment variable  $T$  is a dummy equal to one for the opponent ( $X > 0$ ) and 0 for the runner-up ( $X < 0$ ).<sup>36</sup>

We use a sharp regression discontinuity design and estimate the following equation:

$$Y_{i,l} = \alpha + \tau T_i + \beta X_i + \gamma T_i \times X_i + \epsilon_i. \quad (1)$$

We follow [Imbens and Lemieux \(2008\)](#) and [Calonico, Cattaneo and Titiunik \(2014\)](#) and estimate this specification non-parametrically, fitting a local linear regression on each side of the threshold within an optimal bandwidth selected by the MSERD procedure from [Calonico et al. \(2019\)](#). We cluster standard errors by district  $\times$  year. We report robust p-values ([Calonico, Cattaneo and Farrell, 2020](#)).<sup>37</sup> Our coefficient of interest,  $\tau$ , represents the causal effect, for the leader, of facing opponent  $i$  instead of the runner-up in the second round.

We define the outcome  $Y_{i,l}$  by computing the distance between the leader’s discourse in the first round and candidate  $i$ ’s discourse (also in the first round), the distance between the leader’s discourse in the second round and candidate  $i$ ’s discourse (still in the first round), and taking negative the difference between them:

$$Y_{i,l} = - \left( \underbrace{\left| Y_l^{(2)} - Y_i^{(1)} \right|}_{\text{Distance between leader } l \text{ at 2nd round and opponent } i \text{ at 1st round}} - \underbrace{\left| Y_l^{(1)} - Y_i^{(1)} \right|}_{\text{Distance between leader } l \text{ at 1st round and opponent } i \text{ at 1st round}} \right)$$

$Y_{i,l}$  takes a positive value if the leader moves their discourse toward candidate  $i$ ’s initial position between the first and second rounds.

Our design differs from standard close-election RDDs in two important ways. First, instead of using only one observation per constituency, with some constituencies falling above a threshold and others below, we use two observations per constituency, corresponding to the candidate above the qualification threshold (the opponent) and the candidate below (the runner-up). Second, close-election RDDs generally raise concerns of interpretation and external validity. By contrast, in our design, focusing on close elections is not just useful for identification. It enables us to compare a leader’s adjustment to two potential opponents, in a setting in which these opponents and their discourse did equally well in the first round, so that the only difference between them is that one is present in

<sup>36</sup>In 17 races in France and one in the U.S., the opponent and the runner-up obtained the exact same numbers of votes. We exclude these races from the sample.

<sup>37</sup>We use the R implementation of the rdrobust package ([Calonico et al., 2017](#)).

the second round and the other is not.

**Sample** Our sample includes all races in which we observe the opponent or runner-up’s discourse in the first round and the leader’s discourse in both rounds. More precisely, in France, our sample includes races where both the first and second-round manifestos of the leader, and the first-round manifesto of either the opponent or the runner-up are available. In the U.S., we measure the first-round position of the leader based on captures of their general election website taken before the day of the primary election; and their second-round position based on captures of the same website between the primary and the general election. Remember from Section 2.1 that we used the LoC database to collect the URLs of general election websites. Since the runner-up does not qualify for the general election and is not included in this database, we measure the first-round position of both the runner-up and the opponent based on captures of their primary election websites, whose URLs were collected through the Wayback Machine’s search engine. Therefore, our U.S. sample includes races in which the general election website of the leader and the primary election website of either the opponent or the runner-up are available.

In the U.S., if both the Democratic and Republican parties hold competitive primary races, a unique election and its two primary races can yield four observations, corresponding to the convergence between the Democratic (Republican) leader and their opponent and runner-up in the Republican (Democratic) primary. Since we cluster standard errors by district  $\times$  year, these four observations are included in the same cluster. In both countries, there are races in which we observe the opponent’s discourse but not the runner-up’s, or vice versa, yielding one observation instead of two.

Overall, our sample includes a total of 1,852 observations across 1,225 races in the U.S., and 1,409 observations across 807 races in France. Appendix Tables D.1 and D.2 show the number of races along with the average number of candidates and qualifying margin in the first round. The qualifying margin is 26 percentage points on average in U.S. primary elections and 3.4 percentage points in French first-round elections.

In principle, our analysis could be affected by endogenous sample selection. A first concern is if the first-round manifesto or primary website of opponents qualified for the second-round is observed more often than that of runner-ups. Column 1 of Appendix Table D.3 shows that this is not the case: there is no significant jump in the probability of having a first-round manifesto or website available at the qualification threshold, in either country.

A second important concern is if the leader’s decision to compete in the second round (instead of dropping out of the race) depends on the identity of the opponent. For instance, the leader may decide to stay in the race if the opponent they will face in the second round is extreme, and they may instead drop out if the opponent is moderate and thus deemed very likely to win. The latter

type of races would be excluded from the sample, since we would not observe the leader’s discourse in the second stage of the election. Fortunately, these cases are extremely rare and unlikely to affect our results. In France, no leader in our RDD sample ever drops out between the first and second rounds. In the U.S., only four primary election winners dropped out before the general election.

In Appendix E.1, we discuss these and other situations that could create endogenous sample selection at greater length, and provide empirical evidence that they are not a concern.

**Outcomes** Our main outcome  $Y_{i,l}$  is defined as the *overall* change in similarity between leader  $l$  and opponent or runner-up  $i$ , which aggregates standardized changes in vectorized text similarity as well as changes in similarity in ideology, complexity, and topic distribution. We also consider changes in similarity along each dimension separately.

Vectorized text and topic distribution are multidimensional vectors. Therefore, changes in the leader’s discourse between election rounds on these dimensions will impact their distance with the opponent and the runner-up differently.<sup>38</sup> To measure convergence in text similarity, we define the outcome as the change in cosine similarity between the leader  $l$  and the opponent or runner-up  $i$ ’s text vectors, divided by its standard deviation and averaged across different vector representations of text (see Appendix B for more details). To measure convergence in topics, the outcome is negative the change in Euclidean distance between leader  $l$  and opponent or runner-up  $i$ ’s vectors of topic prevalence (as defined in Section 3.3).

