Group Appeals and Political Mobilization: Evidence from U.S. House Races *

Alonso Roman Amarales[†] Luca Bellodi[‡] Catherine E. De Vries[§]

Abstract

Group membership plays a crucial role in political conflict, with candidates frequently using group appeals to mobilize voters. However, the factors influencing candidates' decisions to employ such appeals remain insufficiently understood. We propose an argument rooted in the notion that candidates selectively deploy group appeals to maximize electoral returns and we introduce a novel method for detecting group appeals in text, which is applied to the universe of tweets posted by U.S. House candidates between 2012-2021. We present three findings. First, candidates' group membership and their district's demographic composition are strongly associated with the frequency of group appeals. Second, protest events associated with specific groups increase appeals to those groups. Third, except for young voters, there is limited evidence that these appeals enhance turnout intentions. These findings offer new insights into how politicians strategically engage underrepresented groups in the electorate, contributing to our understanding of identity-based political mobilization.

Word Count: 9,920

*The authors would like to acknowledge funding by the "Next Generation EU – Age-It – Ageing well in an aging society" project (PE0000015), National Recovery and Resilience Plan (NRRP) - PE8 - Mission 4, C2, Intervention 1.3. The views and opinions expressed are only those of the authors and do not necessarily reflect those of the European Union or the European Commission. Neither the European Union nor the European Commission can be held responsible for them. Arianna Savo provided excellent research assistance.

[†]Bocconi University

[‡]Stanford University

[§]Bocconi University

Introduction

Political mobilization is a cornerstone of democratic governments and a critical determinant of the electoral success of political candidates. Scholarship on voter mobilization emphasizes that turnout is shaped not only by individual political preferences, but also by the strategic choices of politicians in targeting prospective voters (e.g., Rosenstone and Hansen, 2003; Holbrook and McClurg, 2005; McGhee and Sides, 2011). This dynamic is particularly salient in the United States, where competitive elections mean that even small shifts in turnout can prove decisive (for an overview, see Jacobson, 2015). As a result, a key challenge for political candidates is balancing efforts to energize their base while simultaneously broadening their electoral coalition.

While traditional mobilization research has largely emphasized policy positioning, particularly candidates' alignment with the median voter (see, e.g., Adams, 2012), recent scholarship highlights an alternative strategy: direct appeals to social and demographic groups (e.g., Thau, 2018; Horn et al., 2021; Robison et al., 2021; Huber, 2022; Dolinsky, 2023; Haffert et al., 2024; Finseraas et al., 2025). These group appeals (e.g., campaign messages that explicitly reference a social or demographic group) have gained increasing prominence as large-scale societal transformations, including globalization, automation, and immigration, have disrupted traditional political alignments (e.g., Baccini and Sattler, 2024; Dal Bo' et al., 2023; Colantone and Stanig, 2018; Dalton, 2013; Dehdari, 2022). Consequently, candidates can no longer count on the stable support of particular demographic groups, but must instead construct broad and diverse electoral coalitions. This often requires strategically activating, reinforcing, or expanding group-based identities to mobilize voters in a complex political landscape.

Building on classical models of strategic mobilization (Rosenstone and Hansen, 2003) and incorporating recent empirical research on group appeals, we argue that candidates deploy group-based appeals selectively when such strategies maximize electoral returns. Departing from much of the existing literature, which emphasizes party-level incentives (see, e.g., Thau, 2024; Dolinsky, 2023), we focus on how electoral institutions and individual candidate characteristics shape candidates' strategic choices. First, candidates are more likely to appeal to groups that constitute a substantial share of the local electorate, as larger groups offer a more viable electoral coalition

(Posner, 2004; Eifert et al., 2010). Second, the importance of specific groups fluctuates in response to societal events, such as mass protests for example, which can temporarily heighten the perceived salience of certain groups and prompt candidates to adjust their messaging (Wasow, 2020; Gillion, 2012). Third, candidates' own demographic characteristics influence their strategic decisions, as shared identity enhances the credibility and persuasiveness of political messaging, fostering trust and engagement (e.g., on race and gender Kuklinski and Hurley, 1994; Dolan, 1998; Schaffner, 2005). By aligning their rhetoric with these three strategic considerations—structural (group size), contextual (group salience), and personal (candidate identity)—candidates seek to maximize their electoral advantage while adapting to shifting political dynamics.

