Measuring and Understanding Parties' Anti-elite Strategies

Hauke Licht, University of Innsbruck
Tarik Abou-Chadi, University of Oxford
Pablo Barberá, University of Southern California
Whitney Hua, University of Southern California

This article presents a new measure and analysis of parties' anti-elite appeals. In order to measure parties' anti-elite appeals, we apply crowd-sourced coding, supervised machine learning, and novel cross-lingual transfer learning techniques to parties' Twitter posts. Our dataset records quarterly estimates of parties' anti-elite strategies for 20 countries between 2008 and 2021. Based on these indicators, we analyze whether parties' anti-elite rhetoric reflects the potential costs and benefits of this electoral strategy. We find that mainstream parties use anti-elite rhetoric less frequently when they are more likely to be included in the next governing coalition. When challenger parties do well in the polls, they become more anti-elitist. Our article not only contributes to the literature on democratic competition by introducing and applying a new measure of anti-elite strategies, but it also outlines a novel, modular, and scalable procedure to measure party appeals using social media posts.

The years following the global financial, economic, and COVID-19 crises have seen a new phase of public resentment toward political elites as well as increasing vote shares for parties with anti-elite platforms. This type of rhetoric is characterized by the use of sweeping claims that negatively portray elites, defined to include not only politicians but also media outlets, corporations, interest groups, etc. The politicization of anti-elite appeals can be seen across the ideological spectrum as parties attempt to capitalize on people's dissatisfaction and frustration with "politics as usual." The election of political outsider Donald Trump and the decision of the United Kingdom to leave the European Union have generally been interpreted as triumphs for movements protesting against establishment politics. And in the current climate crisis, ecological and left parties, too, increasingly voice general elite criticism. For example, the Green Party of Canada (2019) argued that fighting climate change requires system change and accused established elites for blocking the required change because "you can't do that with the parties who have built and maintained the old system." These recent events and examples highlight the significance of anti-elite rhetoric for political competition in representative democracies.

Although research on anti-system, anti-elite, or populist parties has strongly increased recently (Hunger and Paxton 2022; Zulianello 2020), so far, only little research exists that examines the determinants of parties' anti-elite strategies (de Vries and Hobolt 2020; Kollberg 2024; Polk et al. 2017). Anti-elite appeals are frequently associated with being a core component of populism (Canovan 1999; Mudde 2004) but have rarely been comparatively studied as a party strategy. When researchers have analyzed populist or anti-elite parties, the goal has usually been to identify anti-elite or populist parties per se (e.g., Abedi 2004; Rooduijn et al. 2019) and not to determine under what conditions parties' positions tend to

Hauke Licht (hauke.licht@uibk.ac.at) is an assistant professor of computational political science at the Department of Political Science and the Digital Science Center of the University of Innsbruck, Austria. Tarik Abou-Chadi (tarik.abou-chadi@politics.ox.ac.uk) is a professor of European politics at the Department of Politics and International Relations at the University of Oxford and a professorial fellow at Nuffield College. Pablo Barberá (pbarbera@usc.edu) is an adjunct associate professor of political science and international relations, University of Southern California, Los Angeles, CA 94089. Whitney Hua (huawhitn@usc.edu) is a PhD graduate in political science and international relations, University of Southern California, Los Angeles, CA 94089.

Replication files are available in the *JOP* Data Archive on Dataverse (https://dataverse.harvard.edu/dataverse/jop). The empirical analysis has been successfully replicated by the *JOP* replication analyst. An online appendix with supplementary material is available at https://doi.org/10.1086/730711.

Published online January 10, 2025.

The Journal of Politics, volume 87, number 1, January 2025. Published by The University of Chicago Press for the Southern Political Science Association. https://doi.org/10.1086/730711

become more or less anti-elite focused (but see Polk et al. 2017). However, the fact that we can categorize a party as populist does not mean that it will always chose to emphasize anti-elite positions. Likewise, parties that are generally classified as populist may still employ anti-elite strategies under certain conditions. We identify two main reasons why the strategic use of anti-elitism by political parties so far is understudied. First, we lack an encompassing theoretical approach to when and how parties emphasize anti-elite positions. Second, we do not have measures of parties' anti-elite strategies that vary across parties and time and would thus allow us to empirically test our expectations about this behavior.

In this article, we introduce a new measure of parties' anti-elite strategies based on Twitter data. We combine crowd-sourced coding, supervised machine learning, and novel multilingual transfer learning techniques to quantify the salience of anti-elite rhetoric in the appeals of parties from 20 Western countries between 2008 and 2021 at the party–quarter level.

We provide a novel framework for when parties should turn to more anti-elite rhetoric. Anti-elite appeals can potentially help political parties attract new segments of the electorate. At the same time, they come with the risk of eroding support for governments—something that might be very costly for parties that expect to be part of a government coalition. Within this framework, polling results thus send parties two different kinds of signals. When parties go up in the polls, it is (a) a sign that their strategy is working and (b) that they are more likely to become part of the next government coalition. For challenger parties—parties that have never been in government before (de Vries and Hobolt 2020)—the positive signal of a successful strategy should outweigh the potential risks. They should thus use more anti-elite appeals. Mainstream parties, in contrast, will reduce their anti-elite rhetoric when they have a higher likelihood of being in the next government.

We use our novel granular indicator of parties' anti-elite appeals to test these expectations. We show that when challenger parties gain in the polls, they use more anti-elite appeals. In contrast, mainstream parties that have a higher coalition inclusion probability (see Kayser, Orlowski, and Rehmert 2023) reduce their anti-elite rhetoric.

In this article, we make several contributions. First, we present a new granular time-series cross-section dataset of parties' anti-elite strategies that covers parties from 20 Western democracies. Second, we contribute to the literature on the text-based measurement of parties' electoral strategies. We show that combining established crowd-coding and supervised text classification methods with methodological innovations such as annotation modeling and cross-lingual transfer learning enables quantifying the anti-elite strategies of parties from different countries, for an extended period, and at a high level of temporal granularity. This should facilitate new research into questions that had to remain unaddressed due to a lack of data or were limited in their scope by the time frames of available expert surveys and human-annotated party manifestos. In addition, by providing and testing a novel theoretical framework on parties' anti-elite strategies, our article also contributes to better understanding the ongoing transformation of party politics in advanced democracies. As politics are increasingly affected by anti-elite sentiment, we show how political parties—the most important actors in modern democracies—make strategic use of anti-elite appeals.

PARTIES' ANTI-ELITE STRATEGIES

Populism and anti-elitism

The electoral rise of anti-establishment parties and movements as well as the spread and success of populism throughout Europe suggest that anti-elite strategies play a crucial role in explaining the current transformations of European politics. The anti-elite appeals of political parties have often been discussed in the context of a growing literature on populism. In this literature, anti-elite rhetoric is considered a defining feature of populist appeals, and scholars converge on this view irrespective of whether they conceive of populism as a "thin-centered" ideology (Mudde 2004), a worldview (Müller 2017), a discourse (Hawkins 2009), a discursive style (Bonikowski and Gidron 2016; Jagers and Walgrave 2007), a discursive frame (Aslanidis 2016), or a communication strategy (Reinemann et al. 2019). Mudde (2004), for example, emphasizes that populism conceives "the people" and "the elite" as two monolithic, antagonistic societal blocks, denigrates the elite as "out of touch" and generally corrupt, and portrays their machinations as directed against the interest of the people (see Canovan 1999). Likewise, Aslanidis (2016, 1255) stresses that the portrayal of the elite as an unaccountable minority that uses their power against the popular will is a crucial element of populist discursive frames.[1]

Like populism, anti-elite rhetoric is oftentimes adopted by political actors of diverse ideological persuasions (van Kessel 2015). Left-wing actors, for example, tend to base their elite criticism on the distinction between capital and labor and thus typically target economic elites, the rich and wealthy, and all those that have control over corporations and multinational enterprises (March 2011). In contrast, the prime target of anti-elitism from the right is the "liberal elite" (Mudde 2007). Right-wing actors' enemies within the liberal elite may include anyone who is associated with progressive attitudes, ranging from feminists and left-wing politicians to vegans and hipsters. However, right-wing actors also increasingly include the "mainstream

1. Similar treatments of anti-elitism can be found in the literature on anti–political-establishment and anti-system parties (see Zulianello 2020).

media"; established institutions of liberal democracies, such as the courts; and supranational institutions in this liberal elite.

Much work exists that defines parties as populist or anti-establishment (Abedi 2004; Engler 2020; Rooduijn et al. 2019; Zulianello 2020) but usually does not investigate anti-elite appeals as a strategy that parties can use to varying degrees (but see Kollberg 2024). In fact, as we show empirically later here, all parties we examined use it to a certain extent. We seek to understand what explains this variation, which we theoretically formulate as a function of the costs and benefits associated with the use of anti-elite appeals, which in turn may vary in systematic ways as a function of the party's status and current electoral success.

Understanding parties' anti-elite strategies

A large literature has analyzed parties' programmatic strategies and has derived a number of behavioral patterns that result from changes in their incentives and informational environment. This literature has, for example, investigated if parties react to changes in public opinion (Adams et al. 2004), to election results (Budge 1994; Laver 2005), or to the success of niche and challenger parties (Abou-Chadi and Krause 2020). Generally, this literature postulates that by changing positions and issue emphasis from one election to the next, political parties can change how they are perceived by voters. Although parties generally stay true to their ideological core (Budge 1994) and voters might be slow to update their perceptions of parties' position (Fernandez-Vazquez 2014), the literature overall agrees that parties systematically adjust their programmatic strategies following a cost-benefit analysis of their electoral prospects (Adams 2012).