By contrast, the ideological score and the complexity score are both scalars. This can make comparing convergence to the opponent and to the runner-up challenging. If the leader is on the right of both the opponent and the runner-up in the first round and remains on their right in the second round, the change in similarity to both candidates will be exactly identical even if the leader tries to strategically adjust to the opponent. Therefore, when measuring convergence along these two outcomes, our main specification will restrict the sample to "middle point races": races in which the leader is initially between the opponent and the runner-up, either on the ideology or the complexity scale. In such races, changes in the leader’s discourse will affect their distance to the opponent and the runner-up in opposite ways.<sup>39</sup> For ideology, we set the outcome as negative the change in distance between leader  $l$  and opponent or runner-up  $i$ ’s ideological scores (as defined in Section 3.1). For complexity, the outcome is negative the change in distance between the candidates’ values of

<sup>38</sup>The leader’s movements will only reduce the gap with the opponent and the runner-up in equal amounts if they all fall on the same line in the  $n$  dimension space and if the leader’s position in the first and second round is either to the right or to the left (not between) the opponent and the runner-up. When  $n = 1$ , all three candidates are necessarily aligned. When  $n > 1$ , the leader will in general not be aligned with the opponent and the runner-up. Indeed, in a coordinate base including the opponent and the runner-up, this would mean that the leader’s coordinate is a perfect 0 along the orthogonal dimension to the hyperplane formed by them.

<sup>39</sup>Appendix Table D.7 shows that there is no systematic candidate sorting across the qualification threshold in the subsample of middle point races.

each complexity measure, divided by its standard deviation and averaged across the three measures of complexity (as defined in Section 3.2).


**Identification assumption** The validity of our RDD relies on the key assumption that first-round candidates of a certain type (e.g., candidates who are more similar to the leader) do not systematically sort on the right of the qualification threshold. Such manipulation is unlikely since it would require predicting the outcome of the first election stage with great accuracy, and unpredictable events make electoral outcomes uncertain.

We conduct several tests to bring empirical support for our identification assumption. First, we implement the test proposed by McCrary (2008) and verify that there is no discontinuity in the density of the running variable at the threshold. In our setting, this test would be satisfied by construction if data were available for all candidates, since we would have exactly two observations per race: one to the right of the threshold (the qualified opponent), and one to the left (the runner-up). Since our sample includes races for which we were able to collect the website or manifesto of one candidate but not the other, the McCrary test does not mechanically pass and remains informative. Appendix Figure D.1 shows that the null hypothesis of no sorting at the threshold cannot be rejected at conventional significance levels, whether in the U.S., in France, or when pooling the two samples together. Note that in this graph and in all analyses pooling both samples, we divide the running variable by its standard deviation in each country before pooling them together, in order to account for the fact that vote margins in France are tighter than in the U.S. Hence, the pooled running variable is measured in standard deviations.

Next, we conduct a general balance test by checking whether candidates' *predicted* treatment status jumps at the threshold. To predict treatment status, we first regress actual treatment on the following variables: a set of dummies indicating if the candidate is a Democrat (in the U.S.) or on the left (in France), if they are a woman, if they ran in the previous election in the same constituency, if they won that election, if their party or orientation won that election; the number of tokens in their website or manifesto; and their text similarity, ideological similarity, complexity similarity, and topic similarity to the first-round platform of the leader.

Figure 4 shows the results. Each dot represents the probability of being treated within a given bin of the running variable – i.e., the vote share difference between the qualified opponent and the defeated runner-up. Opponents are located to the right of the threshold, and runner-ups to the left. A quadratic fit on each side of the cutoff is provided as a visual assistance. Figure 4a does not show any discontinuity in the full sample pooling U.S. and French observations. This is confirmed by the point estimate shown in column 1 of Table 1, which is small and non-significant. Figures 4b and 4c do not show any jump in the predicted probability of being treated at the qualification threshold either, in the U.S. or the French sample taken separately, and the corresponding point estimates are

**Figure 4: General balance tests**



Notes: Dots represent the local averages of the predicted treatment status (vertical axis). Averages are calculated within quantile bins of the running variable (horizontal axis). The outcome is the candidate's predicted treatment status based on observable characteristics listed in the text. The treatment variable is a dummy equal to 1 if the candidate qualifies for the second round. The sample is restricted to candidates included in the RDD sample as described in the text. In Figure 4a, both countries' samples are pooled together and the running variable is divided by its standard deviation within each sample. In Figure 4b, the running variable is the vote share difference between the first two candidates in the primary election, and it is measured as percentage points. In Figure 4c, the running variable is the vote share difference between the second- and third-ranked candidates in the first round, and it is measured as percentage points. Continuous lines are a quadratic fit.

**Table 1:** General balance tests

Sample	Pooled (1)	U.S. (2)	France (3)
Treatment	0.005 (0.013)	0.017 (0.023)	0.015 (0.016)
Robust p-value	0.604	0.324	0.344
Observations left	1509	821	688
Observations right	1752	1031	721
Effective obs. left	855	356	427
Effective obs. right	881	375	440
Polyn. order	1	1	1
Bandwidth	0.587	0.205	0.045
Mean, left of threshold	0.481	0.451	0.484

Notes: Standard errors, shown in parentheses, are clustered by district  $\times$  year. We compute statistical significance based on the robust p-value and indicate significance at 1, 5, and 10% with \*\*\*, \*\*, and \*, respectively. The unit of observation is the candidate. The sample is restricted to candidates included in the RDD sample as described in the text. The outcome is the candidate's predicted treatment status based on observable characteristics listed in the text. In column 1, both countries' samples are pooled together and the running variable is divided by its standard deviation within each sample. We use local polynomial regressions: we fit separate polynomials of order 1 on each side of the threshold, using optimal bandwidths from the MSERD procedure. The mean, left of the threshold gives the mean value of the outcome for the runner-ups at the threshold.

not significant (Table 1, columns 2 and 3).

We also test whether there is a discontinuity in any of the individual variables used to predict treatment. Appendix Tables D.3 through D.5 show the results, both for the U.S. and France. All estimates are small and non-significant, except for the probability that the same political orientation – either the candidate themselves or another candidate from that orientation – won the previous election in France, which is significant at the 1% level (Appendix Table D.4, Panel b). Given the large number of tests that we conduct (12 per country, for a total of 24), finding one coefficient significant at the 5% level would be in line with what we would expect. Finding one coefficient significant at the 1% level is more concerning. Reassuringly, our main results are similar when controlling for this covariate (Appendix Table D.6).