Empirically, we assess these expectations by introducing the *group appeals detector*, a novel algorithmic tool for identifying group appeals in text data. Our approach integrates multiple natural language processing techniques to systematically detect explicit references to social and demographic groups, either as active agents (e.g., "Young people deserve a better future") or as recipients of political commitments (e.g., "I will fight for young people's future"). We demonstrate that our measure performs comparably to modern large language models while offering researchers a transparent and scalable approach for analyzing group appeals across time and space. We apply this approach to the universe of tweets posted by U.S. House candidates between 2012 and 2021 and we extract appeals directed at three key demographic groups – young, female, and Black voters – who have historically exhibited lower levels of political participation (Norris, 2004; Mycock and Tonge, 2012; Kitanova, 2020; Weiss, 2020; Burns et al., 1997; Gallego, 2007; Paxton et al., 2007; Leighley and Nagler, 2013).

Our empirical analysis reveals several important findings. First, candidates' group appeals strongly correlate with the demographic composition of their districts. Candidates in districts with a larger share of young people and African Americans make more appeals to these groups, whereas we do not find that a large share of women increases the supply of appeals to female voters. Second, in-group candidates are more likely to appeal to the group to which they belong. Female, Black, and young candidates are significantly more likely to make appeals to female, Black, and young voters, respectively. Third, beyond these structural and individual-level factors, we find that candidates' supply of group appeals responds to major protest events occurring in their districts, which heighten the salience of specific groups. Leveraging the staggered timing of Fridays for Future and Black Lives Matter protests as quasi-random shocks to group salience across districts, we show that candidates in districts experiencing a protest increase their appeals to young and Black voters in the days following the protest. Finally, we match candidates' Twitter appeals to a large public opinion survey to examine, in a real-world, observational setting, whether respondents exposed to a group appeal are more likely to report being willing to turn out in support of the candidate of the party they identify with. Except for young respondents, we find little evidence that social media appeals are effective in mobilizing voters.

This study makes three key contributions to the literature on group appeals and political mobilization. First, it develops a theoretical argument that situates candidates' use of group appeals within the broader scholarship on strategic mobilization and political targeting, highlighting the electoral incentives that shape these rhetorical strategies. Second, it introduces a novel empirical approach for detecting group appeals in candidate rhetoric, offering a transparent and scalable method for systematically identifying explicit group-based messaging in campaign discourse. Finally, it provides new empirical evidence on the determinants and consequences of group appeals, clarifying when and how candidates strategically deploy these messages and assessing their effectiveness in shaping voter participation. By integrating theoretical and methodological innovations, this study advances our understanding of identity-based electoral mobilization and offers broader insights into the strategic calculations that shape the political rhetoric of candidates in contemporary elections.

Candidates' Supply of Group Appeals

While existing work on political opinion formation suggests that group belonging can be activated or deactivated in voter decision-making (Nelson and Kinder, 1996; Kinder, 1998; Kinder and Kam, 2010), only a few studies have examined the conditions under which political candidates rely on group appeals and the extent to which this rhetorical strategy successfully mobilizes voters. A large body of literature, primarily focused on European countries, has documented secular shifts in parties' appeals to social groups. Thau (2018, 2019) shows that social democratic parties in Denmark and the United Kingdom have replaced class-based appeals with appeals to other "non-economic groups" based on content analysis of party programs (see also Robison et al. (2021) and Huber (2022)). By developing a dataset on group and policy appeals in Scandinavia based on party programs, Horn et al. (2021) show that political parties' group and policy appeals, both on the left and right, are increasingly focused on broad demographic group categories beyond class (for datasets on group appeals also see Dolinsky, 2023). Besides providing valuable descriptive trends, this work does not examine the conditions under which candidates or parties make appeals.