Hence, in order to better understand parties' anti-elite strategies, it is necessary to analyze the opportunity structures that affect the incentives that parties have to become more anti-elitist. Anti-elite strategies will allow parties to attract voters that largely base their voting decision not on policy preferences but on a general sentiment of grievance against the elite. Studies have indeed demonstrated that protest voting is an important element of support for populist parties (van der Brug, Fennema, and Tillie 2000). However, employing anti-elite strategies and appealing to protest voters does not come without a cost. Anti-elite appeals have the potential to erode support for specific parts of democratic representation. In contrast to policy appeals, parties need to take into account that anti-elite mobilization might backfire in the future.

The costs associated with parties' anti-elite strategies can be derived from their double role of representing and governing. As Mair (2013) has argued, parties face a dilemma of being responsive to public opinion versus governing responsibly. This dilemma has become more profound in a world that is strongly economically and politically integrated. National governments have less control over policy outcomes that affect their citizens. Anti-elitism strongly mobilizes people based on claims that elites are not responsive enough. Hence, anti-elite mobilization potentially backfires when it increases citizens' demands for responsiveness but parties cannot deliver. In addition, parties that have strongly competed on anti-elite messaging might become less attractive coalition partners. Why would other parties include a party in government that has contributed to eroding the credibility of exactly this institution?

Context conditions will determine how parties weigh these potential costs and benefits of anti-elite strategies. One reason for this is that parties do not have perfect information about electoral preferences. As a result, they may not fully know how strongly their anti-elite strategies will resonate with the electorate. Doing well (or poorly) electorally is the main signal that parties can receive about the appeal of their strategies (Adams et al. 2004; Budge 1994). When parties have electoral success, it increases their incentive to double down on their existing strategy (Laver 2005). In a world of nearly daily polling results, parties will not only focus on election results per se but can take into account how they are doing in the polls to adjust their behavior. Hence, when a party employs an anti-elite strategy and does well in the polls, it should feel encouraged to further emphasize anti-elitism. On the other hand, doing well in the polls is also an indicator that government participation becomes more likely. Vote and seat shares are the prime predictors of becoming part of a government coalition. From this perspective, when parties see their polls rising, they also have an incentive to tone down the anti-elite rhetoric. Anti-elitism might erode the credibility of the government that they could soon be part of themselves or it might reduce the probability of being included in this government in the first place.

In order to understand which perspective is dominant for a political party, we need to distinguish between challenger and mainstream parties (de Vries and Hobolt 2020). Challenger parties can be understood as parties that have never governed before, while mainstream (or "dominant") parties regularly alternate between government and opposition. As de Vries and Hobolt (2020) argue, successful challenger parties combine a strategy of issue entrepreneurship with an anti-elite appeal. Anti-elitism is a crucial part of a challenger party's appeal as it keeps mainstream parties from appropriating their issues. Challenger parties can be more credibly anti-elitist as they have never been part of a government coalition before. For many of these parties, it is also much less likely that they will join a governing coalition. They are often ostracized by other parties and even their own activists are not always necessarily interested in becoming part of a government. Hence, when challenger parties see their vote share rising in the polls, they should feel encouraged to double down on that strategy. The potential costs that are associated with becoming part of a governing coalition will matter less to them.

This is different for mainstream parties. As mainstream parties have governed before, anti-elite strategies are less credible for them and should overall contribute less to their electoral appeal. This is not to say that mainstream parties will never employ anti-elite strategies or that they cannot do it successfully. It just means that anti-elitism plays a less important part in their electoral strategy than for challenger parties. In addition, mainstream parties will anticipate much more what their electoral strategy could mean for government participation. That potential trade-off between vote- and office-seeking strategies is well-documented in the literature (Strøm 1990). When mainstream parties see their expected vote shares rising in the polls, this constitutes a strong signal that they are more likely to be part of the next government. Hence, when mainstream parties increase their vote shares, they should become less anti-elitist.

MEASURING PARTIES' ANTI-ELITE STRATEGIES

In order to test these expectations, we require a dataset that fulfills four criteria. First, the dataset should record time-varying measurements because purely cross-sectional data would not allow analyzing how parties adapt to changing strategic incentives. Second, the time series should be relatively granular because a single data point for each electoral cycle per party is not sufficient to observe the dynamic behavior we theorize. Third, the time series should span several electoral periods so that we can compare party behavior in different political configurations within the same system. Fourth and relatedly, the datasets should include measurements for parties from several countries so that we can compare parties across electoral contexts.

Before describing the measurement strategy we adopted to create such a dataset, below we briefly review existing studies and data sources and outline why the indicators they provide do not meet our needs.

Existing studies and data sources

In line with the literature on party behavior more generally, there exist two different kinds of sources of anti-elite strategy indicators: expert surveys and content-analytic studies. Among expert surveys, the Chapel Hill Expert Survey (CHES) was the first to include an anti-elite salience indicator (Engler, Pytlas, and Deegan-Krause 2019; Polk et al. 2017). Since its 2014 wave, the CHES has included this item twice: once in the general 2019 wave and once in a 2017 flash wave. The Populism and Political Parties Expert Survey (POPPA), conducted in 2018, adds an alternative source of expert ratings of parties' anti-elite strategies (Meijers and Zaslove 2021). Among content-analytic studies, measurements of parties' anti-elite strategies are typically a "byproduct" of scholars' efforts to quantify parties' degree of populism (see Hunger and Paxton 2022; Pauwels 2017).[2] The majority of these studies have adopted a manual content analysis strategy that relies on human coding to obtain these measurements (de Raadt, Hollanders, and Krouwel 2004; Ernst et al. 2017; Jagers and Walgrave 2007; Rooduijn, de Lange, and van der Brug 2014; Zulianello, Albertini, and Ceccobelli 2018). Others have adopted automated strategies like dictionary analysis (Gründl 2022; Oliver and Rahn 2016; Rooduijn and Pauwels 2011) or supervised text classification (Dai and Kustov 2022). The exception to these populism-centered studies is de Vries and Hobolt (2020), who focused on how challenger parties' issue strategies relate to their anti-elite strategies.

Unfortunately, these studies and data sources do not provide the granular time-series cross-national data we require to test our theoretical expectations. On the one had, existing content-analytical studies have a rather limited geographical and/or temporal coverage. Most studies focus on parties in only one country (Jagers and Walgrave 2007; Oliver and Rahn 2016; Reungoat 2010). And studies with a broader geographic coverage quantify parties' anti-elite strategies only over short periods (de Raadt et al. 2004; Rooduijn et al. 2014) or in relatively coarse intervals (Ernst et al. 2017; Zulianello et al. 2018). The CHES and POPPA expert surveys, on the other hand, have relatively broad geographic coverage[3] but provide either only cross-sectional data (POPPA) or relatively short and temporally coarse time series (CHES). Scholars thus currently lack the time-series dataset necessary to study the dynamics of party competition on the anti-elite issue.

We identify three methodological challenges that make measuring parties' anti-elite appeals hard. First, compared to other facets of parties' strategies, such as issue emphasis or position taking, quantifying anti-elite rhetoric is much more challenging. Tapping into this dimension of party competition with content-analytic methods is difficult because anti-elite discourse is context-dependent (see Taggart 2000) and thus manifests in many different expressions, tropes, and rhetorical figures in political text and speech. This implies that whether particular instances of words like "corrupt" indicate anti-elite rhetoric depends on the context in which they are used.

A second challenge that arises with cross-country studies is a rapid increase of cases and very often a language barrier (Licht and Lind 2023). Expert surveys and existing content-analytic strategies for measuring parties' anti-elite strategies commonly

2. Here we do not consider studies that measure populism without explicitly capturing its anti-elitism subdimension (e.g., Hawkins 2009; Bonikowski and Gidron 2016).

3. Both surveys cover many Western and Eastern European party systems (see also Engler et al. 2019).

struggle with one or both of these challenges (Laver 2014). And although dictionary-based strategies scale comparatively well to large amounts of texts (e.g., Gründl 2022; Rooduijn and Pauwels 2011), adapting dictionaries to new country contexts and languages is time-consuming and using machine translation (MT) to speed up this process is error-prone (Proksch et al. 2019). These limitations arguably explain why most content-analytic studies are monolingual and why established experts surveying projects provide only relatively coarse time series.

A third challenge is generating granular and relatively long time series of parties' anti-elite strategies. Conducting expert surveys more than every third or fourth year seems infeasible for logistical and financial reasons (Laver 2014). Traditional text data sources like party manifestos are too infrequently published to provide sufficient granularity. And using temporally more granular text sources like parties' press releases or their social media posts presupposes a measurement strategy that is scalable to large text corpora. Given these conflicting demands, it is not surprising that existing studies either focus on generating long but relatively coarse time series (e.g., de Vries and Hobolt 2020; Rooduijn et al. 2014) or granular but relatively short time series (e.g., Gründl 2022; Zulianello et al. 2018).

Our measurement strategy

Given the lack of suitable data sources and the difficulty of the measurement problem, we developed a novel measurement strategy to produce a granular time-series cross-national dataset of parties' anti-elite strategies. Our dataset covers parties with parliamentary representation in 20 different Western European and English-speaking countries[4] between 2008 and the beginning of 2021[5] and records quarterly estimates of the extent to which parties rely on anti-elite rhetoric in their electoral appeals.

We conceptualize anti-elite rhetoric as statements that negatively evaluate elites and/or their behavior and generalize this criticism to a wider elite group or "the elite" in its entirety. Anti-elite rhetoric is thus very similar to other attack behaviors like "valence attacks" (Jung and Tavits 2021) or negative campaigning (Walter, van der Brug, and van Praag 2014) that may also involve criticism of political elites. However, owing to its degree of generalization, anti-elite rhetoric advances a much broader criticism. The essence of this generalization is that it portrays the elite as a homogeneous, ill-intentioned group of people. Generalized elite criticism thus reflects the opinion that there are no meaningful differences among members of the elite in terms of their negative behaviors and attributes. This conceptualization is in line with the literature on populism discussed in the Parties' Anti-Elite Strategies section that considers anti-elite rhetoric to be a defining feature of populist appeals.