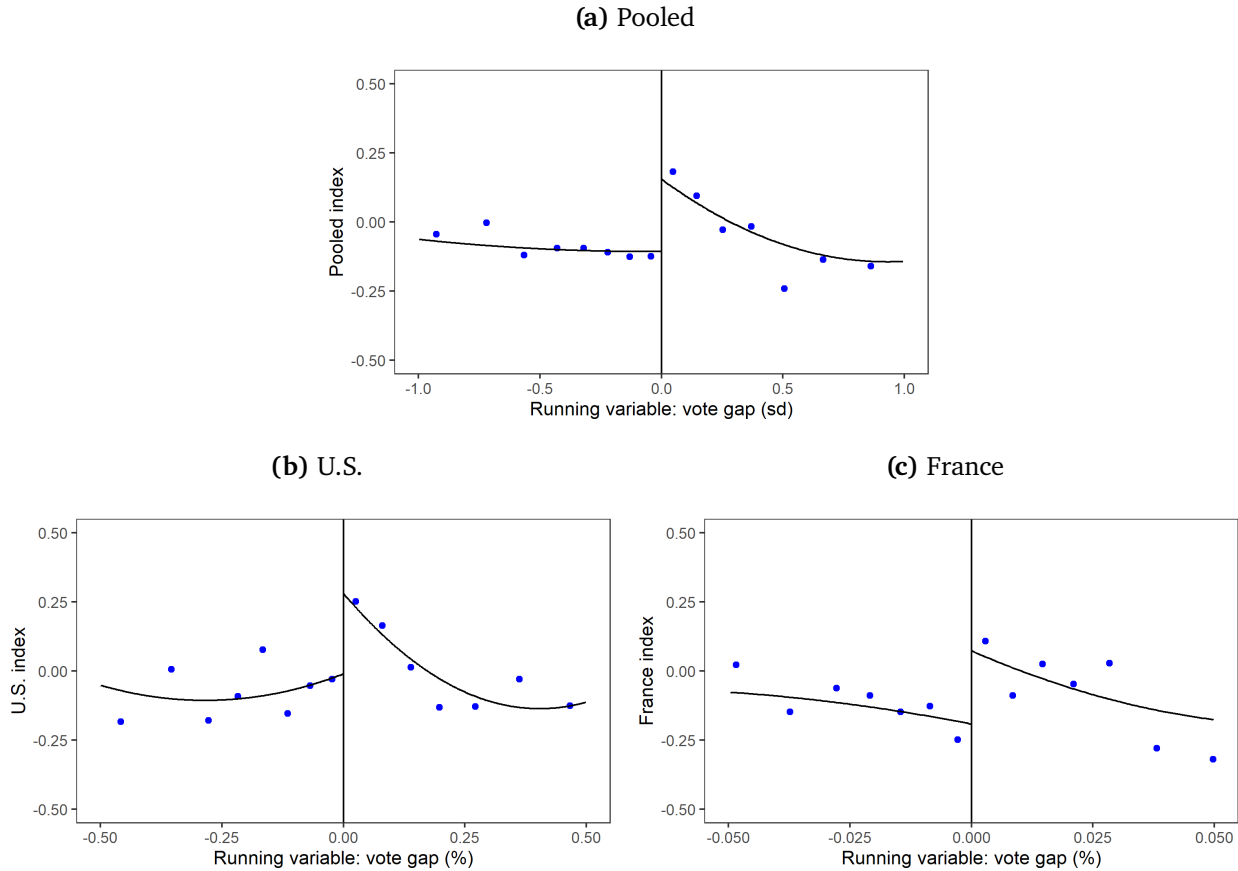
Finally, Appendix Figure D.2 shows that there is no discontinuity in the overall similarity between the candidate's and the leader's discourse in the first round. This provides reassuring evidence that the leader does not initially use language that is systematically more similar to the candidate who ends up qualifying as their opponent, as compared to the runner-up.



## 5.2 Main results

We first test for convergence to the opponent by measuring the impact, for a leader, of facing that opponent as opposed to their runner-up on the between-round change in overall similarity, which aggregates standardized changes in vectorized text similarity as well as changes in similarity in ideology, complexity, and topic distribution. Figure 5 shows that this outcome jumps up at the threshold in the U.S., in France, and in the pooled sample, indicating that leaders close the gap with their actual opponent more than with the runner-up who did not qualify. Table 2 complements the graphical analysis with formal estimates of the effects. Pooling both countries together, leaders' convergence to their actual opponent is 0.36 standard deviation larger than their convergence to the runner-up, which is significant at the 1% level. The effect is equal to 0.41 standard deviation and significant at the 1% level in the U.S. (column 2), and it is equal to 0.30 standard deviation and significant at the 5% level in France (column 3).

**Figure 5: Overall convergence**



Notes: The outcome is the change in overall similarity to the opponent or runner-up between election rounds, defined as the average of the standardized changes in vectorized text similarity as well as similarity in ideology, complexity, and topic distribution. It is constructed separately and divided by its standard deviation within each country. Other notes as in Figure 4.

These results are robust to using a wide range of bandwidths (Appendix Figure D.3). They are similar

**Table 2: Overall convergence**

Sample	Pooled (1)	U.S. (2)	France (3)
Treatment	0.357*** (0.105)	0.414*** (0.153)	0.300** (0.139)
Robust p-value	0.001	0.010	0.049
Observations left	1509	821	688
Observations right	1752	1031	721
Effective obs. left	855	425	417
Effective obs. right	881	439	432
Polyn. order	1	1	1
Bandwidth	0.588	0.253	0.044
Mean, left of threshold	-0.079	-0.049	-0.116

Notes: The outcome is the change in overall similarity to the opponent or runner-up between election rounds, defined as the average of the standardized changes in vectorized text similarity as well as similarity in ideology, complexity, and topic distribution. It is constructed separately and divided by its standard deviation within each country. Other notes as in Table 1.

if we restrict the sample to races in which exactly two candidates compete against each other in the second round, which is the setting closest to the assumptions of the median voter theorem (Appendix Table D.8).<sup>40</sup>

Next, we test for convergence to the opponent on each of the four textual dimensions separately to determine which of these dimensions drive the overall effect. The graphical results are shown in Appendix Figures D.4 and D.5 and the estimates are reported in Table 3.

