

How Right Wing is Right Wing Populism? Evidence from the Manifesto Corpus

ANDREW HALTERMAN, NICHOLAS HALTERMAN, JILL IRVINE

Right wing populist parties in Europe are clearly different from other right wing parties in their rhetoric and electoral appeal. Some observers see substantive differences between right wing populists and other right wing parties, with populists supporting the welfare state and gender equality more than other right wing parties, often as part of an anti-immigration and anti-Muslim agenda. We test this claim using novel data produced by a multilingual convolutional neural net on political party platforms for the years 1990 to 2015 from the Manifesto Corpus. We find no systematic differences between right wing populists and non-populists on support for welfare and gender equality, though there is some evidence that more successful populists are more centrist.

Introduction

Right wing populist parties have become a central part of the political landscape in Europe over the past ten years.¹ Although the success of these parties has varied over time and place, they are likely to be a central feature of European political life for years to come. While it is clear that these parties embrace a variety of positions related to preserving national identity and strengthening borders that few would hesitate to describe as on the right side of the left-right ideological spectrum; it is less clear that all their positions could be described as right wing. FPÖ election

¹We thank Julie Halterman and Ron Halterman for assistance in annotating quasi-sentences in German and French. Andrew acknowledges the financial support of the MIT Department of Political Science and an NSF Graduate Research Fellowship. All replication materials are on a private Github repository and available upon request.

posters proudly proclaim, “We are in favor of women’s right to self determination,” while AfD leaders call for increased penalties against sexual assault. Meanwhile, UKIP appeals to the “left-behind” voters, condemning both left-wing and right wing government parties alike as belonging to the “exam-passing classes” who are out of touch with the needs of the economically disadvantaged (Bogdanor 2016). A main rallying cry of the National Front in the most recent elections was “Neither left nor right,” in order to stake out their new ideological territory, and other right wing populist parties have taken up this stance. Just how right wing then, are right wing populist parties?

Much discussion and debate on right wing populism revolves around its definition and sources of support, but less comprehensive study to date considers the variation among European right wing populist parties and between populist and non-populist right wing parties on issues traditionally associated with the ideological left. We argue that support for welfare state policies and gender equality provide two useful tests of right wing populist parties’ ideological location. Although case studies have revealed the situational dynamics of right wing populist support for gender equality, often as part of anti-immigration, anti-Muslim platforms, the extent to which these positions are being formally adopted across European right wing parties is less clear (Mayer, Ajanovic, and Sauer 2014; Kostadinova and Mikulska 2015; Rashkova and Zankina 2017; Pilkington 2017). Similarly, while scholars point to a greater support for egalitarianism and protection of the common man from the economic forces of globalization among “new right” parties, others argue that this rhetoric does not translate into support for concrete policies (Mudde 2007; De Koster, Achterberg, and Van der Waal 2013; Afonso and Papadopoulos 2015). Building upon this emerging scholarship, our study investigates the extent of shifts on support for gender equality and social welfare policies and how this challenges the orientation of right wing populist parties along the left-right spectrum. By empirically investigating variation along these two lines, we examine the notion of “thin ideology,” and ways of categorizing right wing populist groups and parties.

We select support for the welfare state and for gender equality as our two measures of right wing populists’ ideological position. First, support for the welfare state is one of the canonical left-right distinguishing features between left and right parties. Right wing populists’ support for greater welfare spending would be good evidence that they were less on the right and more supportive of positions historically associated with the left than their label suggests. Gender equality provides a second important measure of left-right ideology, this time on a social dimension as opposed to an economic dimension. Together, gender equality and the welfare state provide two distinct dimensions along which to measure the left-right position of parties. Finally, both are specific platforms that observers have

pointed to as differentiating right wing populists from other right wing parties, a difference which is largely attributed to tactical anti-Muslim and anti-immigration stances.

We begin by exploring the debates about how best to understand right wing populism, arriving at the definition we will use in our own investigation. Drawing on the theoretical literature on populism, we propose a definition of right wing populist parties based on their stance on four issues: globalization, national identity, immigration, and the European Union. Using the Manifesto Project Database (Lehmann et al. 2017), we evaluate which political parties match the populist issue-profile, identifying the set of right wing populist parties in Europe. We then investigate the extent to which right wing populist parties actually adhere to traditionally right wing stances on issues other than these core populist issues by examining their positions on state social welfare services and on women's rights and gender equality.

Because statements in favor of women's rights and gender equality are not coded in the Manifesto Corpus, we hand code around 1,600 party statements from the corpus in English, German, French, and Bosnian-Croatian-Serbian for pro-gender equality and women's rights statements. We then use these labeled sentences to train a convolutional neural net classifier using aligned word embeddings to create a language-agnostic sentence classifier (see Appendix A). We apply this classifier over all documents from those languages in the corpus. Through these processes, we are able to compare party platforms and manifestos across a wide number of cases, using the Manifesto Project data to explore questions of variations in gender and social welfare policies and to examine the extent to which European populist parties can be described as right wing parties. Although we are not able to generate accurate scores for new languages beyond those for which we created training data, we are able to develop very accurate classifiers for gender equality statements with only 100-200 labeled quasi-sentences. This methodological result should be of interest to other scholars who would like their own customized codings of the Manifesto Corpus or other political speech.

The results of our empirical investigation lead to the following substantive conclusions. First, we find that right wing populist parties are not notably supportive of gender equality or the welfare state, either in absolute terms or relative to other right wing parties. In the case of gender equality, these results hold across time and party vote percentage. In the case of welfare, there is evidence that right wing populist parties are becoming increasingly pro-welfare, and that higher vote-share parties are more supportive of the welfare state, but neither of these findings is enough to overcome the negative relationship between right wing populism and support for the welfare state. On the whole, right wing populist parties do not

appear to deviate from traditional right-left lines on either issue. This raises questions about why right wing populist parties are perceived, by the public and some scholars alike, as less right wing than traditional right wing parties, and we conclude by offering some possible answers.

Literature Review and Research Questions

We begin by providing an overview of the debate about concepts of populism as a way of motivating and justifying the measure we propose and use in this paper. We then turn to a discussion of existing literature on populism and our policies of interest: support for the welfare state and for gender equality. These theoretical understandings are then made more concrete in the empirical section that follows.

Defining Right Wing Populism

In order to investigate the variation in positions on social welfare policies and gender equality among right wing populist parties in Europe, we need to establish what we mean by right wing populist, and a method of identifying right wing populist parties in Europe. “Populism” is a notoriously slippery term. It is ill-defined and overused, leading to a lack of consensus on its meaning and a temptation to abandon the term altogether, or to treat it to the same proliferation of adjective pairings that ‘democracy’ has experienced (Sartori 1970). While some scholars insist that it can only be understood as a political method encompassing particular modes of discourse and political practice (Jansen 2011; Aslanidis 2016), others argue that it can be understood as an ideology in its own right, although one that can take many forms and endorse a variety of positions (Mudde 2010; Pelinka 2013).

The discussion of populism as an ideology revolves around the notion of “thin ideology,” the proposition that populism might best be understood as a supra or meta ideology. The thin ideology of populism, based on its strong anti-establishment views and belief that the people should rule, is then layered over “deep ideologies” such as socialism, conservatism, or nationalism (Stanley 2008). The task for those studying populism, then, is to identify its core components and decide whether they fit what we see as the minimum definition of ideology: a simple set of ideas that provides a description of the present, a vision of the future, and a blueprint for getting there.

In contrast to this line of discussion, those who understand populism as a political

method argue that its what might be described as chameleon-like approach to political ideas and positions requires that it be understood *prima facie* as a mode of discourse or a political practice (Aslanidis 2016). Populist ideas simply aren't fixed enough to qualify as a coherent set of values, ideas, and positions. In this view, populist method rests first and foremost on a type of discourse that praises the common man at the same time that it excoriates corrupt elites. Such discourse tends to be highly idiomatic, evocative, and culturally specific, whether it is evoking the "silent majority," "Main Street (vs Wall Street)," or "left-behind voters," and its purpose is to mobilize groups who are discontented with the current political practices and leaders (Albertazzi and McDonnell 2007). While recognizing the importance of discourse, some scholars have found it more useful to think of the populist method as a form of political practice and strategy for gaining power. Jansen, for example, argues that populism is a form of political mobilization, combining the mobilization of marginalized groups—what he calls popular mobilization—with populist discourse, "anti-elite, nationalist, rhetoric that valorizes ordinary people" (2011, 82).

Definitions of what constitutes the ideological right or political right wing have been just as slippery to settle on and just as contested as definitions of populism. Roger Eatwell aptly describes the continually shifting meaning of right wing views as they move among reactionary right, conservative right, radical right, religious right, and populist right (2013). Other scholars have enumerated the views and positions associated with the political right wing; Cas Mudde (2010), for example, has detailed 58 ideological characteristics associated with right wing parties. Still other scholars have adopted a simpler, bifurcated understanding of left and right; in this scheme, the old left (socialism) is distinguished from the new left (post-materialism), while the old right (authoritarian, social hierarchy) is contrasted with the new right (law and order, egalitarian) (De Koster, Achterberg, and Van der Waal 2013). In short, capturing the dimensions of the left/right ideological spectrum over time and place has proven notoriously elusive.