We designed a measurement strategy consisting of seven steps (see fig. 1) that address these challenges. To produce granular party-level time series, we used parties' Twitter posts as a data source. We first compiled a list of the Twitter accounts of parties in our sample and collected all their Twitter posts. To facilitate valid measurement of parties' anti-elite strategies from these very short and linguistically varied documents, we decided to rely on human annotation as primary input to our measurement procedure instead of a dictionary. In the next step, we sampled a subset of the tweets in our corpus and distributed them—together with our custom coding scheme—to crowd workers for human annotation. Crowd-sourced coding creates substantial scalability and cost-efficiency advantages compared with annotation by domain experts or trained coders (Benoit et al. 2016). Yet, crowd workers' judgments likely vary in quality. We hence applied a statistical annotation modeling method to aggregate our crowd-sourced coding into tweet-level labels. These labels are in turn used to train a supervised text classifier. To address the challenge of multilingualism, we applied novel cross-lingual transfer learning strategies developed in the natural language processing literature. Finally, we applied the resulting classifier to all posts and aggregated the resulting predictions into party–quarter estimates of parties' anti-elite strategies.

It is important to emphasize that our measurement strategy builds extensively on existing work in the quantitative text analysis literature (see Grimmer and Stewart 2013). For example, scholars have already used social media texts to quantify parties' electoral strategies (e.g., De Sio, De Angelis, and Emanuele 2018), and they often apply supervised classification (e.g., Barberá et al. 2021), crowd-sourced coding (Benoit et al. 2016), or a combination of these two methods (e.g., Widmann and Wich 2023) to produce measurements for large corpora. Yet, to improve these methods' performance in multilingual social media texts, we adopt methodological innovations like annotation modeling and cross-lingual transfer learning.

We are not aware of existing research that combines these components in the way presented here. However, we believe that researchers can apply our measurement strategy more broadly to quantify other facets of parties' electoral strategies. The crucial steps in our measurement strategy are discussed in greater detail below, where we describe how we constructed our dataset.

The party tweets corpus. We first compiled a list of the official Twitter accounts[6] of the parties in our sample and

4. These countries are Australia, Austria, Belgium, Canada, Denmark, Finland, France, Germany, Greece, Ireland, Italy, Luxembourg, the Netherlands, New Zealand, Norway, Portugal, Spain, Sweden, Switzerland, and the United Kingdom.

5. For more details, see sec. A of the data appendix.

6. There are now many projects compiling and maintaining such lists (e.g., Göbel and Munzert 2022; van Vliet, Törnberg, and Uitermark 2020). Where

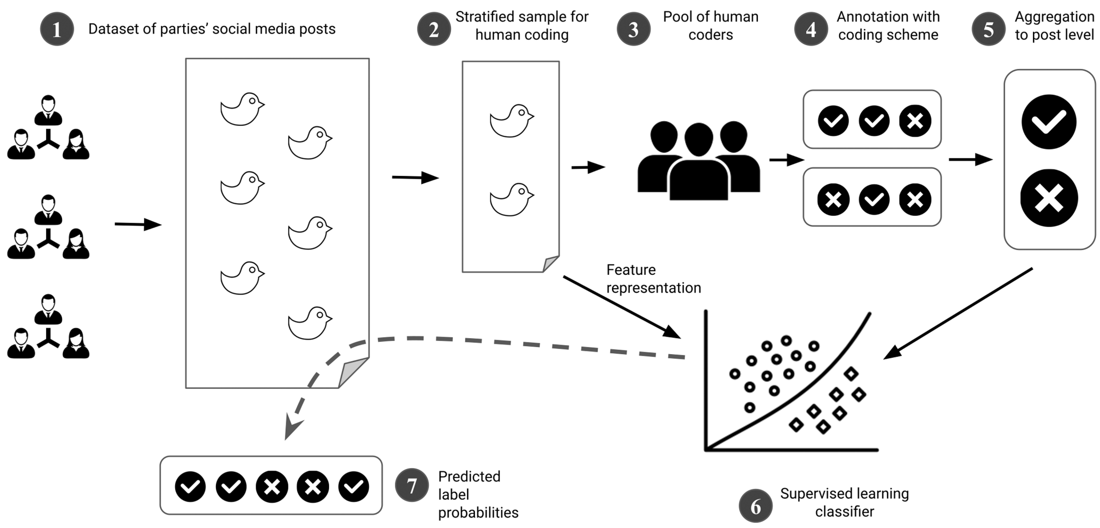

Figure 1. Flow diagram of our measurement strategy.

collected the tweets posted by these accounts during the periods the respective parties were represented in parliament.

For parties that (re)gained parliamentary representation during the period under study, we also collected tweets posted by their accounts in the two years before their parliamentary (re)entry. The resulting dataset contains a total of 786,763 tweets by 204 different parties.[7] We provide a detailed overview of the parties and numbers of tweets per party included in our dataset in sec. A of the Supporting Materials (SM).

Relying on the text of political parties' Twitter posts allows us to track changes in their anti-elite strategies at a high level of temporal granularity across a broad set of actors and countries. Several scholars have already leveraged this advantage to conduct innovative research on intra-party politics (Castanho Silva and Proksch 2022; Ceron 2017; Sältzer 2020), elites' responsiveness to public opinion (Fazekas et al. 2021), agenda setting dynamics (Barberá et al. 2019; Gilardi et al. 2022), parties' issue strategies (De Sio et al. 2018), etc. Further, relying on social media has considerable theoretical appeal. Social media has become the main communication channel by which politicians issue public statements, communicate with constituents, and attempt to influence the public agenda. De Sio et al. (2018) also note that social media both allows and incentivizes parties to adapt flexibly to over-time changes in their political environments. As a consequence, parties are more responsive to electoral incentives on these platforms than in their manifestos. This suggests that parties' social media communication is a nuanced trace of their strategic maneuvering and thus a valid data source for obtaining time-varying indicators of their electoral strategies.

Coding scheme design. Because we use supervised machine learning to quantify parties' anti-elite strategies, we require a sample of tweets categorized according to whether or not they exhibit anti-elite rhetoric. For this reason, we designed a custom coding scheme to collect such annotations from human coders.

Our coding scheme asks coders to determine for each tweet (a) whether it contains or approves of criticism of the elite and, if so, (b) who is the target of this criticism: a specific elite actor or an elite group more generally? Coders could indicate their judgment by selecting one of five different answer categories that combine these distinctions. To ensure that coders made their decisions solely based on the tweet's text, we presented them with the text of tweets one at a time, omitted images or other attached media, removed URLs and other hyperlinks from the tweets' texts, and did not reveal the account that posted the tweet (see Ennser-Jedenastik and Meyer 2018).

Importantly, our coding scheme prompts coders to make the crucial distinctions that characterize our conception of anti-elite rhetoric outlined above. The first distinction determines the presence or absence of *any* form of elite criticism. The second distinction separates "specific" from "general" elite criticism. The "specific" category includes narrowly targeted attacks such as negative campaign messages or other commonly studied valence attacks (Jung and Tavits 2021; Walter et al. 2014). The "general" category, on the other hand, includes broad-brush, generalizing criticisms of "the elite." In line with the conceptualization presented above, we consider only the latter to be indicative of anti-elite rhetoric. In the SM, sec. B.1, we

we could not rely on existing lists, we searched on parties' homepages or used Twitter's "list" feature (e.g., https://twitter.com/Congreso_Es/lists).

7. We thank one of the anonymous reviewers for pointing out a few accounts initially missing from our data. To address this gap, in May and June 2023, we added the accounts listed in table S.3 in sec. A of the SM. Because we added their tweets after our initial data collection, they were not included in the data used to sample tweets for annotation.

provide examples from our corpus that illustrate our distinction between "general" and "specific" elite criticism.

To ensure that the coding decisions our instrument elicits are faithful to our target concept, we tested and evaluated different versions of the coding scheme, and in each round of evaluation, we adapted our coding scheme to improve on its previous versions. Further details are available in the SM, sec. B.

Crowd-sourced coding. We deployed our coding instrument on Amazon Web Service SageMaker Ground Truth, which automates data allocations for annotation to Mechanical Turk (MTurk) crowd workers. Because MTurk crowd workers are English-speaking, we machine-translated all non-English tweets sampled for annotation to English using Google's *Cloud Translation API* before displaying them in the coding scheme.

Following an established best practice for crowd coding (Ipeirotis et al. 2014), we collected six annotations for each tweet. We decided on such a large number of annotations because SageMaker Ground Truth does not allow implementing standard quality control measures such as prescreening or continuous screening (see Benoit et al. 2016). However, the quality of annotations affects the quality of labels and, thus, supervised learning performance (Frenay and Verleysen 2014) and, indirectly, the quality of resulting measurements. We thus anticipated that we would need to remove a number of poor-quality annotations before aggregating them into tweet-level labels (see Hsueh, Melville, and Sindhwani 2009). Removing such annotations, we found that the labels aggregated from multiple crowd coding sets rival the quality of annotations provided by trained coders in our application—corroborating the finding by Benoit et al. (2016).[8]

To select tweets for crowd coding, we adopted a two-step stratified sampling strategy that allowed us to purposefully over-sample tweets likely to feature some form of elite criticism. We adopted this strategy because, based on past literature documenting the prevalence of anti-elite rhetoric in parties' communication (e.g., Jagers and Walgrave 2007), we expected that only a minority of the documents in our corpus would feature generalized elite criticism.