Panel a shows the effects in the U.S. The effect on the change in vectorized text similarity is positive, sizeable (0.41 standard deviation), and significant at the 5% level (column 1), indicating that the leaders adjust their website to use language that is more similar to the language of their actual opponent, as compared to their opponent’s runner-up. By contrast, the effect on ideological convergence is positive but small and non-significant (column 2). Importantly, note that in the U.S., we cannot estimate an effect for the subsample of middle point races by ideology, since there are only 19 elections in which the leader’s primary ideological score falls between the ideological scores of the opponent and the runner-up. Indeed, it is rare for a Republican to use more left-wing language than either of the two Democratic candidates, and vice-versa. This limitation reduces our ability to test for strategic ideological convergence in the U.S. By contrast, races in which the leader’s first-round

<sup>40</sup>In France, we restrict the sample to races where exactly two candidates are present in the runoff, excluding races where either the leader or the qualified opponent drops out. In the U.S., we restrict the analysis to races where only the leader and the qualified opponent are present in the general election, without any other contender. These restrictions could be a source of endogenous sample selection: the identity of the qualified opponent may determine whether that candidate stays in the race to face the leader in the second round or decides to drop out instead, thus affecting the likelihood of a second round featuring exactly two candidates. In the U.S., the identity of the primary winner may also determine how many candidates choose to compete in the general election against that opponent. We discuss and provide evidence against these concerns in Appendix E.1.

complexity is initially in the middle of the opponent and the runner-up are much more common. Using that subsample, we find that leaders' complexity convergence to their actual opponent is 0.46 standard deviation larger than their convergence to the runner-up, which is significant at the 10% level (column 4). This pattern is consistent across all measures of complexity taken separately, although not as precisely estimated (Appendix Table D.9). Using the sample of all races, including those where the leader's initial complexity is not between the opponent and the runner-up yields a slightly smaller estimate (0.35 standard deviation) that remains significant at the 5% level (column 3). Finally, we do not find any significant convergence in topic distribution (column 5). In sum, U.S. candidates' convergence to their opponent is primarily driven by an adjustment of their linguistic complexity.

Panel b shows the effects in France. The effect on the change in vectorized text similarity is positive but smaller than in the U.S. (0.11 standard deviation) and non-significant (column 1). By difference with the U.S., the number of middle point races is sufficient to test for strategic convergence not just in complexity but also in ideology. Indeed, it is common for the ideology of the candidate ranked first in the first round to be between the ideology of the second and third candidates. Think for instance of a race in which these three candidates are in the center, on the left, and on the right, respectively. We find that French candidates' overall convergence to their opponent is primarily driven by convergence on ideology: while the point estimate is not significant for the full sample (column 2), restricting the sample to middle point races, in which we expect the largest effects, yields a large effect of 0.46 standard deviation, significant at the 1% level (column 3). By contrast, we do not find any evidence of convergence in complexity, including when restricting the sample to middle point races (column 5). Finally, column 6 shows a positive effect on convergence in topic distribution, equal to 0.26 standard deviation, and significant at the 10% level.

These results indicate that candidates in both countries are strategic. Rather than running purely on conviction, they adjust their discourse to their opponent. In the U.S., where it is harder to study convergence in ideology, politicians try to appeal to undecided voters by adjusting the complexity of their discourse toward that of their opponent without changing the topics they discuss. This specific type of strategic adjustment – on the style rather than the content – may result from the country's bipartisan setting and high level of polarization. In that context, changing topics may be perceived as flip-flopping and cost politicians votes, whereas adjusting their level of complexity may appeal to new voters without antagonizing the base. In the French multipartisan setting, ideological positions are more malleable and candidates adjust their ideological tone rather than their complexity to their opponent. They also change the relative importance that they give to different topics to match the distribution of topics in their opponent's discourse more closely.

Importantly, note that the convergence remains incomplete. In fact, if candidates fully converged to the median voter, the leader would not have any incentive to converge to the first-round platform

**Table 3: Convergence on different dimensions****(a) U.S.**

Outcome	Text similarity	Ideology	Complexity		Topics
	Full sample (1)	Full sample (2)	Full sample (3)	Middle points (4)	Full sample (5)
Treatment	0.412** (0.153)	0.153 (0.143)	0.353** (0.158)	0.455* (0.250)	0.076 (0.133)
Robust p-value	0.015	0.263	0.040	0.090	0.617
Observations left	821	821	821	207	821
Observations right	1031	1031	1031	207	1031
Effective obs. left	409	408	445	87	360
Effective obs. right	425	422	458	87	378
Polyn. order	1	1	1	1	1
Bandwidth	0.240	0.239	0.268	0.236	0.208
Mean, left of threshold	0.121	0.330	0.090	-0.064	-0.061

**(b) France**

Outcome	Text similarity	Ideology		Complexity		Topics
	Full sample (1)	Full sample (2)	Middle points (3)	Full sample (4)	Middle points (5)	Full sample (6)
Treatment	0.112 (0.129)	0.205 (0.124)	0.463*** (0.162)	0.133 (0.112)	-0.042 (0.174)	0.261* (0.152)
Robust p-value	0.444	0.132	0.007	0.335	0.871	0.095
Observations left	688	688	312	688	172	688
Observations right	721	721	312	721	172	721
Effective obs. left	485	356	186	399	100	415
Effective obs. right	496	364	186	412	100	430
Polyn. order	1	1	1	1	1	1
Bandwidth	0.054	0.033	0.043	0.041	0.036	0.043
Mean, left of threshold	-0.250	0.230	-0.047	0.001	-0.104	-0.020