As far as the consequences of group appeals are concerned, the evidence is mixed. Experimental work from the United States highlights the potential electoral benefits of "identity-based appeals" based on gender in political advertising (Holman et al., 2015), and recent findings show that group appeals made by British parties are indeed effective at strengthening the perceived linkages between voters and parties (Dausgaard and Hjorth, 2025). Similarly, Thau (2021) finds that class voting in the UK responds to class-based appeals and Huber et al. (2024) find that respondents randomly assigned to a policy statement with a clear reference to a group they like are more supportive of the statement, whereas if they dislike the group, their support decreases. Similar findings are reported by Finseraas et al. (2025), who show that working-class voters in Britain and Norway are more likely than middle-class voters to prefer candidates who make group appeals specifically targeted at the working class.

Despite these important insights on the use and consequences of group appeals, the literature to date has not yet fully explored how the strategic use of these appeals by candidates varies across time and space in predictable ways. In the next section, we develop a strategic group appeals argument by fusing insights from the group appeal literature with research on the strategic mobilization of candidates aimed at increasing voter participation.

Group Appeals as a Tool for Strategic Mobilization

Theories of strategic mobilization highlight several obstacles to ordinary citizens' political engagement, as acquiring political information can be costly and is often perceived as having limited personal utility (Rosenstone and Hansen, 2003). When campaigning, candidates employ various strategies to reduce these costs and mobilize their constituents, such as targeted messaging and group appeals (Ansolabehere et al., 1994; Green and Gerber, 2019; Zhuravskaya et al., 2020). By appealing to groups, candidates can increase the perceived stakes in elections, enhance group members' political participation (Kinder and Kam, 2010; Klar, 2013), and strengthen their linkages with voters in an attempt to consolidate their electoral coalitions (Dalton, 2013; Miller et al., 1991; Nelson and Kinder, 1996; Dausgaard and Hjorth, 2025).

However, candidates have limited time and resources, and appealing to *any* group can be counterproductive, especially if appeals to a group may alienate members of other groups. Therefore, which groups do candidates prioritize? Candidates must be strategic in the way they appeal to groups. Specifically, we expect candidates to align their rhetoric with three strategic considerations: structural (group size), contextual (group salience), and personal (candidate identity).

First, we expect the strategic use of group appeals to vary based on electoral rules. If candidates use group appeals as a mobilization strategy, electoral institutions become critical in determining which groups are most advantageous to target (for an overview see Cox, 2015). Electoral rules shape the nature of political competition by defining the incentives for candidates to appeal to specific constituencies. In particular, they influence which voter blocs are necessary for electoral success and, consequently, which groups are most likely to receive targeted appeals (Cox, 1999; Carey and Shugart, 1995).

Majoritarian systems, such as first-past-the-post with single-member districts (SMDs), compel candidates to focus on geographically concentrated voter bases (Powell and Powell Jr, 2000). Candidates in SMD systems must secure a plurality of votes within a localized constituency, incentivizing them to tailor their appeals to the most electorally significant groups within their district. Conversely, candidates in proportional systems are more likely to adopt rhetorical strategies that reflect broader national demographics rather than localized constituencies, as their electoral success depends on appealing to a more diverse and dispersed electorate (Huber and Powell Jr, 1994; Carey and Shugart, 1995). The U.S. electoral system for the House of Representatives, characterized by SMDs and a plurality voting rule, exemplifies how district-level demographics influence candidate strategies. Since candidates need a plurality of votes within their district to win, they are expected to prioritize appeals to groups that are particularly salient within their respective constituencies. This district-centric approach to electoral competition means that group-based appeals will be more frequent when a particular demographic comprises a substantial portion of the local electorate (Fenno et al., 1978; Jacobson, 2015).

The most intuitive way to capture the importance of groups in electoral competition is through their size (Posner, 2004). For example, in districts where young voters constitute a significant proportion of the population, candidates have strong incentives to craft messages that resonate with youth interests. Similarly, in districts with a large share of African Americans, candidates are expected to increase their frequency of appeals to Black voters, recognizing the potential electoral benefits of mobilizing an influential voting bloc (Bobo and Gilliam Jr, 1990; Hajnal, 2006). These strategic considerations are particularly evident in candidates' appeals to female voters. Schaffner (2005) argues that candidates have a structural incentive to appeal to women, given their larger share of the electorate relative to men and their greater likelihood of being influenced by campaign messaging. These insights suggest that candidates' group-based appeals will systematically vary based on the demographic patterns of the district where they run.