Randomly sampling documents for annotation when occurrences of the target phenomenon are rare results in "class imbalance," which makes it more challenging to train classifiers that reliably classify each label class (see Dai and Kustov 2022).[9] Accordingly, we implemented measures to facilitate the label class balance in our final annotated dataset. First, we distributed only tweets that would likely contain political content for annotation, which we determined by applying a pretrained supervised text classifier (Licht 2020).[10] Second, we mitigated repeated annotation of posts with very similar content by clustering the remaining tweets into 500 groups based on their multilingual embeddings and sampling tweets from all these "strata."[11] This heuristic allowed us to facilitate the coverage of our human-annotated sample relative to the underlying semantic space of our corpus. Third, we repeated these methods for stratified sampling in a two-step procedure. We first sampled 3,270 tweets and distributed them for crowd coding. After aggregating coding sets into tweet-level labels (see below), we then used this sample to train a classifier (see the SM, secs. C.3 and C.4). In the second step, we restricted the population of tweets eligible for sampling to thus-far unlabeled tweets this classifier predicted to contain elite criticism. From this subset, we sampled an additional 2,762 tweets and distributed them for crowd coding. Our two-step stratified sampling strategy resulted in a corpus of 6,032 human-coded tweets with six annotations each.

Annotation filtering and aggregation. To be able to train a supervised text classifier in the next step, we needed to assign one label to each tweet in this sample. We thus confronted the problem of aggregating multiple annotations per tweet into tweet-level labels (Chatterjee, Mukhopadhyay, and Bhattacharyya 2019). The most common approach to aggregate coding sets into labels in existing political and communication science studies is to determine labels by means of plurality voting (i.e., assigning the label that is most frequent among a document's annotations; e.g., Benoit et al. 2016). However, our analyses suggest that the crowd workers that contributed to our task varied in their commitment and ability to provide reliable judgments.[12] This led to low average agreement among coders[13] and meant that taking every coding at face value—as the plurality voting approach does—would have resulted in noisy tweet-level labels. We thus first filtered low-quality annotations from our coding data and then adopted an annotation modeling approach (see Passonneau and Carpenter 2014) to further mitigate the negative impact of noisy annotations on the quality of post-level labels. As a consequence, the labels we obtained by aggregating the "cleaned" crowd coding sets strongly correspond

8. See SM, sec. B.

9. In addition, stratifying our corpus by country when sampling tweets for human coding helps avoid bias, because anti-elite rhetoric is context-dependent and can be expressed using different phrases in different political cultures.

10. Elite criticism, like other negative political messages, was likely to be deemed political by crowd coders. Including nonpolitical tweets in the sample distributed for elite criticism coding would have inflated the class imbalance problem.

11. See SM, sec. C.1.

12. See SM, secs. C.3 and C.7.

13. Krippendorff's α is only 0.177 in the raw coding data.

to those assigned by trained coders as well as our "gold standard" judgments.[14]

The first step aimed to exclude annotations contributed by coders who did not have enough time or practice to develop good annotation skills or exhibited signs of low commitment to the task. Accordingly, we excluded coders who (a) contributed five or fewer coding sets in total, (b) exhibited little variation in their coding decision, and (c) removed all coding sets that were made in less than four seconds. Finally, we removed all coding of tweets for which "cannot answer" was the majority decision in the remaining data and all remaining "cannot answer" coding sets.[15]

In the second step, we fitted a Dawid–Skene per-annotator model (Dawid and Skene 1979) to aggregate our crowd-sourced coding set into tweet-level labels. Compared with inducing tweet-level labels by plurality voting, the Dawid–Skene model has the advantage that it estimates individual coders' abilities to correctly detect "true" labels. As a result, this model adjusts for coder-specific judgment biases when estimating tweet-level label class membership probabilities.

Analyzing the estimated label detection abilities confirmed our expectation that they vary considerably across coders.[16] Accordingly, we found that labels aggregated from crowd workers' coding sets agreed more frequently with labels aggregated from expert coding sets when adopting the annotation modeling approach than when applying plurality voting.[17] Given that the labels obtained with these two methods agreed in only 77.9% of tweets in our coding data, we thus concluded that model-based judgment aggregation not only yields systematically different labels than the plurality voting method, but, judging by experts' standards, it also yields systematically more reliable labels.

Inducing tweet-level labels by fitting the per-annotator Dawid–Skene model to the cleaned coding data, general elite criticism tweets make up 21.2% (1,271 tweets) of the labeled sample. Class imbalance in the pooled data is thus still in favor of the "no" label class (55.4% or 3,322 tweets), but our two-step stratified sampling strategy was effective because the share of tweets assigned to the elite criticism classes was higher in the second than in the first iteration (see tables S.5 and S.14 in the SM).

Classifier training and evaluation. The next-to-last step of our measurement strategy is to train a supervised text classifier using the set of labeled tweets. The goal of this step is to obtain a classifier that reliably discriminates between instances of "general" elite criticism and other tweets in held-out data. We thus discarded all tweets with model-induced "unsure" labels[18] and combined the "no" and "specific" elite criticism labels into a single class label. Dichotomizing the labels incentivizes the classifier to learn the distinction we are substantively most interested in, reduces problem difficulty, and, in our case, mitigates the class imbalance problem.[19]

To train our classifier, we adopted a transformer fine-tuning approach. Specifically, we fine-tuned a classification head stacked on top of an "XLM-T" model, a cross-lingual transformer-based neural language model that Barbieri, Espinosa Anke, and Camacho-Collados (2022) pretrained on a large, multilingual Twitter corpus. Transformer fine-tuning, and transfer learning more generally, has powered current advances in political text classification (e.g., Bestvater and Monroe 2023; Widmann and Wich 2023). The central idea of this approach is that a language model learns latent linguistic "knowledge" while pretraining on a large corpus, and a new classification model can then build on this knowledge (i.e., "transfer" it) to more effectively solve a new task. Further, because the XML-T model we build on is multilingual, we can leverage its learned cross-lingual representation capabilities when fine-tuning it to our data. We thus do not need to machine-translate the full corpus (see Licht 2023).

To provide benchmarks for the performance of this classifier, we also implemented two alternative approaches: supervised classification using tweet texts' multilingual sentence embeddings (MSEs) as inputs (Licht 2023) and a classic bag-of-words (BoW) classification approach that relies on MT to align tweets' text representations across languages. We discuss the advantages and limitations of these approaches in section D.1 in the SM.

Following standard practices in the literature, we randomly drew 80% of tweets to use them for classifier training. In this training set, we applied five-fold cross-validation to identify the best-performing hyperparameter values, using the macro F1 score to compare average performances between model runs.[20] Because transformer fine-tuning requires considerable computing resources, we cross-validated only the class weights.[21] In

14. See SM, sec. B, especially B.2.2, B.3.2, and B.4.2.
15. For more details on these decisions, see SM, secs. C.3 and C.7.
16. See SM, secs. C.3 and C.7.
17. See SM, sec. B.4.

18. We assume that the "unsure" coding sets result from coders' occasional lack of confidence in their judgments regarding the general–specific distinction.
19. "Specific" elite criticism is the rarest class. Collapsing it with "no" instances reduces imbalance.
20. We rely on the macro F1 score because it gives equal weight to precision and recall.
21. Increasing a class's relative weight in the loss function increases the cost of misclassifying its instances and thus incentivizes a classifier to correctly label them.

Table 1. Out-of-Sample Predictive Performance of "Best" Classifiers Trained with Different Text Representation and Learning Strategies

	Overall		General Elite Criticism Class		
Model	$F1_{\text{macro}}$	$F1_{\text{micro}}$	Precision	Recall	Specificity
XLM-T classifier	.745	.815	.644	.581	.893
MSEs (MLP)	.685	.737	.482	.660	.763
MT+BoW (w/ MLP)	.600	.681	.385	.460	.754

addition, for the MSE and MT+BoW-based classifiers, we cross-validated text preprocessing choices, the choice of learning algorithm, and a number of algorithm-specific hyperparameters. We refer readers to section D.2 of the SM for details. Finally, we trained models with the best-performing hyperparameters values using the complete training data and evaluated them on the 20% held-out test samples to estimate out-of-sample performance.

Table 1 reports the results of this evaluation. The XLM-T model (top row) performs best in terms of the F1 scores and precision. Overall, it achieves an accuracy (micro F1) of 81.5% and a macro F1 of 74.5%. Notably, however, this overall performance is partly due to the classifier's high reliability in labeling "negative" instances containing no or only specific elite criticism, as indicated by its high specificity. This is to be expected as the classification task we confront is hard. However, we find that up weighting the general elite criticism class during training is helpful. Consequently, the XML-T–based classifier correctly labels about six out of 10 instances of general elite criticism (recall), and it is correct almost two out of three times when assigning this label (precision).

We find this level of performance sufficiently reliable for our task for three reasons. First, as discussed in the SM, section D.3, this performance is comparatively good given that our documents are short and multilingual and that our classification task is hard, as indicated by the levels of inter-coder agreement. Second, in section D.4 of the SM, we present evidence that the XLM-T classifier can recognize country-specific phrases and words that are indicative of anti-elite rhetoric. This finding should mitigate the concern that using a single, cross-lingual classifier to quantify the anti-elite strategies of parties from different countries is unfeasible due to context dependence. Third, in our analyses below, we do not analyze tweets at the document level but aggregate tweet-level predictions at the party–quarter unit. We expect a large share of the classification error to cancel out as a result, which is supported by the validation analyses in the Validating Measurements section.

Aggregating tweet-level predictions into measurements. Having trained a classifier that predicts held-out data relatively reliably, we used it to generate predicted probabilities for the general elite criticism label class for all tweets in our corpus. In the subset of tweets whose content is predicted to be political, 19.1% are predicted to be instances of general elite criticism.

From these tweet-level predictions, we compute our central quantity of interest—the salience of anti-elite rhetoric in parties' electoral appeals—as the average predicted probability of political tweets belonging to the "general elite criticism" class. As described earlier, we focus specifically on this category because we want to exclude from our measurement any criticism of specific actors because such targeted attacks are different from the type of anti-elite rhetoric associated with, for example, populist discourse.