Notes: The outcome is the change in similarity to the opponent or runner-up between election rounds, in terms of text similarity (column 1, Panels a and b), ideological score (column 2, Panel a, and columns 2 and 3, Panel b), complexity score (columns 3 and 4, Panel a, and columns 4 and 5, Panel b), and topic distribution (column 5, Panel a, and column 6, Panel b). Each outcome is divided by its standard deviation. In column 3, Panel b, the sample is restricted to races in which the leader is initially in the middle of the opponent and the runner-up on the ideology scale. In column 4, Panel a, and column 5, Panel b, the sample is restricted to races in which the leader is initially in the middle of the opponent and the runner-up on the complexity scale. Other notes as in Table 2.

of the opponent. No matter who qualifies for the second round, the leader should always converge to the median, expecting that the opponent would do the same. Instead, our design enables us to demonstrate the existence of a *partial* convergence equilibrium. Strategic convergence may be constrained by, for instance, party discipline, the cost of being seen as flip-flopping, and credibility concerns. Since the leader does not expect the opponent to fully converge, getting closer to them may swing voters who would hesitate between the two candidates otherwise.

### 5.3 Heterogeneity analysis

We now investigate the sources of variation in the level of adjustment across candidates. First, if candidates respond to incentives, they may adjust more when their chances of victory really depend on it. Second, some types of candidates may be willing to adjust more than others, e.g., if they care about winning more than about defending their ideas.

To test the first hypothesis, we ask whether leaders converge more to their opponent when the latter poses a greater threat to them. We use an alternative RDD to estimate the effect of facing a more extreme opponent as opposed to a more moderate one. We identify the most extreme of the two candidates that the leader could be opposed to in the second round – e.g., the most extreme of the top two candidates in the Democratic primary that the Republican nominee could have to face – and define the running variable as the difference in vote shares between them and the other potential opponent. The running variable is positive (and the treatment equal to 1) in races where the qualified opponent is more extreme, and it is negative (and the treatment equal to 0) when the qualified opponent is more moderate. Unlike in our main RDD, there is a unique observation per race. We use the leader’s between-round change in overall similarity to the qualified opponent as the outcome. By virtue of the RDD, leaders immediately to the left of the threshold can be expected to be comparable to those immediately to the right, except for the type of opponent that they face. If leaders believe that more extreme opponents have a smaller base and that they can attract fewer new voters in the second round, they may adjust less to them, which would yield a negative effect.

Table 4 shows the impact of facing a more extreme opponent, where extremeness is defined successively in terms of ideology and complexity. See Appendix Figure D.6 for the graphical evidence.<sup>41</sup> Although the effects are not statistically significant, they are negative and large in both countries, indicating that facing a more extreme opponent causes leaders to converge less to them. Specifically, convergence to the more ideologically extreme opponent is 0.46 standard deviation lower in the U.S. and 0.24 standard deviation lower in France (columns 1 and 3). We obtain similar results, although smaller in size, when we define opponent’s extremeness based on their level of complexity (Table 4, columns 2 and 4).

Leaders’ incentives to adjust may be higher not just when they face a moderate opponent but also when they face an incumbent or an opponent whose political orientation received more votes in the previous election. Using separate RDDs, we find that both of these treatments lead candidates to converge more to their opponent, in French elections (Appendix Table D.12).<sup>42</sup> Although sample

<sup>41</sup>We also verify that there is no discontinuity in the density of the running variable (Appendix Table D.10) and in the predicted treatment status (Appendix Table D.11) at the threshold.

<sup>42</sup>Appendix Tables D.13 and D.14 show that there is no discontinuity in the density of the running variable at the threshold and in the predicted treatment status, for either treatment. We cannot run this analysis in the U.S. since too few incumbents run in a competitive primary and since both potential opponents are from the same party, making the vote share of their orientation in the previous election mechanically identical. In France, the sample size is limited because we need

**Table 4:** Convergence to the more extreme opponent

Sample	U.S.		France	
	Extreme ideology (1)	Extreme complexity (2)	Extreme ideology (3)	Extreme complexity (4)
Treatment	-0.455 (0.271)	-0.165 (0.288)	-0.236 (0.221)	-0.208 (0.219)
Robust p-value	0.143	0.556	0.506	0.351
Observations left	327	320	372	309
Observations right	300	307	230	293
Effective obs. left	168	161	184	177
Effective obs. right	169	168	139	169
Polyn. order	1	1	1	1
Bandwidth	0.290	0.279	0.035	0.040
Mean, left of threshold	0.149	0.005	0.006	-0.057

Notes: The outcome is the change in overall similarity between the leader and the opponent between election rounds, defined as the average of the standardized changes in vectorized text similarity as well as similarity in ideology, complexity, and topic distribution. It is constructed separately and divided by its standard deviation within each country. There is one observation per race and the running variable is defined as the difference in vote shares between the more extreme candidate (either in terms of ideology or complexity) and the other potential opponent. It is positive in races where the qualified opponent is more extreme and negative in races where the qualified opponent is more moderate. Other notes as in Table 1.

sizes are much smaller, the effect of facing the incumbent is significant at the 10% level (column 1). Overall, these results confirm that candidates respond to incentives: they are more likely to adjust their discourse when their opponent is a stronger candidate who is likely to gather many votes.

We now turn to our second hypothesis and investigate differences in our main RDD effects across different types of candidates. In Appendix Tables D.15 and D.16, we replicate our main results from Section 5.2 for different subsamples of leaders, using the overall convergence index as the outcome. Although the effects are not all significant, we observe strategic convergence to the opponent across all types of candidates. In particular, the effects are comparable in size among left-wing and right-wing candidates (below and above the median ideological score).

Nevertheless, we do notice interesting differences. In the U.S., convergence tends to be larger among leaders who already ran in the past and incumbents (Appendix Table D.15, Panel a, columns 3 through 6) as well as leaders running in a district where their party received a large support in the previous election (Appendix Table D.16, Panel a, columns 3 and 4).<sup>43</sup> In France, it is larger among leaders who received a large number of votes in the first round (Appendix Table D.16, Panel b, columns 1 and 2). These patterns suggest that some candidates are consistently willing to adjust their discourse more than others, which contributes to their success *ex ante*, and explains why they continue to adjust more *ex post*.

---

to restrict the sample to races where either the second- or third-ranked candidate is the incumbent (column 1) and races where the second- and third-ranked candidates are from different orientations (column 2).