Despite the continuing efforts to arrive at a consensus about what constitutes populism and right wing, a fairly high degree of agreement exists on certain characteristics of right wing populism in contemporary European politics. Whether populism is understood as a political method or political ideology, agreement exists on its two core propositions: that it is aimed against corrupt political elites and that it is committed to defending "the people." How the people is defined is what often distinguishes the left or right prefix that accompanies populism. As Pelinka points out, "[t]he principally radical consequences of populist democracy are based on an extremely ambiguous precondition—the self-evidence of 'the people'" (2013, 3). Where the people is defined by class, left wing populism prevails; where it is defined by nation or ethnicity, right wing populism prevails.

And, indeed, it is precisely the mission of defending the people, understood as the French, Swedish, or Danish nation, that most agree is the ideological driver of right wing populism in Europe today. Defense of the nation is undertaken against the combined threats of globalization, Europeanization, immigration, multiculturalism, and Islamization, all of which are supported and implemented by a group of self interested European elites. In this view, these elites—represented by the establishment and the mainstream political parties—have worked as a kind of 5th column toward the annihilation of their own people.

This consensus about the core propositions of populism informs the definition of right wing populism we use in this study. In the most general sense, right wing populism might be defined as a struggle against corrupt political elites in defense of the nation. According to Cas Mudde (2007), right wing populism is identified in more concrete terms through the fairly stable set of views that describe the struggle for national survival: namely, opposition to the European Union, to multiculturalism and internationalism, and, perhaps most importantly, to immigration. As Mény and Surel (2002) caution, by identifying populism with only with specific issues or positions we may miss out on its analytical purchase. Nevertheless, we argue that the way in which these positions are articulated—the mode of elaborating ideas and defining situations—contains clear elements of populism. According to right wing populist views, it is precisely because the current political elites have embraced multiculturalism, Europeanization, and globalization, while opening the doors to immigration, that these forces threaten to overwhelm the national identity, cultural integrity and welfare of the native citizens of these countries. We do not use the term radical right here, because although right wing populist parties attempt to mobilize the populace on behalf of their political project, they are not revolutionary in their aims (Golder 2016). Rather, they are most concerned with restoring the status quo ante—before mass migration, Europeanization and globalization (Pelinka 2013).

Right Wing Populism and Social Welfare

While right wing populist parties can be described as right wing because of the strong element of ethno-nationalism that infuses their values and views, it is less clear where these parties fall on the left/right spectrum when it comes to other issues. Studies of right wing populist parties suggest that at least some right wing populist parties seek to distance themselves from traditional right wing parties, particularly when it comes to issues relating to social welfare policies and gender roles. Earlier work on the economic positions of far right parties put forth two views: either these parties' economic policies were seen as underdeveloped and

rhetorically de-emphasized or, in the more common view, they were seen to differ little from other right wing parties in their support for neo-liberal approaches to the economy and the state (Betz 1994; Cheles, Ferguson, and Vaughan 1995). According to this view, far right parties garner support when they adopt the “winning strategy” of support for dismantling the overgrown state bureaucracy underpinning social welfare and replacing it with a lean and mean (law and order oriented) one (McGann and Kitschelt 2005).

More recent developments have caused some scholars to reevaluate this view. They argue that the economic downturn of 2008 and increasing opposition to immigration has caused right wing populist parties to shift their views. These parties now endorse social welfare benefits for the native population, while making it more difficult for these benefits to be given to immigrant outsiders (Mudde 2007; Betz 2013; De Koster, Achterberg, and Van der Waal 2013; Golder 2016). A common theme in right wing populist party’ campaigns is that vulnerable native citizens, especially the elderly, are being denied the social welfare benefits they need while newly arrived immigrants receive generous stipends from the government. Moreover, ordinary citizens are being hurt by the forces of globalization and Europeanization directed by corrupt political elites bent on self enrichment. As Marine Le Pen repeated frequently on the campaign trail, “there is no more left or right, only globalists and nationalists.”

Thus, the winning formula may no longer include support for the neo-liberal policies of the traditional right, but a shift toward more left wing policies of egalitarianism and support for the welfare state (Kitschelt 2007). In short, since “the economic interests of the potential electorate of these parties are not in line with traditional right wing stances on economic redistribution and the welfare state... the economic agenda of new-rightist populist parties is often markedly different from that of traditional right wing parties” (De Koster, Achterberg, and Van der Waal 2013, 4). In the view of some scholars, current right wing populist views can be characterized by two dimensions, what is labeled welfare chauvinism—that social welfare benefits should be reserved for the native population—and welfare populism—that the current social welfare policies have been corrupted by political elites and do not benefit the common man (De Koster, Achterberg, and Van der Waal 2013). Rovny, however, challenges this view by arguing that right wing populist parties may intentionally engage in position blurring on economic issues, calling into question the utility of assessing right wing populist parties’ economic positions (Rovny 2013).

In sum, there are two competing positions in the literature about right wing populist support for social welfare that we shed light on. One, which might be called the conventional view, holds that right wing populists are very similar

to other right wing parties, merely adopting populism as a rhetorical strategy. This view predicts that populists will show a similar lack of support for welfare programs as other right wing parties do. A contrasting approach in the literature, one which has gained popularity recently, points to right wing populists as a special breed, willing to cross traditional party lines to strengthen their “common people” platforms. This view predicts that right wing populists may adopt social welfare as an issue, and may deviate significantly from non-populist right wing parties in their support for social welfare policies. In order to investigate these competing views, we ask three specific questions: First, how supportive of the welfare state are right wing populists? Second, is that level of support changing over time? And, finally, how do right wing populist parties compare to other right wing parties on welfare issues?

Right Wing Populism and Gender

Numerous studies have suggested that in addition to the apparent increase in their support of social welfare benefits, right wing populist parties can also be distinguished from traditional right wing parties and earlier radical right parties by their increasing support for women’s rights and gender equality (Betz and Meret 2009; Betz 2013; Pelinka 2013; Akkerman 2015). Perhaps the most striking example of this growing support for women’s rights was seen in the AfD’s vocal advocacy of increasing legal penalties against sexual assault in the aftermath of the sexual assaults on new year’s eve in 2015 in Cologne and other German cities. Rigorous cross national examination of right wing populist and far right parties’ positions on gender are scant, however (Mudde 2007).

Research analyzing the role of women in far right parties in Europe in the 1990s and 2000s, generally focused on women’s participation in and support for these parties. Studies described this role as one of subordination and exclusion, as women were relegated to caretaking tasks related to their domestic functions. Moreover, women were described as much less likely to support the views of right wing parties resulting in a gender voting gap for these parties. When women did join far right and right wing populist parties, they did so because of the participation of male relatives rather than through their own alignment with the parties’ positions (Blee and Deutsch 2012; Bacchetta and Power 2013). Nevertheless, a few scholars underscored that ethno-nationalist and right wing movements had successfully exploited support for gender equality as a way of distinguishing themselves from perceived Muslim enemies. For example, a study of the extreme right Serbian Radical Party revealed that while the party insisted on a return to a “traditional and martial” Serbian masculinity that had been weakened by

decades of communism, it generally supported women's "modern and equal" role in public life as a counterpoint to what it claimed were the backward, oppressive gender attitudes of Serbia's Muslim (and irredentist) minority population (Irvine and Lilly 2007).

More recent research has explored this instrumental support for women's equal rights, arguing that it has become an increasingly prominent feature of right wing populist parties in Europe, particularly in response to terrorist attacks in the US and Europe and the increasing level of immigration into Europe from Africa and the Middle East (Betz and Meret 2009; Akkerman and Hagelund 2007; Spierings et al. 2015; Mayer, Ajanovic, and Sauer 2014). As these right wing populist parties have embraced identitarian concerns related to European culture, values, and civilization, their advocacy of women's equal rights has become a central theme in their claim to be defending these values (Betz and Meret 2009). By raising concerns about the treatment of women in immigrant communities, right wing populist parties have been able to shift support away from multiculturalism in countries like Norway and the Netherlands, toward a greater emphasis on unicultural and assimilationist approaches to these communities[a]. New alliances between feminists and right wing populists emerged over purported high levels of genital mutilation, domestic violence, and forced marriages in immigrant communities, with some claiming that these cultural and religious practices were evidence of a fundamental "clash of civilizations[b]." As Akkerman and Hagelund (2007) put it, this discourse on gender "helped to legitimate the shift away from multiculturalism among left parties, while it has provided a potentially respectable anti-immigration position for radical-right parties" (213). As a result, the boundaries between left and right parties began to shift and blur.

Given that support for women's rights has become a discursive strategy of at least some right wing populist parties, to what extent has support for women's rights been adopted as an official party position of right wing populist parties in Europe? How have these parties' positions on women's rights and gender roles changed over time? Is there significant variation among right wing populist parties when it comes to views on gender? How do right wing populist parties compare to mainstream conservative parties?