One could use various time units to obtain time-series estimates of this quantity of interest. If not stated otherwise, we select the quarter (i.e., a three-month interval) as the temporal unit of aggregation because in more granular temporal units (e.g., months or weeks), the number of tweets per unit is relatively low and more frequently zero, which would yield noisier estimates.[22]

Validating measurements

To demonstrate that our measurement strategy produces useful estimates of parties' anti-elite strategies, we summarize the results of different validity assessments. First, we show how our estimates vary across political parties. Figure 2 reports the average level of anti-elitism for Western European countries that have seen the emergence of populist parties on

22. Where the number of political tweets per party–quarter is zero, we indicate this with missing values.

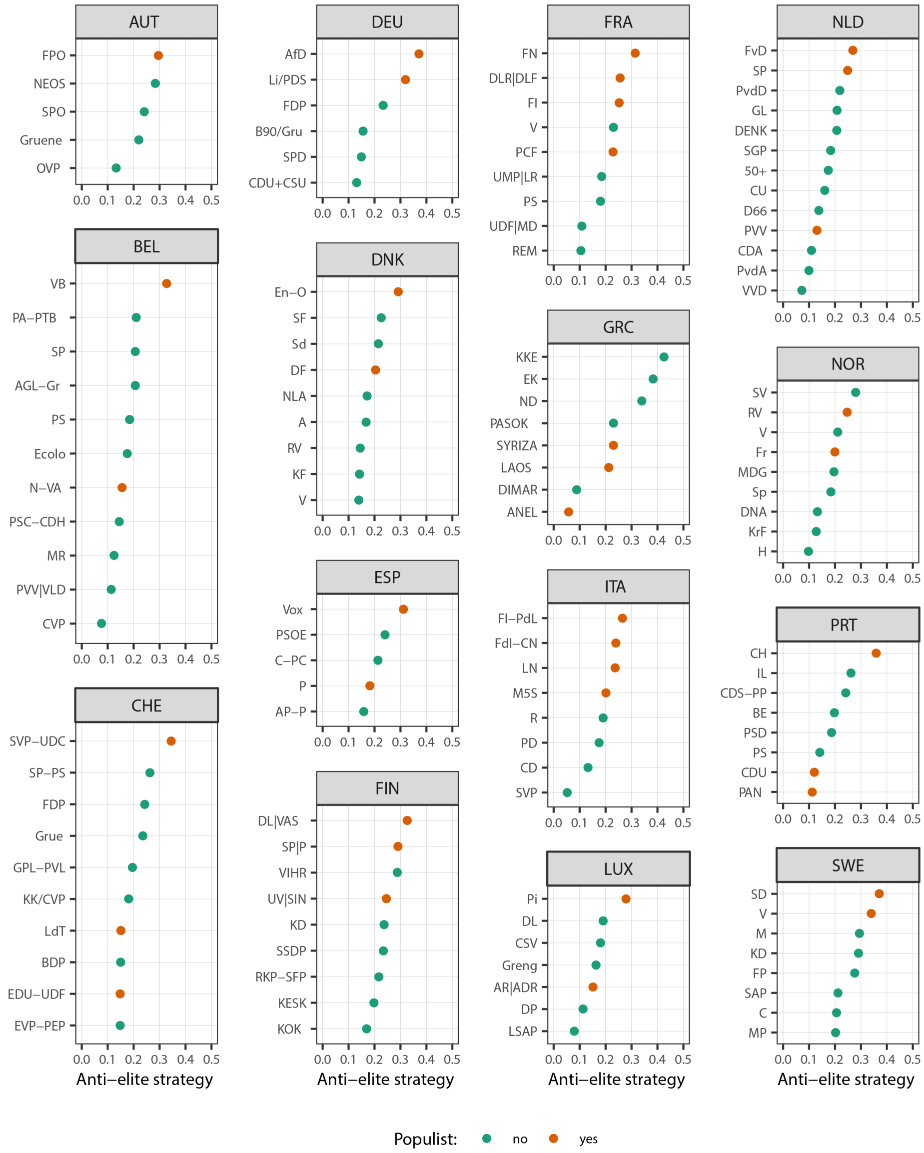

Figure 2. Estimates of parties' anti-elite strategies. Points indicate the mean predicted probability of "general" elite criticism across a party's tweets. Coloring distinguishes between populist (gray) and non-populist parties (black). For presentational purposes, we excluded regional parties in Spain and other minor parties.

the left and right.[23] We can see that the party that emphasizes anti-elite rhetoric the most is a populist party in all but two countries. In fact, in the sample of parties included in figure 3, populist parties are predicted to be significantly more likely to use anti-elite rhetoric than other parties (the two-sided t statistic is 4.02, $p < .01$). For example, we measure high values for the radical right Swiss Peoples Party (SVP) and the Dutch Forum for Democracy (FvD). Anti-elite rhetoric is not limited to radical right actors, however, as can be seen with the German Left Party (Li/PDS) or the French populist left La France Insoumise (FI). And although some populist parties do not rank highest in their countries, it is notable that these are often parties that have governed or supported a government during our period of investigation.[24]

Figure 3 plots our party time-series estimates of anti-elite strategies for three countries over time. Our new measure allows for much more temporally fine-grained analyses than many other measures of party positions. For example, we can document how anti-elite rhetoric became much more prevalent in

23. We distinguish between populist and nonpopulist parties following the literature and the PopuList project (Rooduijn et al. 2019).

24. For example, the Greek ANEL (January to September 2015), the Norwegian Fremskrittspartiet (Fr; September 2013 to January 2018), or the Finish UV–SIN (April 2015 to June 2017).

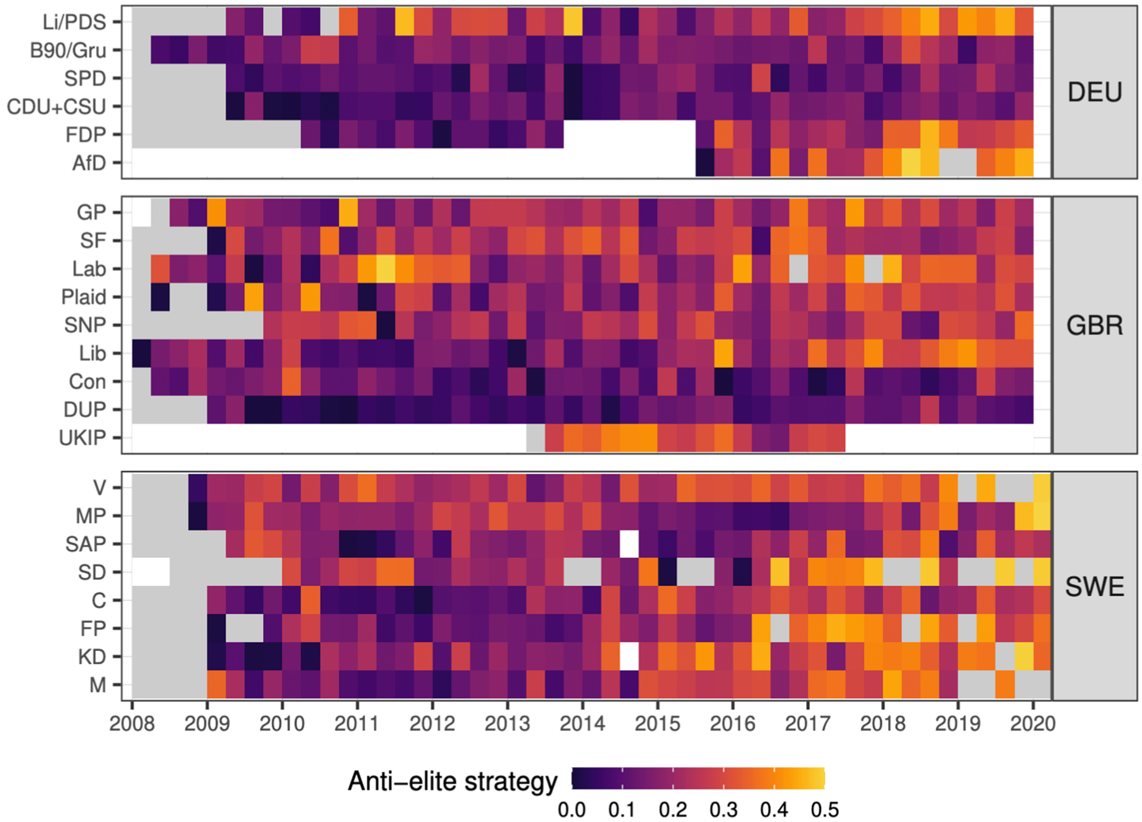

Figure 3. Quarterly estimates of German, British, and Swedish parties' anti-elite strategies. One square plotted per party–quarter estimate. Coloring of squares indicates tweets' mean predicted probability of "general" elite criticism. Lighter values correspond to higher values. Squares left blank (white) mark periods for which we did not include a given party in our dataset. Gray-shaded squares mark party–quarters for which no tweets are available to compute an estimate.

the British party system after Brexit[25] and that the Labour Party adopted a much stronger anti-elite strategy after the election of Jeremy Corbyn as party leader. In Germany and Sweden, we also observe that anti-elite discourse has become more salient – although in Germany to a lesser extent than in Sweden. These changes coincide with the success of the radical-right challengers Alternative for Germany (AfD) and Sweden Democrats (SD), respectively. This is a first indication that political parties indeed react to signals from the electorate about a demand for anti-elite rhetoric and adjust their strategies accordingly.

Next, we cross-validate our estimates against CHES data (Bakker et al. 2020; Polk et al. 2017). Specifically, we collected CHES anti-elite salience estimates for the parties contained in the 2014 and 2019 waves that are also recorded in our data. We then compared them to our estimates obtained by aggregating tweet-level predicted probabilities in the time window spanning the 12 months (four quarters) before the given CHES wave's field end date.[26] Figure 4 depicts this comparison in a scatter plot. The correlations of CHES indicators with our estimates are 0.47 for the 2014 wave and 0.37 for the 2019 wave, respectively.[27] These correlations confirm a positive association between these two sources of party strategy indicators.[28] Given that our estimates' goal is to pick up fine-grained variation over time, we should expect to see differences with respect to a measure based on expert judgments collected every couple of years. Further, in the SM, section E.2, we show that the convergence of dictionary-based measurements with the CHES indicators is usually lower. For example, applying the German anti-elitism dictionary compiled by Gründl (2022) results in lower convergence in both waves.