> ***Prediction 1 – structural:*** *The supply of candidates' group appeals increases with the size of the group in their districts.*

Appeals to larger groups enhance the potential returns of targeted mobilization efforts, making these groups prime targets for campaign outreach. However, the electoral significance of groups may not only be a matter of group size *per se* but may change over time due to other factors. Rapid societal change, mass protests, and major political events can suddenly create political demands from certain groups, and candidates may see this as an opportunity to consolidate or gain support among them. For example, climate change and global warming have sparked new discussions around intergenerational justice, with young people being particularly vocal in demanding change. Events such as the climate protests organized by the Fridays for Future movement can heighten the electoral return of appealing to young voters. Similarly, protests like Black Lives Matter bring to the fore African Americans' demands for change regarding police violence, systemic racism, and racial discrimination, temporarily increasing candidates' incentives to appeal to their Black voter base.

Research on the politics of protests suggests that even though only a small fraction of the population participates in protests, their impact can be widespread, affecting agenda-setting and the positioning of political elites (Gause, 2022; Valentim, 2023). Protests can also create lasting shifts in public support for particular issues, especially when exposure is local and repeated (Mazumder, 2018; Tertytchnaya and Lankina, 2020). Moreover, protests may also serve as informational signals to politicians, helping them gauge which groups are more politically mobilized and, therefore, more likely to turn out and vote. When candidates face imperfect information about the political engagement of different demographic groups in their district, visible and sustained protests can serve as cues about voter priorities and issue salience (Wouters and Walgrave, 2017; Gillion, 2012).

Empirical research on historical protests supports these claims. Wasow (2020) examines Black-led protests in the 1960s and finds that peaceful demonstrations successfully garnered public sympathy and influenced political elites, while violent protests led to a backlash that shifted public opinion toward more conservative candidates. This phenomenon – which the author labels "agenda seeding" – highlights how the tactics employed by protest movements determine their broader political impact. In this sense, protest movements not only amplify the electoral importance of appealing to groups but also shape the strategic responses of political candidates. Candidates are expected to adjust their messaging in ways that either align with or distance themselves from these movements, based on their perceived political risks and opportunities.

Against this backdrop, we expect that sudden events, such as protests with a clear group connection, will increase candidates' supply of appeals to these groups, particularly when these protests are highly visible, sustained, and align with broader societal concerns.

> ***Prediction 2 – contextual:*** *Events increasing the perceived importance of a group are associated with an increase in candidates' appeals to the group.*

The effects of group importance – whether derived from the group's size or from events that heighten the salience of certain groups – highlight the demand-side factors of group appeals. However, one might expect that candidates' own characteristics also influence their strategic calculus in deciding which groups to appeal to. Minority candidates, for example, might believe their appeals to in-group minority voters will be more effective in mobilizing them. Similarly, young candidates might view their appeals to younger voters as more credible and, therefore, more likely to resonate and mobilize the young electorate.

Evidence from the literature on affective polarization and racially polarized voting suggests that in-group candidates – i.e., those belonging to the same demographic group they appeal to – are perceived as more trustworthy and effective in their appeals (Key, 1948; Kousser, 1974; Davidson and Grofman, 1994). For instance, Bobo and Gilliam (1990) find that African Americans in cities with African American mayors displayed higher rates of political efficacy and participation compared to African Americans in cities with white mayors. Similarly, Banducci et al. (2004) find that descriptive representation enhances minority political engagement, and Finseraas et al. (2025) show that working-class voters in Britain prefer working-class candidates.

Further supporting the importance of the messenger, Kuklinski and Hurley (1994) find that Black respondents were more likely to agree with an identical statement when it was attributed to an in-group member. This aligns with broader findings that appeals from in-group members tend to be more persuasive (e.g., Druckman, 2001). Hersh and Schaffner (2013) shows that the sense of linked fate stemming from shared identities between candidates and voters can increase support for the candidate. Besides race, some evidence shows that female voters see female candidates more favorably than male candidates (Dolan, 1998; Ansolabehere and Iyengar, 1994; la Cour Dabelko and Herrnson, 1997; Kahn, 1993; Schaffner, 2005). This implies that candidates will be more likely to appeal to groups they belong to, as their message will be perceived as more credible, thereby further enhancing voter mobilization.