Recent work by Tikke Akkerman (2015) suggests some preliminary answers to these questions. Akkerman examines the positions of six right wing populist parties in order to discover the extent to which this party family has discarded its conservative stance and adopted more liberal views of gender relations: the Party for Freedom (PVV); the National Front (FN); the Freedom Party of Austria (FPO); Swiss People's Party (SVP); the Danish People's Party (DF); and Vlams Belang. Drawing upon party manifestos, she develops a gender score based on

party positions in two policy domains: family policy and immigration/integration policy. If party manifestos express support for women’s participation in the labor market, political participation and equal right, public childcare, family planning and same sex partnerships they are liberal—if they express support for wages for motherhood, private child care, large families and their privileged position, and opposition to reproductive rights, they are conservative (Akkerman 2015, 41). The immigration/integration policy domain is based on party levels of support for the rights of immigrant women in general (including wearing hijab) their protection from violence, mutilation or trafficking, and their religious, social and family reunification rights. She also considers support for restrictions of gay immigrants rights.

Akkerman concludes that, contrary to what we might expect, right wing populist parties are more conservative than traditional right wing parties on gender and they have remained so over time. “This indicates that support for gender equality and women’s rights has now become widely spread over the whole political spectrum in these countries, with only the radical-right parties left to defend the last vestiges of (modern) conservative family relations” (Akkerman 2015, 52). Nevertheless, significant variation among right wing populist parties exists. While some parties hold traditional conservative views on women—that their place is in the home; others hold modern conservative views on women—that their primary role is as mothers, but they also have an important part to play in the labor force. In four of the parties, excluding the FPO and FN, gender has become less, not more, salient over time.

We propose to extend this and other studies in the following ways. First, while we include the FPO, the SVP and the FN in our analysis, we also include UKIP and the AfD, two right wing populist parties that have received a great deal of attention in the past three years, as well as others including the Alliance for the Future of Austria (BZÖ), the Democratic Unionist Party (DUP), the Croatian Party of Rights (HSP), the Socialist Party of Serbia (SPS), and the Serbian Radical Party (SRS). Ultimately, we aim to include all European right wing populist parties in our empirical analysis through the neural net classifier model described in Appendix 1. Second, we have devised a measure of support for women’s rights and gender equality that avoids the complexity of family policies such as motherhood wages or child care, whose meaning in terms of gender equality is often context specific. Instead, we propose a measure that includes general statements of support for gender equality and women’s rights, women’s participation in political life and the labor market. Moreover, we measure levels of support for combating violence against women, since this has been a central issue in right wing populist rhetoric in defense of liberal, European values against Muslim immigrants. Third, we exclude support for same sex marriage or gay rights in our gender equality

measure, since right wing populist parties may hold divergent views on women’s rights and LGBTQ rights. Rather, we propose to tackle this issue separately in future work. Finally, our method of machine coding allows us to develop gender equality scores for all parties, not just for right wing parties. We are thus able to make observations about the level of consensus on women’s rights in a particular country, measure the generalized relationship between right wing populism and gender equality across different countries and make comparative evaluations between right wing populist parties and other parties.

Empirical Approach

We use the Manifesto Corpus to create measures of our four variables of interest: parties’ degree of right wing populism, position on a left-right scale, support for the welfare state, and rhetorical support for gender equality. Two measures, left-right ideology and support for welfare, have already been extensively studied with the project’s data. We use code `per504` to measure parties’ support for the welfare state, which shows a share of quasi-sentences in the standard approach taken by the project.² Rather than examine these raw proportions, we calculate a country-specific vote-weighted z-score for each party. In order to do so, we take the mean `per504` score for each country weighted by the vote share received by each party, giving large successful parties more impact on the average country position than smaller fringe ones. We then use this weighted mean and a similarly weighted standard deviation to calculate z-scores for each party, showing each party’s level of deviation from the centrist position on welfare in their country. By performing calculations by country, we prevent relatively-anti-welfare parties in unusually pro-welfare political contexts from appearing as outliers against a Europe-wide average.

Placing parties on a left-right scale on the basis of their manifestos is a bread and butter undertaking for scholars of manifestos and party platforms. We use the method described by Lowe et al. (2011) to generate left-right scores for each manifesto. We chose this method primarily because of its ease of implementation in the ManifestoR package (Lewandowski et al. 2015), its ability to generate scores through the most recent elections, and because we could easily modify the ideology

²Manifesto Project Welfare Wording: “Favorable mentions of the need to introduce, maintain, or expand any public social service or social security scheme. This includes, for example, government funding of: Health Care Child Care Elder care and pensions Social housing Note: This category excludes education”

scores to remove support for the welfare state from the calculation. Because we are interested in the differences between populists and non-populists on support for the welfare state, controlling for left-right position, it was inappropriate to include welfare state support in the ideology score.

Our third and fourth measures required more novel work. To measure right wing populism, we constructed a composite score of four existing manifesto codings based on our definition of right wing populism (described above). Those four dimensions are: European Community - Negative (**per110**), National Way of Life - Positive (**per601**), Immigration - Negative (**per601-2**), and Multiculturalism - Negative (**per608**), normalized as the proportion of quasi-sentences in the manifesto. As with welfare, we translate these raw proportions into z-scores for each dimension; in order to prevent more frequently discussed issues from dominating the scores. These z-scores are calculated from the entire set of parties, rather than by country, in order to allow us to examine right wing populism as an independent phenomenon—we are interested in the real degree to which a party is right wing populist, not whether its country sports an unusual number of right wing populists. We do weight the mean and standard deviation by vote share, for the reasons mentioned with welfare coding.

The final new measure we construct is a measure of parties' support of gender equality, as measured by their written manifestos. Existing literature and anecdotal observations indicate that right wing populist parties may be more supportive of gender equality than other right wing parties, largely out of anti-Muslim tactical considerations. The existing format of the Manifesto Project corpus does not code for support of gender equality as distinct from the other forms of social equality in **per503**; studying the relationship between gender equality support and right wing populism therefore requires a new measure of support for gender equality in parties' manifestos. Specifically, we code the following kinds of statements as pro-gender equality:

1. General statements affirming gender equality and women's equal rights
2. Positive reference to increasing women's participation in political life and increasing women's political representation
3. Positive statements about increasing women's participation in the labor force and economic decision making
4. Statements condemning obstacles to women's equality, including violence against women, lack of access to child care, etc.
5. Positive statements about measures to combat violence against women
6. Statements noting the disadvantages faced by women, including gaps in pay, representation on corporate boards, etc.

Two major obstacles stand in the way of creating a new coding of quasi-sentences:

first, the manifestos are written in many languages, each of which requires separate processing or coding, and second, even within a language, a process that requires reading each quasi-sentence would be prohibitively time consuming for a single research project. We address these two problems (the latter more effectively than the former) by developing a statistical classifier that can generate labels for an entire language's quasi-sentences after being trained on only a small subset (see Appendix A). All of our measures are affected to some degree by missing data in the manifesto corpus. Some manifesto text is missing altogether, while other text is present but not available in quasi-sentence form. After producing scores for each quasi-sentence, we create a weighted z-score for each manifesto with the same method used for welfare scores.

Substantive Findings

We now turn to a discussion of what we have been able to learn from the empirical strategy we describe above. First, we confirm that our four-part measurement of right wing populism matches the consensus understanding of which parties in Europe are populist. We then turn to our findings on the stance of right wing populists toward the welfare state and gender equality.

Identifying Right Wing Populist Parties

Before examining welfare and gender, we first checked to see how well our measure of right wing populism functioned. By taking the sum of manifesto's z-scores in each issue (immigration, EU community, multiculturalism, and national way of life) we produced a continuous "right wing populism" score for each manifesto. The units of this score are number of vote-weighted standard deviations from the vote-weighted mean. The rankings produced by this method match well common expert classification of parties (Table 1). Notable right wing populist parties all feature prominently, including the FPO, AfD, UKIP, SVP, and to a lesser extent, the FN. Much of our analysis focuses on this continuous measure of right wing populism, which offers more information for analyzing the general relationship between right wing populism and policy issues than a binary indicator would. In instances in which a party classification is needed, we set the cutoff at 3 standard deviations above the mean for right wing populist parties. This score represents a notable difference from the average party, captures all of the major right wing populist parties except for FN's 2002 document, and excludes recognized center-right parties such as the CDU.

table 1 The top 20 most right wing populist parties

	Party	Country	Election Date	Populism Score
1	Swiss People's Party	Switzerland	23/10/2011	19.91
2	Danish People's Party	Denmark	20/11/2001	18.30
3	United Kingdom Independence Party	United Kingdom	07/06/2001	17.62
4	Party of Freedom	Netherlands	12/09/2012	16.40
5	Danish People's Party	Denmark	13/11/2007	15.04
6	Independent Greeks	Greece	25/01/2015	14.68
7	Alternative for Germany	Germany	24/09/2017	14.40
8	Austrian Freedom Party	Austria	29/09/2013	13.96
9	Danish People's Party	Denmark	15/09/2011	13.00
10	Party of Freedom	Netherlands	22/11/2006	12.80
11	National Union Attack	Bulgaria	25/06/2005	11.97
12	Croatian Party of Rights	Croatia	23/11/2003	11.07
13	Party of Freedom	Netherlands	09/06/2010	10.64
14	League of Polish Families	Poland	23/09/2001	10.26
15	United Kingdom Independence Party	United Kingdom	07/05/2015	9.36
16	National Alliance	Latvia	17/09/2011	8.82
17	Communist Party of Greece	Greece	17/06/2012	8.50
18	True Finns	Finland	18/03/2007	8.38
19	United Patriots	Bulgaria	26/03/2017	8.24
20	Sweden Democrats	Sweden	19/09/2010	8.14

Social Welfare Results

To unpack the relationship between right wing populism and support for the welfare state, we began by answering three main questions.