25. Given that our case selection strategy focuses on parties with parliamentary representation, we did not collect data for the Twitter account of UKIP as it lost parliamentary representation on June 8, 2017.

26. November 30, 2014, for the 2014 wave and January 31, 2020, for the 2019 wave.

27. Changing the number of quarters used to aggregate tweet-level estimates before a CHES wave's field end date does not substantively change these figures (see the SM, sec. E.1).

28. The strength of the correlation is comparable to that between manifesto-based measures and those on expert survey data or voting advice applications (da Silva et al. 2023).

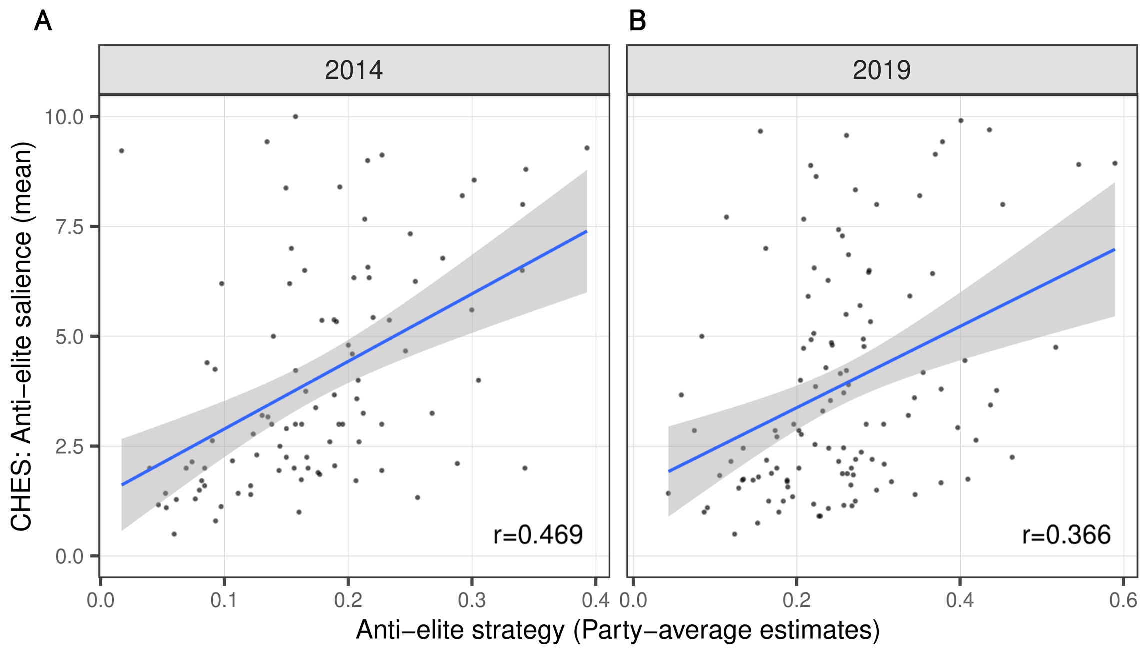

Figure 4. Estimates of parties' anti-elite strategies plotted against Chapel Hill Expert Survey (CHES) anti-elite salience indicators. Plot panel columns indicate CHES waves. Estimates of parties' anti-elite strategies were obtained by aggregating parties' tweets posted in the 12 months before the field end date of a given CHES wave. Note: Parties with less than 100 tweets in these date ranges are omitted.

PARTIES' STRATEGIC USE OF ANTI-ELITE APPEALS

Having discussed in detail how we measure parties' anti-elite strategies, we next outline how we tested our expectations about parties' strategic use of anti-elite appeals. We use two different main independent variables. First, we use monthly polling data collected by Kayser et al. (2023). We also included additional polls based on poll-of-polls estimates of POLITICO Europe (n.d.) to extend time coverage. Combined, these data sources cover 17 countries in our dataset (see the SM, sec. E.1, for details). Second, we use the *coalition inclusion probability* measure developed by Kayser et al. (2023).[29] This measure captures how likely it is that a party would be included in a governing coalition if a government was formed at a given moment in time. This measure thus specifically captures how much parties anticipate being part of the next governing coalition. Accordingly, we use this measure to address our mechanism about the costs and benefits of anti-elite strategies more directly. When the probability of being part of the next government coalition increases, mainstream parties should employ less anti-elite rhetoric.

We follow de Vries and Hobolt (2020) and operationalize challenger parties as parties that have never been part of a government coalition before. Mainstream parties, in contrast, are parties that currently govern or are in opposition but have governed before. We first show some descriptive evidence of how mainstream and challenger parties differ in their use of anti-elite appeals.

Figure 5 shows that challenger parties use significantly more anti-elite appeals than mainstream parties. This is the case in all but six countries. In four of the countries where the expected pattern does not hold (Austria, Italy, Greece, and Switzerland), radical right and radical left parties have governed for a significant time during the period under investigation. Hence, in line with de Vries and Hobolt (2020), we find that challenger parties generally rely more on anti-elite appeals than mainstream parties (or "dominant" parties in their terminology).

Next, we are interested in how challenger and mainstream parties react differently to the informational incentives from polling results and coalition inclusion probabilities. We test our theoretical expectations in a time-series cross-section dataset that records quarterly aggregates for each party.[30] We include the polling results and coalition inclusion probabilities lagged by one quarter. To deal with some of the challenges with time-series cross-sectional data, all our models include a lagged dependent variable and party fixed effects. We use panel-corrected standard errors. Given our different

29. Coalition inclusion probabilities are available for parties from 12 countries in our sample for varying periods (see the SM, sec. E.1).

30. We averaged polling numbers and coalition inclusion probability estimates at the party–quarter level to match our anti-elite strategy estimates.

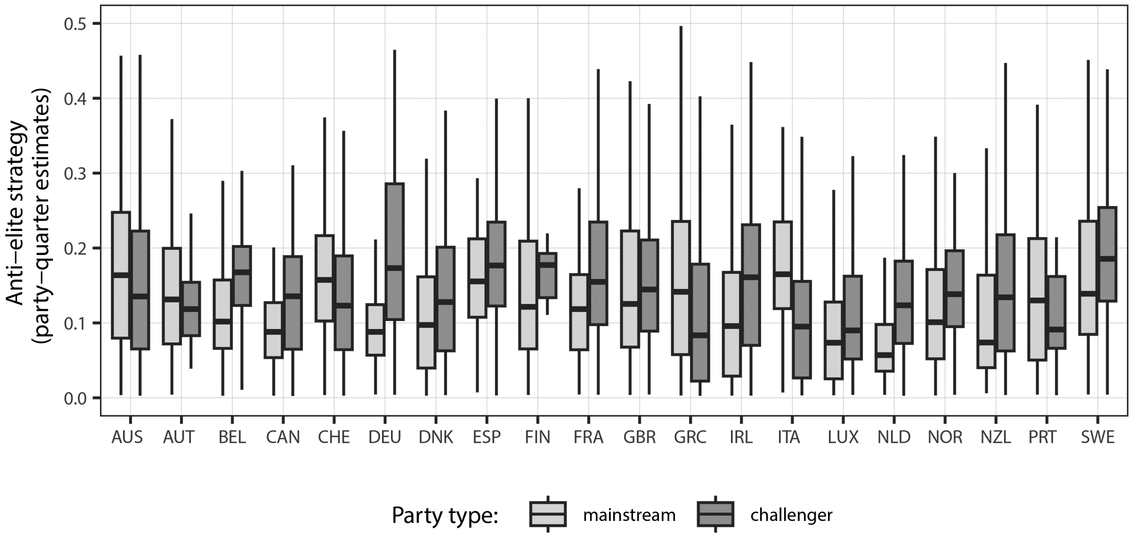

Figure 5. Distribution of party–quarter level estimates of parties' anti-elite strategies by party type. Data points outside of 1.5 × interquartile ranges not plotted.

expectations for challenger and mainstream parties, we fit separate models for them.[31]

Figure 6 shows the predicted values of changes in the polls on parties' anti-elite strategies for mainstream parties (left) and challenger parties (right), holding all other variables at their means. For mainstream parties, we find a negative effect of polling results on the use of anti-elite appeals. When mainstream parties did better in the polls in the last quarter, they reduced their use of anti-elitism. We argue that polling results send mixed signals to mainstream parties. On the one hand, they signal that political strategies are working out. On the other hand, they make it more likely that a party will be part of the next government. In line with our expectations, mainstream parties react more strongly to a signal of government participation and reduce their use of anti-elite appeals when they do better in the polls. This effect, however, is rather small and thus reflects the potentially conflicting signals that mainstream parties receive from the polls.

In contrast, the right panel shows that when challenger parties are doing better in the polls in the previous quarter, they make more use of anti-elite appeals. Doing well in the polls signals to challenger parties that anti-elite appeals resonate well with the electorate, which leads them to more strongly emphasize anti-elitism. In line with our expectations, considerations about coalition inclusion seem to be less relevant for challenger parties.

In figure 7, we more directly test how the likelihood of being included in the next coalition affects parties' use of anti-elite appeals. We show the predicted values of anti-elite appeals conditional on the coalition inclusion probability for mainstream (left) and challenger parties (right). Mainstream parties significantly reduce their use of anti-elite appeals when they are more likely to be included in a government coalition. In contrast, challenger parties do not significantly change their anti-elite appeals in reaction to changes in the likelihood of being included in a coalition.