> ***Prediction 3 – personal:*** *In-group candidates are more likely to make appeals to that group compared to candidates who do not belong to the group.*

By synthesizing insights from electoral institutions, political mobilization, and candidate identity, our argument provides a comprehensive framework for understanding the supply of group appeals in electoral politics.

Data & Measurement

We test our predictions by collecting novel data on politicians' appeals to groups on social media. Specifically, our empirical analysis focuses on the universe of Twitter content posted by Republican and Democratic candidates running for the U.S. House of Representatives from 2012 to 2021. The institutional setup of the U.S. House elections lends itself well to testing the predictions. First, candidates' campaigns rely heavily on social media, which are cheap and high-frequency channels for candidates to reach a diverse group of constituents. On average, 31% of respondents to the 2020 American National Election Study report visiting Twitter at least once a day in the past year (ANES, 2021). Second, the majoritarian electoral system allows districts and voters to be precisely matched with candidates, who, in turn, have electoral incentives to pander to the voters in their district.

District Data. We collect data on district-level characteristics from the American Community Survey (ACS), which provides yearly data on the U.S. population's social, economic, and demographic characteristics. In particular, we use the 5-year estimates, which enhance statistical reliability for smaller areas. For every district-year observation, we collect data on the age, gender, and racial composition of the population. We also use this source of data to collect other district-level socio-economic variables, including employment and income data.

Candidates & Elections Data. We obtain data on the names and parties of candidates, the districts where they ran, and their incumbency status from the *Candidates in American General Elections* database compiled by Cha et al. (2021).

We make significant original integrations to this dataset by adding data on the age, gender, and race of candidates. For winning candidates, we collect this information from the Biographical Directory of the U.S. Congress, which maintains biographic information for all present and former members of Congress.[1] For losing candidates, demographic information was mainly collected from Porter and Treul (2024) and Fraga et al. (2021). Additionally, data on candidates' ethnicity were integrated from Deshpande (2022b), Deshpande (2022a), and Barney (2017), as well as from extensive manual searches.

[1]We scraped gender and race data from the repository *Women Members by Congress* and *Black-American Members by Congress* available at `https://history.house.gov/`.

Table 1: Percentage of Candidates with a Twitter Account.

	Democrats			Republicans		
Election	Candidates	% Twitter	% Tweets > 0	Candidates	% Twitter	% Tweets > 0
2012	415	80.7	67.7	420	83.6	68.8
2014	406	85.2	71.7	397	88.4	73.6
2016	418	84.4	73.2	405	86.2	72.3
2018	442	97.5	85.3	399	93.0	78.7
2020	437	96.8	88.1	419	93.3	85.4

Notes: Total number of candidates running in the U.S. House Elections, percentage of candidates with an active Twitter account, and percentage of candidates with at least one tweet downloaded.

We obtain a total of 2,401 candidate-election observations. We recover the year of birth for 47% of the candidates, the race for 89%, and the gender for 99%. The median birth year of candidates is 1960, and the average age of candidates in the dataset is 56. 27% of candidates are women, 10% are black, 6% Hispanic, and 83% white.

Twitter Data. Starting from the list of candidates running for the U.S. House, we obtain the Twitter handle. Our extensive data collection follows Bellodi et al. (2023) and is described in greater detail in Section A in the Appendix.

From the 2,401 unique Democratic and Republican candidates that ran for office between 2012 and 2020, we obtained at least one Twitter account for 82.3%. This percentage increases over time and exceeds 95% when considering candidates running in the 2018 and 2020 elections.[2] We then passed a total of 2,560 unique Twitter handles belonging to 1,970 unique candidates to the Twitter API. For 81% of the 2,560 accounts, we successfully download tweets. Eventually, we obtain at least one tweet for 65% of the candidates in our data. Because many candidates run for office multiple times, when looking at election years separately, our coverage rate of candidates ranges between 67% in 2012 to 88% in 2020. Table 1 shows the share of candidates by party and election year with an active Twitter account and available tweets at the time of data collection.