1. What is the relationship between right wing populism and support for welfare?
2. Is that relationship changing over time?
3. How do right wing populists compare to other right wing parties on welfare issues?

To preview our findings, we conclude that right wing populism is negatively correlated with support for the welfare state. We also find some evidence that average support for welfare programs is increasing among right wing populists, but not enough to change the negative relationship between right wing populism and welfare support. We find no evidence that right wing populist parties are more pro-welfare than their non-populist counterparts. In fact, we find populism to be a negative predictor of welfare scores even when controlling for left/right party position. Answering these questions raises a fourth: if right wing-populists are not systematically pro-welfare, either in isolation or relative to other right wing

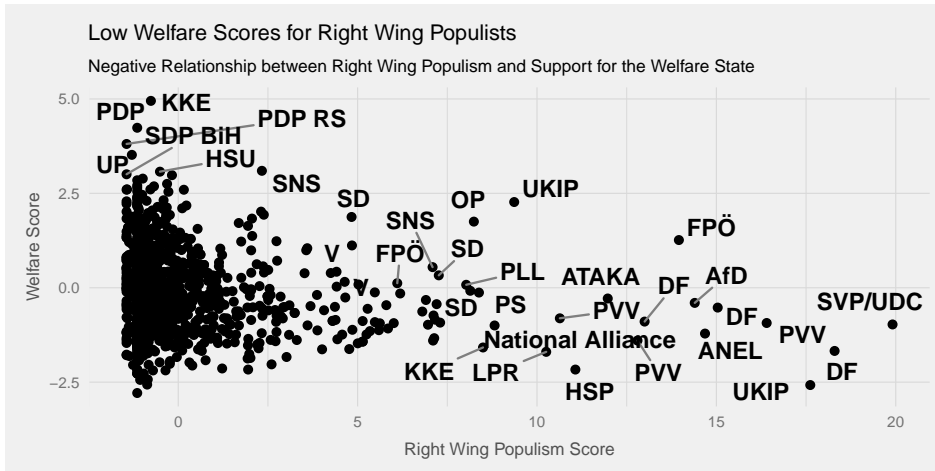


Figure 1. Populism vs. support for the welfare state

parties, why have they often been characterized as such, both in media accounts and in the contemporary literature? We conclude by arguing that an availability heuristic is at work, distorting observers' perceptions of right wing populists' stances on the welfare state.

Our first question is the relationship between right wing populism and support for the welfare state. In isolation, how supportive of the welfare state are right wing populists? We find a negative correlation between right wing populism scores and welfare scores Figure 1. In a linear model, each standard deviation from the mean for right wing populism scores decreases a party's welfare score by .07 standard deviations (Table ??). The effect is not very large, but certainly denies that more populist parties are on the whole more supportive of welfare than less right wing populist ones. To some extent, this result is not surprising. Right wing populists are right wing, even if we do not use their stance on the welfare state to calculate their ideology score.

On the second question, even if right wing populist parties are not generally pro-welfare, it could be the case, as some scholars have argued, that their support for social welfare policies is increasing over time (Mudde 2007; Betz 2013; De Koster, Achterberg, and Van der Waal 2013; Golder 2016). In fact, we do find evidence that the average welfare scores are increasing for right wing populist parties, both unweighted and weighted for party vote share Figure 1. The coefficient of the effect is not very large, but it is statistically significant in a linear model. In a linear regression, each year since 2000 increases the weighted average welfare support

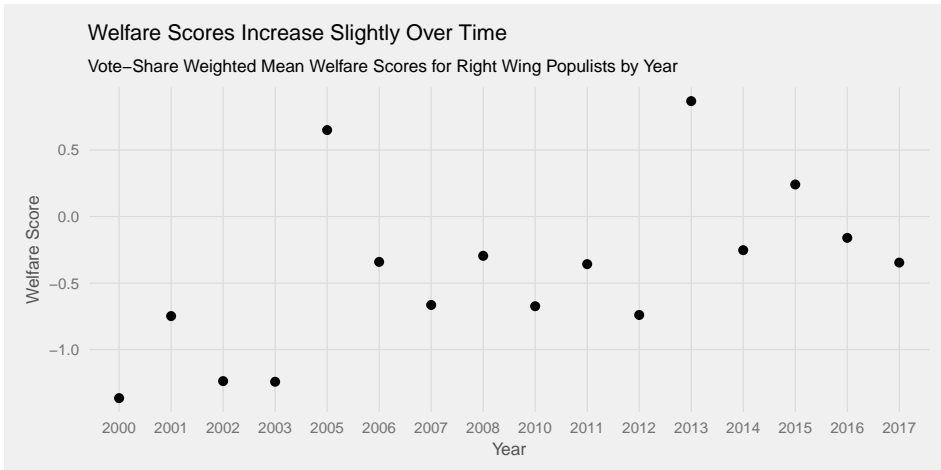


Figure 2. Populist welfare scores over time

score for right wing populists by 0.06 standard deviations from the country mean (Table ??). Nevertheless, even in recent years right wing populists have not been particularly pro-welfare, although if the existing trend holds this may change over time.

Finally, we test the claim that right wing populists are more pro welfare relative to other right wing parties. According to this position, even though right wing populists may not appear particularly pro welfare in absolute terms, they may be significantly more pro welfare than their right-left positioning would suggest. Initially, if we use a simple populist/non populist split for right wing parties, the populists are consistently less pro-welfare (Figure 3). In line with the findings from question 2, however, three of the five years in which populists outscore non-populist right wing parties occurred since 2010, pointing towards a potential change.

Moving beyond a simple binary split, we can test the relationship between right wing populism measured continuously and support for the welfare state, while controlling for left/right ideological position. If right wing populist parties are unusually pro-welfare relative to how right wing they are on other issues, we would expect right wing populism to positively correlate with support for the welfare state when controlling for left/right ideology. In fact, in a multivariate linear model, even when controlling for left/right ideology, right wing populism negatively predicts support for the welfare state (Table ??). If anything, it seems that right wing populists are less pro-welfare than their non-populist counterparts;

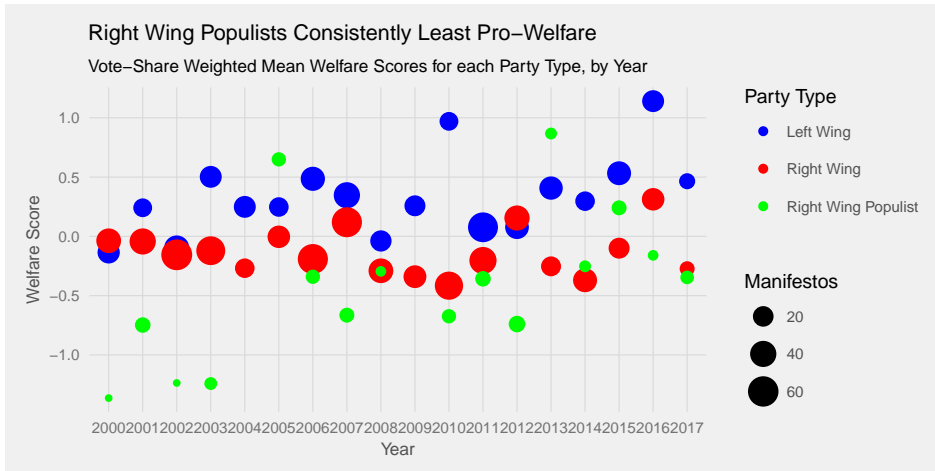


Figure 3. Welfare scores by party type

there is certainly no evidence to suggest that they are systematically more pro-welfare.

Taken as a whole, our results do not support the characterization of populist parties as a new type of pro-welfare right wing party, though there is some evidence that right wing populist parties may be heading in that direction.

Faced with these results, a new question arises: why might people think that populist parties are pro-welfare if they are not? We contend that this impression is essentially the result of an availability heuristic (Tversky and Kahneman 1973) fueled by two factors: that larger right wing populist parties tend to be more pro-welfare, and that several high-profile parties have had unusually high pro-welfare scores and higher scores than center-right parties.