CONCLUSION AND DISCUSSION

Anti-establishment parties and anti-establishment movements have shaken up the politics of many advanced democracies. Their popularity has led to anti-elite rhetoric becoming a constant feature of these polities. In this article, we advance our understanding of this type of rhetoric by developing a new measure of parties' anti-elite appeals in 20 different Western European and English-speaking countries from 2008 to 2020. We create this measure using political parties' tweets, which allows us to generate fine-grained time-series data about when parties use anti-elite appeals.

Here, we apply this novel measure to analyze if parties change their anti-elite strategies in response to informational incentives. We show that when challenger parties gain in the polls, they become more likely to use anti-elite appeals. When mainstream parties are more likely to be included in the next governing coalition, they use fewer anti-elite appeals. Most importantly, our findings demonstrate that parties indeed use anti-elite appeals strategically. We should thus think of

31. Please refer to the SM, sec. E.1, for the ordinary least squares coefficient estimates of these models.

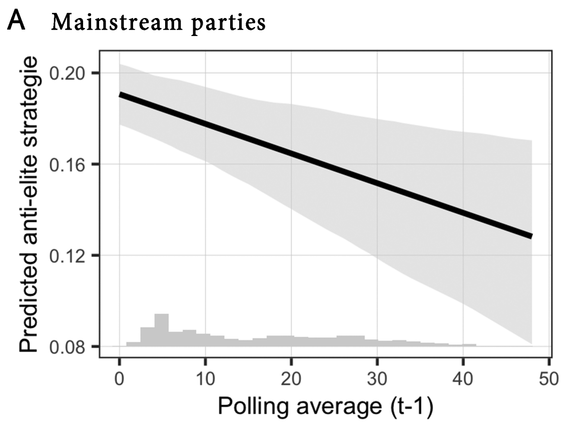

Figure 6. Predicted anti-elitism conditional on polling ($t - 1$).

anti-elitism less as an unchangeable, binary characteristic of political parties and more as an option in parties' strategic repertoire. Treating anti-elite appeals more like other political positions and issues will allow researchers to integrate them into theories of party competition and develop them further.

Our empirical strategy does not allow us to fully unpack the information feedback mechanisms that underlie the relationship between parties' polling results and their use of anti-elite rhetoric or how any changes in rhetoric may lead to changes in public support. However, a natural next step for our work would be additional research that takes full advantage of the temporally granular nature of our data, combined with individual poll results. This type of detailed time series analysis could be used to develop a deeper understanding of what drives the connection between parties' polling results and their use of anti-elite appeals. This analysis could potentially be complemented with other sources of text, such as press releases, media interviews, and social media posts on other platforms, in order to obtain a fuller picture of parties' communication strategies.

Our contribution is equally methodological. Research on democratic representation and party competition is often hindered by a lack of time-series cross-section data on parties' electoral strategies. This article combines crowd-sourced annotation, supervised text classification, and cross-lingual transfer learning to obtain such estimates from political parties' social media posts. Our results demonstrate that our approach is methodologically sound and leads to valid estimates of parties' communication choices in ways that may not have been feasible until recently, particularly regarding our ability to apply this method to multiple languages. Our measure will thus allow researchers to study in more detail when parties across different countries use anti-elite appeals, when these appeals are more successful, and how they generally affect democratic politics and representation.

Although we focus here on quantifying parties' anti-elite strategies, future work could apply our measurement strategy more widely to study other facets of parties' electoral strategies, such as issue attention and agenda-setting dynamics. Although these dimensions have been studied extensively

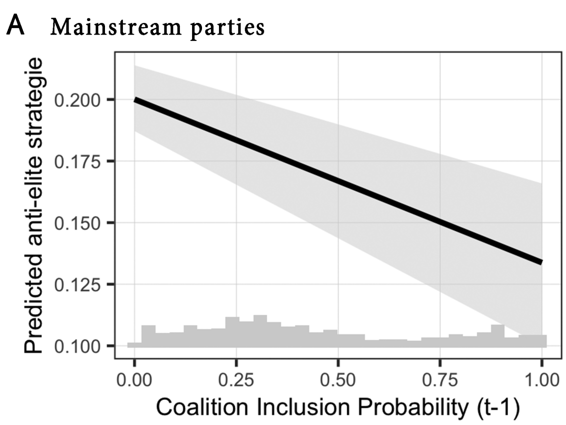

Figure 7. Predicted anti-elitism conditional on coalition inclusion probabilities ($t - 1$).

before, our approach introduces innovations that should improve the validity of past measurement approaches and reduce the costs associated with their application, particularly in cross-country analyses. Our method thus substantially expands the tool kit of researchers who study political parties and other actors.

REFERENCES

Abedi, Amir. 2004. *Anti-Political-Establishment Parties: A Comparative Analysis*. Routledge Studies in Extremism and Democracy. London: Routledge.

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): 829–47.

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 (1): 401–19.

Adams, James, Michael Clark, Lawrence Ezrow, and Garrett Glasgow. 2004. "Understanding Change and Stability in Party Ideologies: Do Parties Respond to Public Opinion or to Past Election Results?" *British Journal of Political Science* 34 (4): 589–610.

Aslanidis, Paris. 2016. "Is Populism an Ideology? A Refutation and a New Perspective." *Political Studies* 64 (1): 88–104.

Bakker, Ryan, Liesbet Hooghe, Seth Jolly, Gary Marks, Jonathan Polk, Jan Rovny, Marco Steenbergen, and Milada Anna Vachudova. 2020. *2019 Chapel Hill Expert Survey*. https://www.chesdata.eu/2019-chapel-hill-expert-survey.

Barberá, Pablo, Amber E. Boydstun, Suzanna Linn, Ryan McMahon, and Jonathan Nagler. 2021. "Automated Text Classification of News Articles: A Practical Guide." *Political Analysis* 29 (1): 19–42.

Barberá, Pablo, Andreu Casas, Jonathan Nagler, Patrick J. Egan, Richard Bonneau, John T. Jost, and Joshua A. Tucker. 2019. "Who Leads? Who Follows? Measuring Issue Attention and Agenda Setting by Legislators and the Mass Public Using Social Media Data." *American Political Science Review* 113 (4): 883–901.

Barbieri, Francesco, Luis Espinosa Anke, and Jose Camacho-Collados. 2022. "XLM-T: Multilingual Language Models in Twitter for Sentiment Analysis and Beyond." In *Proceedings of the Thirteenth Language Resources and Evaluation Conference*. Paris: European Language Resources Association, 258–66.

Benoit, Kenneth, Drew Conway, Benjamin E. Lauderdale, Michael Laver, and Slava Mikhaylov. 2016. "Crowd-Sourced Text Analysis: Reproducible and Agile Production of Political Data." *American Political Science Review* 110 (2): 278–95.

Bestvater, Samuel E., and Burt L. Monroe. 2023. "Sentiment Is Not Stance: Target-Aware Opinion Classification for Political Text Analysis." *Political Analysis* 31 (2): 235–56.

Bonikowski, Bart, and Noam Gidron. 2016. "The Populist Style in American Politics: Presidential Campaign Discourse, 1952–1996." *Social Forces* 94 (4): 1593–621.

Budge, Ian. 1994. "A New Spatial Theory of Party Competition: Uncertainty, Ideology and Policy Equilibria Viewed Comparatively and Temporally." *British Journal of Political Science* 24 (4): 443–67.

Canovan, Margaret. 1999. "Trust the People! Populism and the Two Faces of Democracy." *Political Studies* 47 (1): 2–16.

Castanho Silva, Bruno, and Sven-Oliver Proksch. 2021. "Politicians Unleashed? Political Communication on Twitter and in Parliament in Western Europe." *Political Science Research and Methods* 10 (4): 776–92.

Ceron, Andrea. 2017. "Intra-Party Politics in 140 Characters." *Party Politics* 23 (1): 7–17.

Chatterjee, Sujoy, Anirban Mukhopadhyay, and Malay Bhattacharyya. 2019. "A Review of Judgment Analysis Algorithms for Crowdsourced Opinions." *IEEE Transactions on Knowledge and Data Engineering* 32 (7): 1234–48.

Dai, Yaoyao, and Alexander Kustov. 2022. "When Do Politicians Use Populist Rhetoric? Populism as a Campaign Gamble." *Political Communication* 39 (3): 383–404.

da Silva, Frederico Ferreira, Andres Reiljan, Lorenzo Cicchi, Alexander H. Trechsel, and Diego Garzia. 2023. "Three Sides of the Same Coin? Comparing Party Positions in VAAs, Expert Surveys and Manifesto Data." *Journal of European Public Policy* 30 (1): 150–73.

Dawid, Alexander Philip, and Allan M Skene. 1979. "Maximum Likelihood Estimation of Observer Error-Rates Using the EM Algorithm." *Applied Statistics* 28 (1): 20–8.

de Raadt, Jasper, David Hollanders, and André Krouwel. 2004. "Varieties of Populism: An Analysis of the Programmatic Character of Six European Parties." Working Papers Political Science No. 2004/4. Vrije Universiteit Amsterdam, August.

de Vries, Catherine E., and Sara B. Hobolt. 2020. *Political Entrepreneurs: The Rise of Challenger Parties in Europe*. Princeton, NJ: Princeton University Press.

De Sio, Lorenzo, Andrea De Angelis, and Vincenzo Emanuele. 2018. "Issue Yield and Party Strategy in Multiparty Competition." *Comparative Political Studies* 51 (9): 1208–38.

Engler, Sarah. 2020. "Centrist Anti-Establishment Parties and Their Protest Voters: More than a Superficial Romance?" *European Political Science Review* 12 (3): 307–25.

Engler, Sarah, Bartek Pytlas, and Kevin Deegan-Krause. 2019. "Assessing the Diversity of Anti-Establishment and Populist Politics in Central and Eastern Europe." *West European Politics* 42 (6): 1310–36.

Ennser-Jedenastik, Laurenz, and Thomas M. Meyer. 2018. "The Impact of Party Cues on Manual Coding of Political Texts." *Political Science Research and Methods* 6 (3): 625–33.