<sup>43</sup>Large support is defined as receiving more than the median vote share in the previous general election.

In sum, the extent to which candidates adjust their discourse to their opponent reflects both their type and the electoral incentives that they face. We finally check whether candidates' propensity to strategically converge to their opponent has changed over time and find that it is larger today. Splitting the sample period in two halves reveals that the effect is already substantial before 2008 in both countries (Appendix Table D.17, columns 1 and 3), but it is larger and only statistically significant after 2008 (columns 2 and 4).

## 6 Conclusion

This paper provides the first direct empirical test of the convergence mechanism underlying the median voter theorem. Using a novel database containing 9,000 candidate websites for the primary and general election of U.S. House of Representatives between 2002 and 2016, as well as 57,000 candidate manifestos issued for the first and second round of French parliamentary and local elections between 1958 and 2022, we derive two sets of results.

First, we find that, as their target electorate broadens from the first to second round, candidates present in both stages adjust their discourse and "move to the center." They do so by moderating their ideological tone as well as the complexity of their discourse, and by diversifying the topics they talk about.

Second, we show that this convergence to the center often results from candidates adjusting to the rival they will face in the second round, once they learn their identity: politicians "follow each other." Our RDD exploits races in which the identity of the rival is quasi-random. In some circumstances, the rival emerges after a close contest between very different politicians, so that strategic adjustments to the winner have the opposite sign to adjustments that would be made in response to the platform presented by the loser. An example from France involves a centrist politician who ranked first in the first round and, in the runoff, must face a left-wing politician who barely qualified ahead of a right-wing candidate. On average, we observe that the centrist candidate moves to their left between the first and second rounds in that case.

In the U.S., the system of primaries prevents identifying ideological adjustment to the opponent as the two possible rivals are always on the same side of the ideological spectrum. For instance, the Republican nominee always has both possible rivals in the general election emerge on their left, from the Democratic primary, so closing the gap with the candidate who won the primary would also close it with the candidate who lost. Accordingly we study an alternative dimension to ideology: the complexity of the language used. In many races, a politician a) faces a rival that narrowly defeated their opponent in the primary, and b) has a measure of language complexity that is between that of the two possible contenders. Then, the politician will generally change the complexity of their

discourse to match that of their opponent, simplifying it if the less complex opponent qualified, and moving in the opposite direction otherwise.

Pooling the U.S. and French samples and using an index that summarizes convergence in ideology, complexity, as well as topics and vectorized text, we find that the change in similarity to the qualified opponent is 0.36 standard deviation larger than the change in similarity to the runner-up.

Although politicians do not fully converge to the median voter, they do strategically get closer to each other. They may have their own ideas and convictions, but they certainly act on incentives to win elections nonetheless, just like Downsian models of electoral competition predict. Furthermore, strategic convergence is not limited to ideological tone and policy platform: politicians also adapt the complexity of their language. Beyond providing evidence on the strategies that candidates use to defeat their competitors, these results show that deciphering politicians' behavior requires studying not just the substance of their discourse but also its form.



## 7 References

- Abou-Chadi, Tarik. 2016. "Niche party success and mainstream party policy shifts—how green and radical right parties differ in their impact." *British Journal of Political Science* 46(2):417–436.
- Abou-Chadi, Tarik and Werner Krause. 2020. "The causal effect of radical right success on mainstream parties' policy positions: A regression discontinuity approach." *British Journal of Political Science* 50(3):829–847.
- Acree, Brice DL, Justin H Gross, Noah A Smith, Yanchuan Sim and Amber E Boydston. 2020. "Etch-a-Sketching: Evaluating the post-primary rhetorical moderation hypothesis." *American Politics Research* 48(1):99–131.
- Adams, James. 2012. "Causes and electoral consequences of party policy shifts in multiparty elections: Theoretical results and empirical evidence." *Annual Review of Political Science* 15:401–419.
- Adams, James and Samuel Merrill. 2003. "Voter turnout and candidate strategies in American elections." *The Journal of Politics* 65(1):161–189.
- Adena, Maja and Anselm Hager. 2022. Does online fundraising increase charitable giving? A nationwide field experiment on Facebook. Technical report WZB Discussion Paper.
- Anagol, Santosh and Thomas Fujiwara. 2016. "The runner-up effect." *Journal of Political Economy* 124(4):927–991.
- Ansolabehere, Stephen, James Snyder and Charles Stewart. 2001. "Candidate positioning in US House elections." *American Journal of Political Science* pp. 136–159.
- Benoit, Kenneth, Kevin Munger and Arthur Spirling. 2019. "Measuring and Explaining Political Sophistication through Textual Complexity." *American Journal of Political Science* 63(2):491–508.
- Benoit, Kenneth, Kohei Watanabe, Haiyan Wang, Paul Nulty, Adam Obeng, Stefan Müller and Akitaka Matsuo. 2018. "quanteda: An R package for the quantitative analysis of textual data." *Journal of Open Source Software* 3(30):774.  
**URL:** <https://quanteda.io>
- Besley, Timothy and Stephen Coate. 1997. "An economic model of representative democracy." *The quarterly journal of economics* 112(1):85–114.
- Bouton, Laurent, Julia Cagé, Edgard Dewitte and Vincent Pons. 2022. Small Campaign Donors. NBER Working Paper 30050.
- Burden, B. 2001. "The Polarizing Effects of Congressional Elections, Kin Galderisi, Ezra and Lyons, ed., Congressional Primaries and the Politics of Representation."