On the first effect, there is some evidence that higher vote-share right wing populist parties are more supportive of welfare state (Figure 4). In a linear model, right wing populist parties' welfare z-scores increase by around 0.03 for each additional percentage share of the vote they receive (Table ??). This relationship may help explain the mischaracterization of right wing populist parties' positions: if higher-profile high vote-share parties tend to be more populist, they may obscure the real average. The causal nature of this relationship is unclear. Populist parties that adopt more pro-welfare stances may enjoy electoral benefits. Alternatively, populist parties that enter the mainstream by receiving a higher vote share may shift their positions after the fact to be more centrist. In either event, the rela-

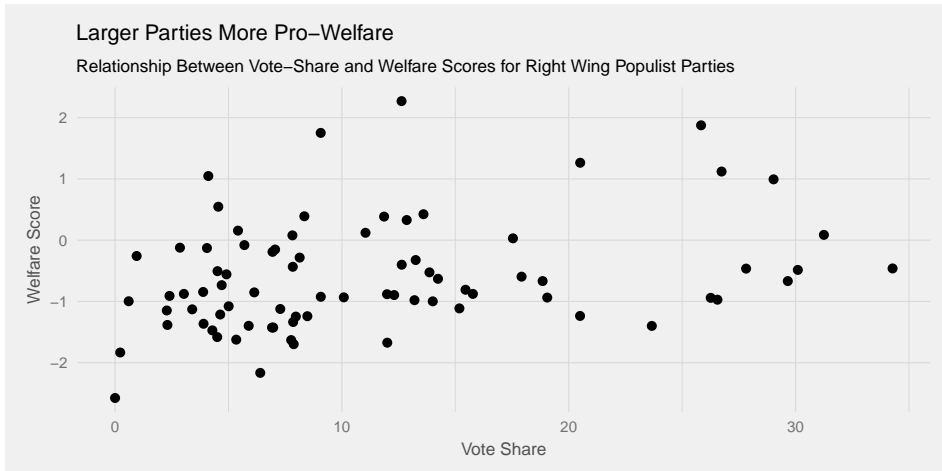


Figure 4. Vote-Share vs Support for Welfare State

tionship between right wing populist success and degree of welfare-support merits further investigation.

A second reason right wing populist parties might appear more pro-welfare as a group is that several high-profile parties have had unusually high pro-welfare scores and higher scores than center-right parties. In fact, a small number of highly visible outliers have been heavily referenced in accounts of pro-welfare right wing populists. Even beyond the general trend of more successful right wing populist parties being more pro-welfare, these parties warp common perception both through unusually high scores and through unexpectedly higher scores than their center-right or left-wing rivals.

The most striking, though perhaps not the most discussed, of these outliers is UKIP in 2015. With a country-specific z-score of 2.27, they were the second most pro-welfare party in the election, only behind the Social Democratic and Labour Party which received 0.3% of the vote. UKIP was nearly a full standard deviation higher than the Labour party on welfare support, and more than 2 above the Conservatives. The FPÖ in 2013 are a classic outlier, scoring as the second most pro-welfare party in the election with a score of 1.26 standard deviations above the country mean. Notably, they scored over a standard deviation higher than the center-right Austrian People's Party. Perhaps most discussed, the FN in 2017 scored higher than Macron's Republic Onwards! coalition, a discrepancy that seems to show a reversal in traditional party divides.

Taken together, these parties might seem to present a compelling case of pro-welfare right wing populists across multiple countries. Such a conclusion is misleading. UKIP's 2015 manifesto and the FPÖ's 2013 manifesto are massive outliers among the rest of their manifestos. FN's higher score than Macron's coalition is surprising, but the gap was less than 0.15 standard deviations. Furthermore, FN's real score was below the country average, and lower than 5 other parties in the election. The AfD in Germany has been one of the least pro-welfare parties in both of its elections. These isolated examples are striking and theoretically interesting, but belie the overall status of right wing populist parties as not particularly pro-welfare.

Gender Equality results

We aim to answer the same basic questions for right wing populist parties' positions on gender equality as we did for welfare. Namely: what is the relationship between right wing populism and support for gender equality; how is it changing over time; and what is the relationship relative to other right wing parties? Due to the unreliability of cross-language modeling, we perform the gender equality analysis only on manifestos with hand-coded language support. This limits our countries of analysis to Austria, Croatia, France, Germany, Ireland, Serbia, Switzerland, and the United Kingdom. With this more limited sample, our results on gender are less reliable than for welfare, particularly for longitudinal work. On the whole, however, we find that right wing populists are not notably pro-gender equality in any of the three aspects.

As with welfare, right wing populists are not especially pro-gender equality in absolute terms. As a group, right wing populists do not score highly on support for gender equality compared to most other parties (Figure 5). In a linear model, right wing populism is a significant negative predictor of pro-gender equality scores, with each standard deviation from the mean reducing gender scores by 0.09 standard deviations. This effect is larger than with welfare, and predicts a firmly right wing populist party like the FPÖ to be more than half a standard deviation below the mean gender score.

In contrast to our findings concerning right wing populist parties' support for social welfare, there is no real evidence to suggest that right wing populists are increasing their support for women's rights gender equality over time. The existing trend is negative if existent at all, but the smaller sample size for gender scoring makes it difficult to establish a trend with a high degree of confidence Figure 6. On the whole, there is no solid evidence that pro-gender equality scores are increasing among right wing populist parties, and it may be the case that they are decreasing.

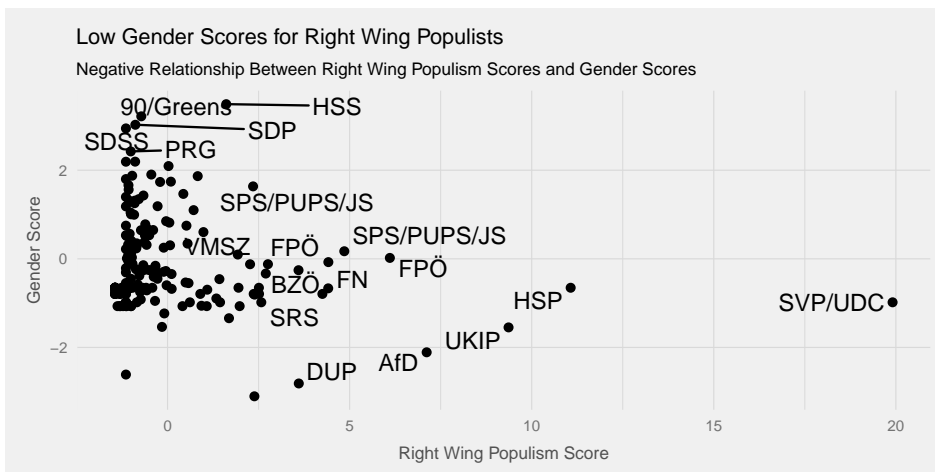


Figure 5. Populism vs support for gender equality

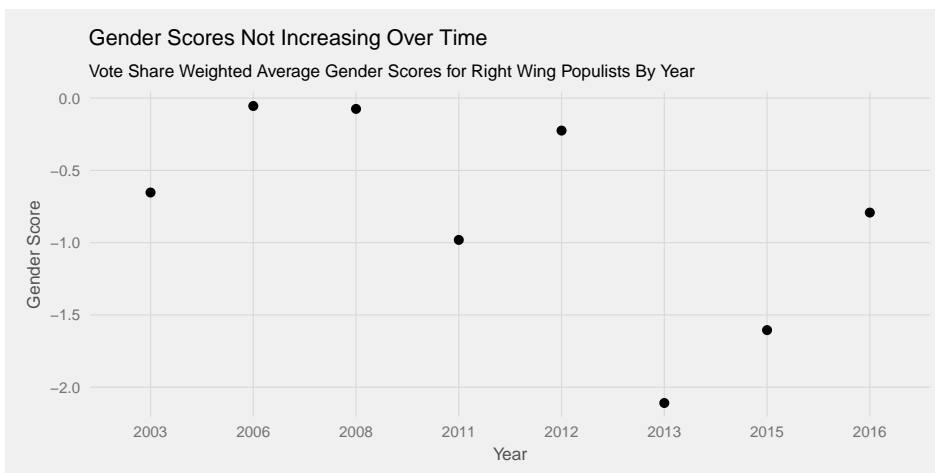


Figure 6. Gender-equality scores over time

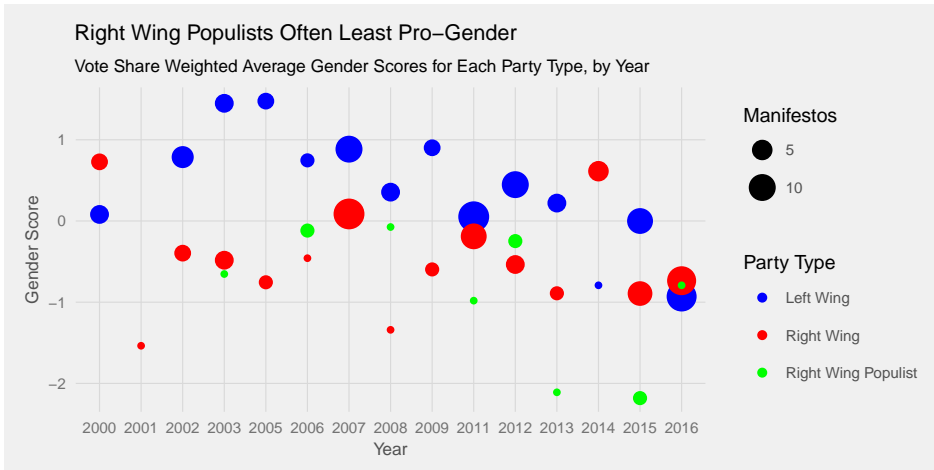


Figure 7. Gender-equality scores by party type

Similar to the comparison of right wing populist parties and traditional right wing parties concerning social welfare issues, there is no real evidence to support the claim that right wing populist parties favor women’s rights and gender equality more than their non-populist right wing counterparts. When measured year by year, right wing populists score lower than non-populist right wing parties in the majority of years, although again the limited sample set leads to some missing years for populists Figure 7. However, lower gender scores for right wing populists are mostly explained by higher average right wing scores for right wing populists than non-populist right wing parties. In a linear regression, controlling for left-right position leaves populism as a weak and insignificant predictor of gender scores, though the sign is still negative (Table ??). For gender, then, the evidence seems to suggest that right wing populists are generally in line with their non-populist equivalents on the right. On the whole, the case for right wing populists as a new type of more pro-gender equality, right wing party is not supported by the evidence, and lacks even the small concessions present with welfare scores.