Ernst, Nicole, Sven Engesser, Florin Büchel, Sina Blassnig, and Frank Esser. 2017. "Extreme Parties and Populism: An Analysis of Facebook and Twitter across Six Countries." *Information, Communication & Society* 20 (9): 1347–64.

Fazekas, Zoltán, Sebastian Adrian Popa, Hermann Schmitt, Pablo Barberá, and Yannis Theocharis. 2021. "Elite-Public Interaction on Twitter: EU Issue Expansion in the Campaign." *European Journal of Political Research* 60 (2): 376–96.

Fernandez-Vazquez, Pablo. 2014. "And Yet It Moves: The Effect of Election Platforms on Party Policy Images." *Comparative Political Studies* 47 (14): 1919–44.

Frenay, B., and M. Verleysen. 2014. "Classification in the Presence of Label Noise: A Survey." *IEEE Transactions on Neural Networks and Learning Systems* 25 (5): 845–69.

Gilardi, Fabrizio, Theresa Gessler, Maël Kubli, and Stefan Müller. 2022. "Social Media and Political Agenda Setting." *Political Communication* 39 (1): 39–60.

Göbel, Sascha, and Simon Munzert. 2022. "The Comparative Legislators Database." *British Journal of Political Science* 52 (3): 1398–408.

Green Party of Canada. 2019. "To fight climate change, you need to change the system." *Twitter*, March 15. https://twitter.com/CanadianGreens/status/1106644168632238080.

Grimmer, Justin, and Brandon M. Stewart. 2013. "Text as Data: The Promise and Pitfalls of Automatic Content Analysis Methods for Political Texts." *Political Analysis* 21 (3): 267–97.

Gründl, Johann. 2022. "Populist Ideas on Social Media: A Dictionary-Based Measurement of Populist Communication." *New Media & Society* 24 (6): 1481–99.

Hawkins, Kirk A. 2009. "Is Chávez Populist? Measuring Populist Discourse in Comparative Perspective." *Comparative Political Studies* 42 (8): 1040–67.

Hsueh, Pei-Yun, Prem Melville, and Vikas Sindhwani. 2009. "Data Quality from Crowdsourcing: A Study of Annotation Selection Criteria." In Eric Ringger, Robbie Hertel, and Katrin Tomanek, eds., *Proceedings of the NAACL HLT 2009 Workshop on Active Learning for Natural Language Processing*. Stroudsburg, PA: Association for Computational Linguistics, 27–35.

Hunger, Sophia, and Fred Paxton. 2022. "What's in a Buzzword? A Systematic Review of the State of Populism Research in Political Science." *Political Science Research and Methods* 10 (3): 617–33.

Ipeirotis, Panagiotis G., Foster Provost, Victor S. Sheng, and Jing Wang. 2014. "Repeated Labeling Using Multiple Noisy Labelers." *Data Mining and Knowledge Discovery* 28 (2): 402–41.

Jagers, Jan, and Stefaan Walgrave. 2007. "Populism as Political Communication Style: An Empirical Study of Political Parties' Discourse in Belgium." *European Journal of Political Research* 46 (3): 319–45.

Jung, Jae-Hee, and Margit Tavits. 2021. "Valence Attacks Harm the Electoral Performance of the Left but Not the Right." *The Journal of Politics* 83 (1): 277–90.

Kayser, Mark A., Matthias Orlowski, and Jochen Rehmert. 2023. "Coalition Inclusion Probabilities: A Party-Strategic Measure for Predicting Policy and Politics." *Political Science Research and Methods* 11 (2): 328–46.

Kollberg, Markus. 2024. "The Challenger Advantage – How Challenger Parties Disrupt Mainstream Party Dominance in the European Parliament." *Journal of European Public Policy*. https://doi.org/10.1080/13501763.2024.2391510.

Laver, Michael. 2005. "Policy and the Dynamics of Political Competition." *American Political Science Review* 99 (2): 263–81.

Laver, Michael. 2014. "Measuring Policy Positions in Political Space." *Annual Review of Political Science* 17 (1): 207–23.

Licht, Hauke. 2020. "Mitigating the Undue Influence of Domain-Irrelevant Content on Measurements of Elite Political Behavior When Using Social Media Data." Paper presented at Digital Democracy Workshop, Zurich, Switzerland, November.

Licht, Hauke. 2023. "Cross-Lingual Classification of Political Texts Using Multilingual Sentence Embeddings." *Political Analysis* 31 (3): 366–79.

Licht, Hauke, and Fabienne Lind. 2023. "Going Cross-Lingual: A Guide to Multilingual Text Analysis." *Computational Communication Research* 5 (2): 2.

Mair, Peter. 2013. *Ruling the Void: The Hollowing of Western Democracy*. London: Verso.

March, Luke. 2011. *Radical Left Parties in Europe*. London: Routledge.

Meijers, Maurits J., and Andrej Zaslove. 2021. "Measuring Populism in Political Parties: Appraisal of a New Approach." *Comparative Political Studies* 54 (2): 372–407.

Mudde, Cas. 2004. "The Populist Zeitgeist." *Government and Opposition* 39 (4): 542–63.

Mudde, Cas. 2007. *Populist Radical Right Parties in Europe*. Cambridge: Cambridge University Press.

Müller, Jan-Werner. 2017. *What Is Populism?* London: Penguin.

Oliver, J. Eric, and Wendy M Rahn. 2016. "Rise of the Trumpenvolk: Populism in the 2016 Election." *The Annals of the American Academy of Political and Social Science* 667 (1): 189–206.

Passonneau, Rebecca J, and Bob Carpenter. 2014. "The Benefits of a Model of Annotation." *Transactions of the Association for Computational Linguistics* 2: 311–26.

Pauwels, Teun. 2017. "Measuring Populism: A Review of Current Approaches." In Reinhard C. Heinisch, Christina Holtz-Bacha, and Oscar Mazzoleni, eds., *Political Populism: A Handbook*. Baden-Baden, Germany: Nomos, 123–136.

POLITICO Europe. n.d. "Poll of Polls." https://www.politico.eu/europe-poll-of-polls/ (accessed May 20, 2020).

Polk, Jonathan, Jan Rovny, Ryan Bakker, Erica Edwards, Liesbet Hooghe, Seth Jolly, Jelle Koedam, et al. 2017. "Explaining the Salience of Anti-Elitism and Reducing Political Corruption for Political Parties in Europe with the 2014 Chapel Hill Expert Survey Data." *Research & Politics* 4 (1): 1–9.

Proksch, Sven-Oliver, Will Lowe, Jens Wäckerle, and Stuart Soroka. 2019. "Multilingual Sentiment Analysis: A New Approach to Measuring Conflict in Legislative Speeches." *Legislative Studies Quarterly* 44 (1): 97–131.

Reinemann, Carsten, James Stanyer, Toril Aalberg, Frank Esser, and Claes H. de Vreese, eds. 2019. *Communicating Populism: Comparing Actor Perceptions, Media Coverage, and Effects on Citizens in Europe*. New York: Routledge.

Reungoat, Emmanuelle. 2010. "Anti-EU Parties and the People: An Analysis of Populism in French Euromanifestos." *Perspectives on European Politics and Society* 11 (3): 292–312.

Rooduijn, Matthijs, Sarah de Lange, and Wouter van der Brug. 2014. "A Populist Zeitgeist? Programmatic Contagion by Populist Parties in Western Europe." *Party Politics* 20 (4): 563–75.

Rooduijn, Matthijs, and Teun Pauwels. 2011. "Measuring Populism: Comparing Two Methods of Content Analysis." *West European Politics* 34 (6): 1272–83.

Rooduijn, Matthijs, Stijn van Kessel, Caterina Froio, Andrea Pirro, Sarah de Lange, Daphne Halikiopoulou, Paul Lewis, Cas Mudde, and Paul Taggart. 2019. *The PopuList: An Overview of Populist, Far Right, Far Left and Eurosceptic Parties in Europe*. https://popu-list.org/.

Sältzer, Marius. 2020. "Finding the Bird's Wings: Dimensions of Factional Conflict on Twitter." *Party Politics* 28 (1): 61–70.

Strøm, Kaare. 1990. "A Behavioral Theory of Competitive Political Parties." *American Journal of Political Science* 34 (2): 565–98.

Taggart, Paul. 2000. *Populism: Concepts in the Social Sciences*. Philadelphia: Open University Press.

van der Brug, Wouter, Meindert Fennema, and Jean Tillie. 2000. "Anti-Immigrant Parties in Europe: Ideological or Protest Vote?" *European Journal of Political Research* 37 (1): 77–102.

van Kessel, Stijn. 2015. *Populist Parties in Europe: Agents of Discontent?* London: Palgrave Macmillan.

van Vliet, Livia, Petter Törnberg, and Justus Uitermark. 2020. "The Twitter Parliamentarian Database: Analyzing Twitter Politics across 26 Countries." *PLOS ONE* 15 (9): e0237073.

Walter, Annemarie S., Wouter van der Brug, and Philip van Praag. 2014. "When the Stakes Are High: Party Competition and Negative Campaigning." *Comparative Political Studies* 47 (4): 550–73.

Widmann, Tobias, and Maximilian Wich. 2023. "Creating and Comparing Dictionary, Word Embedding, and Transformer-Based Models to Measure Discrete Emotions in German Political Text." *Political Analysis* 31 (4): 626–41.

Zulianello, Mattia. 2020. "Varieties of Populist Parties and Party Systems in Europe: From State-of-the-Art to the Application of a Novel Classification Scheme to 66 Parties in 33 Countries." *Government and Opposition* 55 (2): 327–47.

Zulianello, Mattia, Alessandro Albertini, and Diego Ceccobelli. 2018. "A Populist Zeitgeist? The Communication Strategies of Western and Latin American Political Leaders on Facebook." *The International Journal of Press/Politics* 23 (4): 439–57.