- Cagé, Julia, Caroline Le Pennec and Elisa Mougin. Forthcoming. “Firm Donations and Political Rhetoric: Evidence from a National Ban.” *American Economic Journal: Economic Policy* .
- Calonico, Sebastian, Matias D Cattaneo and Max H Farrell. 2020. “Optimal bandwidth choice for robust bias-corrected inference in regression discontinuity designs.” *The Econometrics Journal* 23(2):192–210.
- Calonico, Sebastian, Matias D Cattaneo, Max H Farrell and Rocio Titiunik. 2017. “rdrbust: Software for regression-discontinuity designs.” *The Stata Journal* 17(2):372–404.
- Calonico, Sebastian, Matias D Cattaneo, Max H Farrell and Rocio Titiunik. 2019. “Regression discontinuity designs using covariates.” *Review of Economics and Statistics* 101(3):442–451.
- Calonico, Sebastian, Matias D Cattaneo and Rocio Titiunik. 2014. “Robust nonparametric confidence intervals for regression-discontinuity designs.” *Econometrica* 82(6):2295–2326.
- Calvert, Randall L. 1985. “Robustness of the multidimensional voting model: Candidate motivations, uncertainty, and convergence.” *American Journal of Political Science* pp. 69–95.
- Canes-Wrone, Brandice, David W Brady and John F Cogan. 2002. “Out of step, out of office: Electoral accountability and House members’ voting.” *American Political Science Review* 96(1):127–140.
- Cascio, Elizabeth U and Ebonya Washington. 2014. “Valuing the vote: The redistribution of voting rights and state funds following the voting rights act of 1965.” *The Quarterly Journal of Economics* 129(1):379–433.
- Cattaneo, Matias D, Michael Jansson and Xinwei Ma. 2018. “Manipulation testing based on density discontinuity.” *The Stata Journal* 18(1):234–261.
- Covington, Michael A. and Joe D. McFall. 2010. “Cutting the Gordian Knot: The Moving-Average Type-Token Ratio (MATTR).” *Journal of quantitative linguistics* 17(2):94–100.
- Dale, Edgar and Jeanne S Chall. 1948. *A formula for predicting readability*. Columbus, O.: Bureau of Educational Research, Ohio State University.
- Dano, Kevin, Francesco Ferlenga, Vincenzo Galasso, Caroline Le Pennec and Vincent Pons. 2023. Co-ordination and Incumbency Advantage in Multi-Party Systems - Evidence from French Elections. Technical report Working Paper.
- Davis, Otto A and Melvin J Hinich. 1968. “On the power and importance of the mean preference in a mathematical model of democratic choice.” *Public Choice* pp. 59–72.
- Davis, Otto A., Melvin J. Hinich and Peter C. Ordeshook. 1970. “An Expository Development of a Mathematical Model of the Electoral Process.” *The American Political Science Review* 64(2):426–448.

- Devlin, Jacob, Ming-Wei Chang, Kenton Lee and Kristina Toutanova. 2018. "BERT: Pre-training of Deep Bidirectional Transformers for Language Understanding." *CoRR* abs/1810.04805.
- Downs, Anthony. 1957. "An economic theory of political action in a democracy." *Journal of political economy* 65(2):135–150.
- Druckman, James N., Martin J. Kifer and Michael Parkin. 2010. "Timeless Strategy Meets New Medium: Going Negative on Congressional Campaign Web Sites, 2002–2006." *Political Communication* 27(1):88–103.
- Druckman, James N., Martin J. Kifer and Michael Parkin. 2018. "Resisting the Opportunity for Change: How Congressional Campaign Insiders Viewed and Used the Web in 2016." *Social Science Computer Review* 36(4):392–405.
- Erikson, Robert S and Gerald C Wright. 1980. "Policy representation of constituency interests." *Political Behavior* 2(1):91–106.
- Fauconnier, Jean-Philippe. 2015. "French Word Embeddings."  
**URL:** <http://fauconnier.github.io>
- Flesch, Rudolph. 1948. "A new readability yardstick." *Journal of applied psychology* 32(3):221–233.
- Foarta, Dana and Massimo Morelli. 2021. "Complexity and the reform process."
- Fowler, Anthony. 2013. "Electoral and policy consequences of voter turnout: Evidence from compulsory voting in Australia." *Quarterly Journal of Political Science* 8(2):159–182.
- Fujiwara, Thomas. 2015. "Voting technology, political responsiveness, and infant health: Evidence from Brazil." *Econometrica* 83(2):423–464.
- Gaultier-Voituriez, Odile. 2016. "Archelec, les archives électorales de la Ve République, du papier au numérique." *Histoire@ Politique* (3):213–220.
- Gentzkow, Matthew, Jesse M Shapiro and Matt Taddy. 2019. "Measuring Group Differences in High-Dimensional Choices: Method and Application to Congressional Speech." *Econometrica* 87(4):1307–1340.
- Gerber, Alan S and Donald P Green. 2000. "The effects of canvassing, telephone calls, and direct mail on voter turnout: A field experiment." *American political science review* 94(3):653–663.
- Granzier, Riako, Vincent Pons and Clémence Tricaud. 2023. The Large Effects of a Small Win: How Past Rankings Shape the Behavior of Voters and Candidates. Technical report National Bureau of Economic Research.
- Grofman, Bernard. 2004. "Downs and two-party convergence." *Annual Review of Political Science* 7(1):25–46.

- Grossman, Emiliano. 2019. "The French Agendas Project." *Comparative Policy Agendas: Theory, Tools, Data* p. 90.
- Hahn, Lance W. and Robert M. Sivley. 2011. "Entropy, semantic relatedness and proximity." *Behavior Research Methods* 43(3):746–760.
- Hall, Andrew B. 2015. "What happens when extremists win primaries?" *American Political Science Review* 109(1):18–42.
- Hotelling, Harold. 1929. "Stability in Competition." *The Economic Journal* 39(153):41–57.
- Hurka, Steffen and Maximilian Haag. 2020. "Policy complexity and legislative duration in the European Union." *European Union Politics* 21(1):87–108.
- Husted, Thomas A and Lawrence W Kenny. 1997. "The Effect of the Expansion of the Voting Franchise on the Size of Government." *Journal of Political Economy* 105(1):54–82.
- Imbens, G. and K. Kalyanaraman. 2012. "Optimal Bandwidth Choice for the Regression Discontinuity Estimator." *The Review of economic studies* 79(3):933–959.
- Imbens, Guido W and Thomas Lemieux. 2008. "Regression discontinuity designs: A guide to practice." *Journal of econometrics* 142(2):615–635.
- Katz, Daniel Martin. 2013. "Measuring the Complexity of the Law: The United States Code." p. 41.
- Koh, Allison, Daniel Kai Sheng Boey and Hannah Béchara. 2021. Predicting Policy Domains from Party Manifestos with BERT and Convolutional Neural Networks. preprint SocArXiv.
- Laver, Michael, Kenneth Benoit and John Garry. 2003. "Extracting policy positions from political texts using words as data." *American Political Science Review* 97(2):311–331.
- Le Pennec, Caroline. 2020. Strategic Campaign Communication: Evidence from 30,000 Candidate Manifestos. SoDa Laboratories Working Paper Series 2020-05 Monash University, SoDa Laboratories.
- Lee, David S, Enrico Moretti and Matthew J Butler. 2004. "Do voters affect or elect policies? Evidence from the US House." *The Quarterly Journal of Economics* 119(3):807–859.
- Lehmann, Pola, Theres Matthieß, Nicolas Merz, Sven Regel and Annika Werner. 2021. "Manifesto Corpus." Version: 2021a.
- Levitt, Steven D. 1996. "How do senators vote? Disentangling the role of voter preferences, party affiliation, and senator ideology." *The American Economic Review* pp. 425–441.
- Levy, Gilat, Ronny Razin and Alwyn Young. 2022. "Misspecified Politics and the Recurrence of Populism." *American Economic Review* 112(3):928–962.