This brings us to the question posed earlier: if right wing populists are fairly similar to non-populist right wing parties in their positions on gender equality, what leads to the impression that they are different? Unlike with welfare, there does not appear to be a meaningful relationship between vote share and gender scores, though such a relationship could be obscured by the limited sample size. Rather, it seems likely that public perception of right wing populist support for gender equality may be driven by campaign rhetoric emphasizing these parties’ defense of

European values, including gender equality, against Islam in general and Muslim immigrants in particular. Especially in the case of gender, it is important to note the potential discrepancies between party rhetoric and actual party manifesto content. “Pseudo-feminist” rhetoric may be increasingly common in right populist campaigns, as with the FPÖ’s campaign posters proclaiming “women’s right to self-determination”, but these rhetorical tactics may not necessarily translate into substantive policy material in manifestos. Particularly if pro-gender-equality stances are used merely as a vehicle for anti-Muslim messages, they may be difficult to detect through manifestos.

Conclusion

This paper brings new empirical findings to the discussion of how right wing right wing populism in Europe actually is. We provide a new measure of right wing populism based on our theoretical understanding of right wing populism in Europe and implement it using four existing codings in the Manifesto Corpus. Having defined the universe of right wing populist parties in Europe, we can then empirically and systematically examine previous claims that right wing populists are more pro-welfare and pro-gender-equality than their non-populist counterparts. We have chosen these issues because they feature prominently in the scholarly literature and in public perception about how these right wing populists are different from other right wing parties. We are particularly interested in these parties’ support for women’s rights and gender equality because this has received little study to date.

To measure gender equality, we needed to construct a new measure that was not in the existing Manifesto Corpus. To do this, we turned to a convolutional neural network and a new set of hand-annotated quasi-sentences from the corpus to produce a very accurate statistical classifier. The main methodological finding of this paper is that highly accurate custom classifiers for specialized manifesto statements can be very rapidly produced using this kind of model and specialized active learning annotation software. We encourage other researchers to adopt this method. While we hoped to generate accurate scores across languages with the use of aligned word embeddings, we were unable to generate good cross-language predictions. Our experiments on countries that share a language show that at least a portion of this failure is attributable to differences in countries’ manifesto speech that is separate from their language, which is an area of methodological research that will require specialized knowledge in political science, not just in computer science and natural language processing.

Our new measure allow us to answer the question of whether right wing populists are more centrist than other right wing parties on the two key issues of the welfare state and gender equality. Our results demonstrate that no major differences exist between populist and non-populist parties on the right for these two issues, as measured by their manifestos. If anything, populists are further to the right. We argue that the perception that they are more centrist on welfare results from an availability heuristic, where more prominent populist parties and populist parties with larger vote shares are more supportive of the welfare state than most populist parties, causing people to misjudge the stances of populist parties as a whole. Our more comprehensive measure of right wing populist parties in Europe reveals that several of the most prominent populist parties such as UKIP and the FPÖ are actually outliers on welfare support for populists as whole. Distinguishing these outliers from the rest of right wing populist parties thus appears to be the most fruitful avenue for future work. The picture for gender is slightly different: populist parties as a whole are quite far on the right. A second avenue of future research would rely not only on manifesto text, but also on transcripts of campaign and legislative speeches and on press releases to measure parties' stances on the welfare state and gender. It could be that populists' positions are not different from other right wing parties in their manifestos, but are quite different in the rhetorical appeals they make to voters, especially on gender issues. Studying this question will require new data collection efforts, but our methodological work indicates that this may be more feasible than it was previously.

References

- Afonso, Alexandre, and Yannis Papadopoulos. 2015. "How the Populist Radical Right Transformed Swiss Welfare Politics: From Compromises to Polarization." *Swiss Political Science Review* 21 (4). Wiley Online Library: 617–35.
- Akkerman, Tjitske. 2015. "Gender and the Radical Right in Western Europe: A Comparative Analysis of Policy Agendas." *Patterns of Prejudice* 49 (1-2). Taylor & Francis: 37–60.
- Akkerman, Tjitske, and Anniken Hagelund. 2007. "'Women and Children First!' Anti-Immigration Parties and Gender in Norway and the Netherlands." *Patterns of Prejudice* 41 (2). Taylor & Francis: 197–214.
- Albertazzi, Daniele, and Duncan McDonnell. 2007. *Twenty-First Century Populism: The Spectre of Western European Democracy*. Springer.
- Aslanidis, Paris. 2016. "Is Populism an Ideology? A Refutation and a New Perspective." *Political Studies* 64 (1_suppl). SAGE Publications Sage UK: London, England: 88–104.
- Bacchetta, Paola, and Margaret Power. 2013. *Right-Wing Women: From Conservatives to Extremists Around the World*. Routledge.
- Betz, Hans-Georg. 1994. *Radical Right-Wing Populism in Western Europe*. Springer.
- . 2013. "Minarets, Burqas, and Other Central Threats." In *Right-Wing Populism in Europe: Politics and Discourse*, edited by Ruth Wodak, Brigitte Mral, and Majid KhosraviNik. A&C Black.
- Betz, Hans-Georg, and Susi Meret. 2009. "Revisiting Lepanto: The Political Mobilization Against Islam in Contemporary Western Europe." *Patterns of Prejudice* 43 (3-4). Taylor & Francis: 313–34.
- Blee, Kathleen M, and Sandra McGee Deutsch. 2012. *Women of the Right: Comparisons and Interplay Across Borders*. Penn State Press.
- Bogdanor, Vernon. 2016. "The British General Election of 2015 and the Rise of the Meritocracy." *The Political Quarterly* 87 (1). Wiley Online Library: 39–45.
- Bojanowski, Piotr, Edouard Grave, Armand Joulin, and Tomas Mikolov. 2016. "Enriching Word Vectors with Subword Information." arXiv Preprint arXiv:1607.04606.
- Cheles, Luciano, Ronnie Ferguson, and Michalina Vaughan. 1995. *The Far Right*

in Western and Eastern Europe. Longman Pub Group.

De Koster, Willem, Peter Achterberg, and Jeroen Van der Waal. 2013. "The New Right and the Welfare State: The Electoral Relevance of Welfare Chauvinism and Welfare Populism in the Netherlands." *International Political Science Review* 34 (1). Sage Publications Sage UK: London, England: 3–20.

Dyer, Chris, Victor Chahuneau, and Noah A Smith. 2013. "A Simple, Fast, and Effective Reparameterization of IBM Model 2." In. *Association for Computational Linguistics*.

Eatwell, Roger. 2013. *Fascism and the Extreme Right*. Routledge.

Golder, Matt. 2016. "Far Right Parties in Europe." *Annual Review of Political Science* 19. *Annual Reviews*: 477–97.

Irvine, Jill A, and Carol S Lilly. 2007. "Boys Must Be Boys: Gender and the Serbian Radical Party, 1991–2000." *Nationalities Papers* 35 (1). Taylor & Francis: 93–120.

Jansen, Robert S. 2011. "Populist Mobilization: A New Theoretical Approach to Populism." *Sociological Theory* 29 (2). Wiley Online Library: 75–96.

Kitschelt, Herbert. 2007. "Growth and Persistence of the Radical Right in Post-industrial Democracies: Advances and Challenges in Comparative Research." *West European Politics* 30 (5). Taylor & Francis: 1176–1206.

Kostadinova, Tatiana, and Anna Mikulska. 2015. "The Puzzling Success of Populist Parties in Promoting Women's Political Representation." *Party Politics*. SAGE Publications Sage UK: London, England, 1354068815601604.

Lehmann, Pola, Theres Matthieß, Nicolas Merz, Sven Regel, and Annika Werner. 2017. "Manifesto Corpus. Version: 2017b." Berlin: WZB Berlin Social Science Center.

Levy, Omer, and Yoav Goldberg. 2014. "Neural Word Embedding as Implicit Matrix Factorization." In *Advances in Neural Information Processing Systems*, 2177–85.

Levy, Omer, Anders Søgaard, and Yoav Goldberg. 2017. "A Strong Baseline for Learning Cross-Lingual Word Embeddings from Sentence Alignments." In *Proceedings of the 15th Conference of the European Chapter of the Association for Computational Linguistics: Volume 1, Long Papers*, 1:765–74.

Lewandowski, Jirka, Nicolas Merz, Sven Regel, and Pola Lehmann. 2015. "ManifestoR: Access and Process Data and Documents of the Manifesto Project." R

package version.

Lowe, Will, Kenneth Benoit, Slava Mikhaylov, and Michael Laver. 2011. "Scaling Policy Preferences from Coded Political Texts." *Legislative Studies Quarterly* 36 (1). Wiley Online Library: 123–55.

Mayer, Stefanie, Edma Ajanovic, and Birgit Sauer. 2014. "Intersections and Inconsistencies. Framing Gender in Right-Wing Populist Discourses in Austria." *NORA-Nordic Journal of Feminist and Gender Research* 22 (4). Taylor & Francis: 250–66.

McGann, Anthony J, and Herbert Kitschelt. 2005. "The Radical Right in the Alps: Evolution of Support for the Swiss Svp and Austrian FpÖ." *Party Politics* 11 (2). Sage Publications London, Thousand Oaks, New Delhi: 147–71.

Mény, Yves, and Yves Surel. 2002. "The Constitutive Ambiguity of Populism." In *Democracies and the Populist Challenge*, 1–21. Springer.

Mikolov, Tomas, Ilya Sutskever, Kai Chen, Greg S Corrado, and Jeff Dean. 2013. "Distributed Representations of Words and Phrases and Their Compositionality." In *Advances in Neural Information Processing Systems*, 3111–9.

Montani, Ines, and Matthew Honnibal. 2018. "Prodigy: A New Annotation Tool for Radically Efficient Machine Teaching." *Artificial Intelligence* to appear.

Mudde, Cas. 2007. *Populist Radical Right Parties in Europe*. Vol. 22. 8. Cambridge University Press Cambridge.

———. 2010. *The Ideology of the Extreme Right*. Manchester University Press.

Pelinka, Anton. 2013. "Right-Wing Populism: Concept and Typology." In *Right Wing Populism in Europe: Politics and Discourses*, edited by Ruth Wodak, Brigitte Mral, and Majid KhosraviNik, 149–62. A&C Black.

Pilkington, Hilary. 2017. "'EDL Angels Stand Beside Their Men... Not Behind Them': The Politics of Gender and Sexuality in an Anti-Islam (Ist) Movement." *Gender and Education* 29 (2). Taylor & Francis: 238–57.

Rashkova, Ekaterina R, and Emilia Zankina. 2017. "Are (Populist) Radical Right Parties Männerparteien? Evidence from Bulgaria." *West European Politics* 40 (4). Taylor & Francis: 848–68.

Rovny, Jan. 2013. "Where Do Radical Right Parties Stand? Position Blurring in Multidimensional Competition." *European Political Science Review* 5 (1). Cambridge University Press: 1–26.

Ruder, Sebastian, Vulić, and Anders Søgaard. 2017. "A Survey of Cross-Lingual

Word Embedding Models.” arXiv <https://arxiv.org/abs/1706.04902>.

Sartori, Giovanni. 1970. “Concept Misformation in Comparative Politics.” *American Political Science Review* 64 (4). Cambridge University Press: 1033–53.

Smith, Samuel L, David HP Turban, Steven Hamblin, and Nils Y Hammerla. 2017. “Offline Bilingual Word Vectors, Orthogonal Transformations and the Inverted Softmax.” arXiv Preprint arXiv:1702.03859.

Spierings, Niels, Andrej Zaslove, Liza M Mügge, and Sarah L de Lange. 2015. “Gender and Populist Radical-Right Politics: An Introduction.” *Patterns of Prejudice* 49 (1-2). Taylor & Francis: 3–15.

Srivastava, Nitish, Geoffrey E Hinton, Alex Krizhevsky, Ilya Sutskever, and Ruslan Salakhutdinov. 2014. “Dropout: A Simple Way to Prevent Neural Networks from Overfitting.” *Journal of Machine Learning Research* 15 (1): 1929–58.

Stanley, Ben. 2008. “The Thin Ideology of Populism.” *Journal of Political Ideologies* 13 (1). Taylor & Francis: 95–110.

Tversky, Amos, and Daniel Kahneman. 1973. “Availability: A Heuristic for Judging Frequency and Probability.” *Cognitive Psychology* 5 (2): 207–32.

Yang, Zichao, Diyi Yang, Chris Dyer, Xiaodong He, Alexander J Smola, and Eduard H Hovy. 2016. “Hierarchical Attention Networks for Document Classification.” In *HLT-Naacl*, 1480–9.

Appendix A: Multilingual Document Classification Model

Statistical Classifier for Pro-Gender Equality Statements

Reading and coding every quasi-sentence in the Manifesto Corpus would have been entirely infeasible for this project. Instead, we turned to a machine learning approach that would allow us to annotate only a small fraction of a country’s quasi-sentences and then automatically generate labels for the rest. In addition to generating labels for the languages we annotated in, we also experimented with a system that could generate predictions even for languages without training data. We describe the model below and then then turn to the results. Document classification tasks require three decisions: how the document will be represented, how labels for the documents will be represented, and what model to use. Our classifier is a multilayered convolutional neural net that takes as input quasi-sentences represented as a sequence of word embeddings. The output of the model is a binary decision for whether the quasi-sentence is a statement in favor of gender equality.

Word embeddings. Traditional text analysis assigns each word in the vocabulary to a number, and represents words in statistical models as a vector the size of the vocabulary that is entirely 0s except for a 1 in the location of that word’s vocabulary number. In the past four years, the state of the art in text classification and natural language processing more generally has shifted to an approach that represents words as dense, low-dimensional vectors called word embeddings (Mikolov et al. 2013; Levy and Goldberg 2014). Word embeddings are a method of learning a low-dimensional representation of words from text based on local co-occurrence with other words. Word embeddings represent words as much shorter vectors (300 is standard) that are dense, meaning each dimension has a continuous value. The embeddings can be learned as a hidden layer in a shallow neural net that tries to predict a word based on its neighbors (Mikolov et al. 2013) or as a matrix factorization on word co-occurrence (Levy and Goldberg 2014). These dimensions, like the “topics” in topic models, have no inherent meaning, but are learned automatically from the text and may in some situations have an interpretation.

The advantages of these distributed representations for document classification are two-fold: first, by representing words as dense vectors rather than one-hot vectors, the dimensionality of vectors needed to represent a language’s vocabulary can be reduced from hundreds of thousands to merely hundreds. Second, by making the dimensions and words’ locations in them meaningful, models can learn much

more efficiently and can handle previously unseen words by being able to compare them to previously seen words. The similarity between any one-hot vectors is 0, meaning a model needs to see each word before it can learn what weight to give it. Continuous embeddings, on the other hand, can look similar, so that a model that has learned a weight for “woman” would be able to infer a weight for the previously unseen word “girl”.

Recent work has examined the possibility of aligning word vectors across languages to make words embeddings to be proximate with their translations in other languages in embedding space (see Ruder, Vulić, and Søgaard 2017 for a recent survey). Work had shown that learning multilingual aligned vectors was possible during the training phase, but this approach requires large computational resources and access to very large quantities of text aligned text (text with known translations across languages) (Dyer, Chahuneau, and Smith 2013; Levy, Søgaard, and Goldberg 2017). S. L. Smith et al. (2017) show that multilingual embeddings can be learned very quickly on pre-trained word embeddings using singular value decomposition. We use their method to align the pre-trained Fasttext vectors provided by Bojanowski et al. (2016) for use in our model. Initially, “woman” and “frau” would have vectors with no apparent similarity or dissimilarity, since they were learned in two separate processes with two separate text corpora. After alignment, the two words should have similar vectors since their meanings are equivalent in the two languages.

Aligning embeddings across languages using the method described in S. L. Smith et al. (2017) requires a small number of known translated word pairs in order to learn the alignment across languages. We use three sources of translations for our alignment list. First, following S. L. Smith et al. (2017), we use shared word spellings across languages, on the intuition that they are often used in the same way. This often works (“Berlin”, across most Latin alphabet languages), but other times does not (“gift” across German and English). Second, we take the most common 500 words from the UK manifesto corpus, minus 25 stopwords, and use Google Translate to translate them into each target language. Finally, we use a hand constructed list of gender-related terms and also use Google Translate to translate them into each European language.

Using these word embeddings, we transform each document (quasi-sentence) into a matrix with the shape (words, embedding_dim). Most existing text analysis in political science represents documents as a single vector, consisting of the sum of all words’ one-hot vector. This “bag of words” representation has the advantage of being more compact than a matrix representation and performs remarkably well on many tasks, but loses any information that is contained in word order. By representing each document as a matrix in our model, we are able to preserve the

information that word order conveys.