- Lindbeck, Assar and Jörgen W Weibull. 1987. "Balanced-budget redistribution as the outcome of political competition." *Public choice* 52:273–297.
- Martin, Lanny W and Georg Vanberg. 2007. "A robust transformation procedure for interpreting political text." *Political Analysis* 16(1):93–100.
- Martin, Louis, Benjamin Muller, Pedro Javier Ortiz Suárez, Yoann Dupont, Laurent Romary, Éric Villemonte de la Clergerie, Djamé Seddah and Benoît Sagot. 2020. CamemBERT: a Tasty French Language Model. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*.
- McCrary, Justin. 2008. "Manipulation of the running variable in the regression discontinuity design: A density test." *Journal of econometrics* 142(2):698–714.
- Meltzer, Allan H and Scott F Richard. 1983. "Tests of a rational theory of the size of government." *Public Choice* 41(3):403–418.
- Merz, Nicolas, Sven Regel and Jirka Lewandowski. 2016. "The Manifesto Corpus: A new resource for research on political parties and quantitative text analysis." *Research & Politics* 3(2):2053168016643346.
- Michel, Jean-Baptiste, Yuan Kui Shen, Steven Nowak, Martin A. Aiden, Erez Lieberman Aiden, Aviva Presser Veres, Adrian Gray, Matthew K. Pickett, Joseph P. Hoiberg, Dale Clancy, Dan Norvig, Peter and Jon Orwant. 2011. "Quantitative Analysis of Culture Using Millions of Digitized Books." *Science (American Association for the Advancement of Science)* 331(6014):176–182.
- Miller, Grant. 2008. "Women's suffrage, political responsiveness, and child survival in American history." *The Quarterly Journal of Economics* 123(3):1287–1327.
- Miller, Michael and Nicki Camberg. 2020. "U.S. House Primary Election Results (2012-2018).".
- MIT Election Data and Science Lab. 2017a. "U.S. House 1976–2020.".
- MIT Election Data and Science Lab. 2017b. "U.S. President 1976–2020.".
- Montoro, Rocío and Dan McIntyre. 2019. "Subordination as a potential marker of complexity in serious and popular fiction: a corpus stylistic approach to the testing of literary critical claims." *Corpora* 14(3):275–299.
- OpinionWay. 2017. Les Français et les programmes électoraux. Sondage OpinionWay pour Le Printemps de l'Economie.
- Osborne, Martin J and Al Slivinski. 1996. "A model of political competition with citizen-candidates." *The Quarterly Journal of Economics* 111(1):65–96.

- Palfrey, Thomas R. 1984. "Spatial equilibrium with entry." *The Review of Economic Studies* 51(1):139–156.
- Pettigrew, Stephen, Karen Owen and Emily Wanless. 2014. "U.S. House Primary Election Results (1956-2010).".
- Pons, Vincent. 2018. "Will a five-minute discussion change your mind? A countrywide experiment on voter choice in France." *American Economic Review* 108(6):1322–63.
- Pons, Vincent and Clémence Tricaud. 2018. "Expressive voting and its cost: Evidence from runoffs with two or three candidates." *Econometrica* 86(5):1621–1649.
- Poole, Keith T and Howard Rosenthal. 1985. "A spatial model for legislative roll call analysis." *American journal of political science* pp. 357–384.
- Shannon, Claude Elwood. 1948. *The mathematical theory of communication*. Urbana: University of Illinois Press.
- Spenkuch, Jörg L and David Toniatti. 2018. "Political advertising and election results." *The Quarterly Journal of Economics* 133(4):1981–2036.
- Spirling, Arthur. 2016. "Democratization and Linguistic Complexity: The Effect of Franchise Extension on Parliamentary Discourse, 1832–1915." *The Journal of Politics* 78(1):120–136.
- Tolochko, Petro and Hajo G. Boomgaarden. 2018. "Analysis of Linguistic Complexity in Professional and Citizen Media." *Journalism studies (London, England)* 19(12):1786–1803.
- Underhill, Wendy. 2017. "Primary runoff elections.". Accessed: 2022-07-29.  
**URL:** <https://www.ncsl.org/research/elections-and-campaigns/primary-runoff-elections.aspx>
- Wittman, Donald. 1977. "Candidates with policy preferences: A dynamic model." *Journal of economic Theory* 14(1):180–189.
- Wittman, Donald A. 1973. "Parties as utility maximizers." *American Political Science Review* 67(2):490–498.
- Woolley, John T and Gerhard Peters. 2017. "The American presidency project." Santa Barbara, CA. Available from World Wide Web: <http://www.presidency.ucsb.edu/ws>.
- Yamada, Ikuya, Akari Asai, Jin Sakuma, Hiroyuki Shindo, Hideaki Takeda, Yoshiyasu Takefuji and Yuji Matsumoto. 2020. Wikipedia2Vec: An Efficient Toolkit for Learning and Visualizing the Embeddings of Words and Entities from Wikipedia. In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing: System Demonstrations*. Association for Computational Linguistics pp. 23–30.