Model

To generate a binary classification for whether the quasi-sentence contains a pro-gender equality statement, we feed each document into a convolutional neural net. Convolutional neural nets (CNNs) are a form of neural net that look for features of local regions in an input, learning which features are important and sharing local model weights across the entire input. In computer vision, where CNNs are the state of the art, a convolutional layer will learn to detect shapes or more complicated features such as faces regardless of where in the image they occur. By stacking convolutional layers, information can be shared across input units, such that “gender equality” will generate a feature that is distinct from “gender” and “equality” in other parts of the sentence.³ The neural network we use is publically available through the spaCy library and is a modified implementation of Z. Yang et al. (2016), using CNNs instead of recurrent neural nets, but maintaining the novel attention mechanism.

Annotation procedure

To generate a classifier that works well across languages and to evaluate its cross-lingual performance, we labeled several hundred examples per language of pro-gender equality statements (and statements that are not) in English, French, German, Croatian, Bosnian, and Serbian. We used an annotation tool that uses active learning (Montani and Honnibal 2018) to determine which pieces of text to

³Why not a recurrent neural net? Awareness of local features and input order is a hallmark of recurrent neural networks, another approach to modeling inputs from convolutional neural nets. Recurrent neural nets operate on sequence data, making them ideal for modeling speech and text. We opt for a convolutional neural net, though, for three reasons. First, the property of location invariance that CNNs provide is useful when modeling text across languages, where word order and grammar are often completely different. CNNs provide some local pooling of information, but do not model the entire sequence at once. Second, CNNs are much faster to train than recurrent neural nets, which can require specialized hardware and a great deal of time. Finally, we use an existing CNN architecture that is optimized for short text classification and that integrates well with the natural language processing and data annotation tools we use.

display to the annotator. By favoring examples where the model most unsure over examples where it is very confident, the tool maximizes the expected information gain per labeled example. The use of active learning and specialized annotation software meant that we could collect 1,764 labeled quasi-sentences in a only a few hours of annotation. The number of annotations (both positive and negative) per country are presented in the table below.


table 2 Number of labeled quasi-sentences per country

Country	Labeled Quasi-Sentences
Austria	384
Bosnia-Herzegovina	70
Croatia	334
France	152
Germany	264
Ireland	123
Luxembourg	134
Serbia	63
United Kingdom	104

The first experiment we ran was to verify that we could get good classification performance within each country. To do so, we randomly sampled 20% of each country’s annotations and set them aside as an evaluation set. We trained an individual model for each country and evaluated it on the held out set. Figure ?? and the table below show that we achieved quite good accuracy for most countries, as measured by the model’s F1 score.⁴ The only two consistently poor performers were Bosnia-Herzegovina and Serbia, each of which had fewer than 60 quasi-sentences to train on. Luxembourg was inconsistent across runs, ranging from 94% to less than 50%, perhaps because its manifestos also include some German text. Smaller languages may also perform worse than larger languages because the quality of their word embeddings will be poorer: word embeddings are best when trained on hundreds of millions of words, and the language-specific Wikipedias ours are trained on vary greatly in size.⁵

⁴F1 scores are the harmonic mean of precision and recall: $2 \cdot \frac{\text{precision} \cdot \text{recall}}{\text{precision} + \text{recall}}$. This metric is preferred over accuracy for class-imbalanced classification tasks like this one. If gender equality statements make up 0.5% of the corpus, the model could achieve 99.5% accuracy by always predicting 0. Because this model’s recall would be 0, its F1 score would also be, better reflecting an intuitive sense of model quality.

⁵The small sample size for Bosnia-Herzegovina and Serbia also leads to another problem that becomes apparent in Figure ??. Because precision and recall are



	Country	F1 score	training n
1	Austria	0.97	308
2	Bosnia-Herzegovina	0.29	56
3	Croatia	0.92	268
4	France	0.89	122
5	Germany	0.85	212
6	Ireland	0.95	99
7	Luxembourg	0.94	108
8	Serbia	0.40	51
9	United Kingdom	0.95	84

The heartening conclusion we draw from these results is that CNNs with pre-trained word embeddings can learn to very accurately classify statements in support of gender equality with only 100-200 training examples. Using specialized annotation software, it is feasible to collect this amount of data in less than one hour per language.⁶ Because each country is learned separately, aligned word embeddings are not needed and a non-aligned model (the baseline model) matches or outperforms models with aligned embeddings.

Next, we wanted to test how well the model performed across countries but within a given language. Figure 8 shows that for our multi-country languages, the perfor-

defined in terms of true positives, evaluation samples that do not contain a true positive will result in precision and recall being undefined and the result dropping out.

⁶5 seconds per annotation \times 200 = 16.6 minutes.

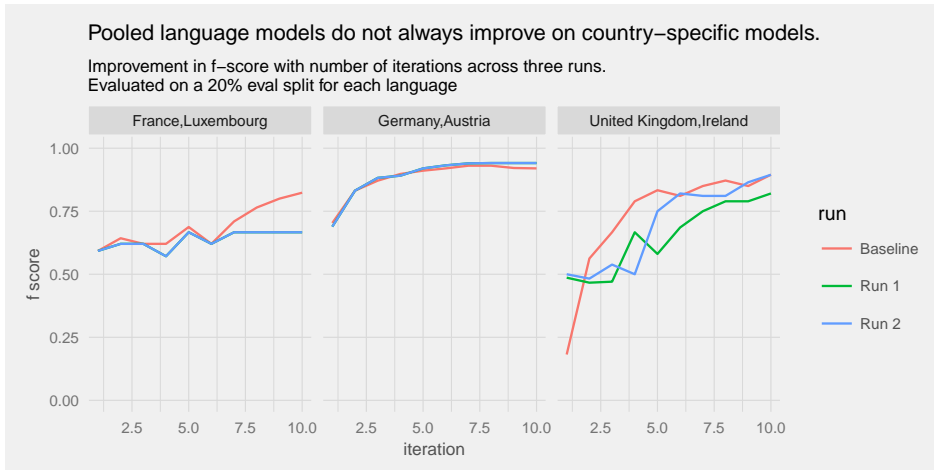


Figure 8. Accuracy for countries pooled by language

mance on a 20% eval sample is about the same or worse than when trained on one country only. This is our first indication of the difficulty of generalizing across countries, even when keeping language constant. This is confirmed by a direct test Figure 9 of training on one country in a language and testing on a second country sharing its language. The evaluation set is the model’s accuracy on a held-out 20% sample and matches the results from the single country training: the model is able to very accurately classify pro-gender equality statements from a country it has trained on. The second test accuracy line shows the accuracy on a the second country, which the model did not see during training.

What this shows is that there is some country-specific effect that limits how well models can travel. Here, the evaluation accuracy (accuracy on a random sample of the training country) is higher than the test accuracy (the accuracy on the target language). Even though the two countries share a language, the model performs worse on the second country than on unseen data from the training country. The test country is different enough from the training/evaluation country that, despite sharing a language, the model cannot generalize well to it from the training data.

Our final experiment was to train our classifier on a set of languages and then apply it to a wholly unseen language during testing. If the aligned embedding alignment were good enough, word order was relatively consistent, and country-specific terms relatively rare or unimportant, this approach could have promise. Instead, we find very poor cross-language performance. Remarkably, the model trained on all countries except Germany and Austria and then tested on Germany




Figure 9. Testing the model within language, across countries

performs quite well, but this is likely to be a fluke. The results shown here use a model with a greater dropout proportion (0.5) than previous models (0.2), which led to slightly better accuracy as a result of less overfitting.⁷ As we would expect, the model is quite good at learning the sample of text given to it as training data, even across many languages (the “eval” line in Figure 10), but the test accuracy is very poor.

Classifier conclusion

Our research on issues classifiers for the Manifesto Corpus thus shows both reasons for optimism and pessimism for future automated classification work. We show that a convolutional neural network with pre-trained word embeddings can learn to classify pro-gender equality statements in the manifesto corpus with very high accuracy with only 100-200 training examples. These classifications are reproducible, have confidence/uncertainty scores attached, and can be generated instantly for new manifestos in each trained language, giving them a marked advantage over manual annotation. At the same time, our work shows reasons for

⁷Dropout is a form of regularization in neural nets to limit overfitting. A random proportion of model weights during each backpropagation step are “frozen” to prevent them from being updated (Srivastava et al. 2014).




Figure 10. Testing model on new languages

pessimism about a generalized, cross-language classifier. Our results from cross-country classification within the same language (Austria/Germany, Ireland/UK) reveal a notable degree of country-specificity. Even within one language, countries have different ways of discussing the same concept, limiting the applicability of models across countries. Imperfect embedding alignment, differences in sentence construction and word order across languages, and greater political differences mean that the prospects for cross-language classification using this approach appear limited. Future work could draw on aligned text (text for which a direct translation is provided) to improve the model. A multitask neural net could generate sentence classification labels for each of the matched sentences with a single set of shared weights, operating on the aligned word embeddings. At the same time, the model could be penalized for the difference in the model's representation of each language. By penalizing loss on multiple tasks, we should be able to reduce what appears to be the model's overfitting within each language.