Once we obtained at least one Twitter handle for each candidate, we downloaded the universe of tweets (in English and excluding retweets) posted by candidates from 2012 to 2021 through the Twitter API. For any candidate-election pair, we retrieve the universe of tweets posted during the

[2] Given the extensive scraping performed and the numerous iterations of manual checks, we are confident that the vast majority of the remaining candidates did not have an active Twitter account at the time of data collection.

year of the election and the following, for a total of 3,7 million tweets (61% posted by Democrats and 39% by Republicans).

Detecting Group Appeals

Our definition of group appeal is the rhetorical strategy through which a politician tailors her message to resonate with specific demographic groups. In spoken or written language, we assume that statements can appeal to groups in two ways. Groups can be portrayed as active agents (subjects of a statement, e.g., "Young people deserve a better future!"), or can be passive agents (recipients of a statement, e.g., "I will fight for the future of younger people!").

Based on this conceptualization, there are (at least) four ways in which researchers could extract group appeals from text corpora. First, researchers could hand-code documents and extract appeals to groups, which has been the predominant approach in the literature (Thau, 2018, 2019; Horn et al., 2021; Huber, 2022; Dolinsky, 2023). Second, appeals to specific groups could be detected by specifying a dictionary or list of words that capture group appeals (see e.g., Dausgaard and Hjorth, 2025). Third, researchers could rely on supervised classifiers trained on smaller portions of the corpus and then used to detect group appeals on out-of-sample documents (Licht and Sczepanski, 2023). Fourth, researchers could rely on large language models and zero-shot classification and simply "ask" the model whether the tweet contains an appeal to a pre-defined list of groups.

These different approaches have advantages and disadvantages. While enjoying high internal validity, human coding is labor-intensive and cannot be replicated or easily extended to other contexts. Dictionary methods are easy to implement, but they need to assume the correct specification of the dictionary, are not flexible enough to match various instances in which candidates might appeal to a group, and do not retain information about the context in which group-related words are used in the text. Supervised classifiers can be versatile and effective at detecting group appeals, but are generally costly to scale up for different groups. Finally, although large language models are considered to achieve the best performance for classification tasks, they are generally expensive to use for large datasets like ours, lack transparency, and are not efficient if researchers decide to add new groups to the set of target groups Barrie et al. (2024).

In the following section, we introduce and validate the *group appeal detector*, a novel measurement strategy that leverages various natural language processing techniques to overcome the primary limitations of the approaches discussed above, while achieving performance comparable to that of existing large language models.

The Group Appeal Detector

We propose a measurement strategy that aims to maximize flexibility, transparency, and performance when detecting group appeals within a text corpus. First, we extract all subjects and recipients from each tweet using dependency parsing. Second, we build, in a data-driven way, a set of words that captures references made on Twitter to the target groups and represent this as an average group-specific embedding. Third, we produce a vector representation of each subject/recipient from the first step and compute the semantic similarity to the group-specific embedding. If the similarity exceeds a certain threshold, we classify the tweet as containing an appeal to the group.

Step 1: Dependency parsing. Words in a text sequence, like a tweet, have specific syntactic features. They can represent different parts of speech (POS) (e.g., verbs, adjectives, or nouns), and they have dependency relations between one another (e.g., subjects, propositions, or direct objects). By combining the POS information of each word with its dependency relation, we can isolate the subjects and recipients in every tweet. For this task, we use `SpaCy` (Honnibal and Johnson, 2015), a supervised language model that achieves state-of-the-art performance on several syntactic tasks like POS tagging and dependency parsing.

To illustrate how POS tagging and dependency parsing work, consider the tweet “Young people deserve a better future!”. `SpaCy` processes the text and returns the part-of-speech for each word (“young” is an adjective, “people” a noun, and “deserve” a verb) along with the dependency relations (“people” as the nominal subject of the verb “deserve” and “future” as the direct object), as shown in Figure 1. In this first step, we, therefore, extract the subjects and direct objects for

Figure 1: SpaCy POS tagging and dependency parsing.

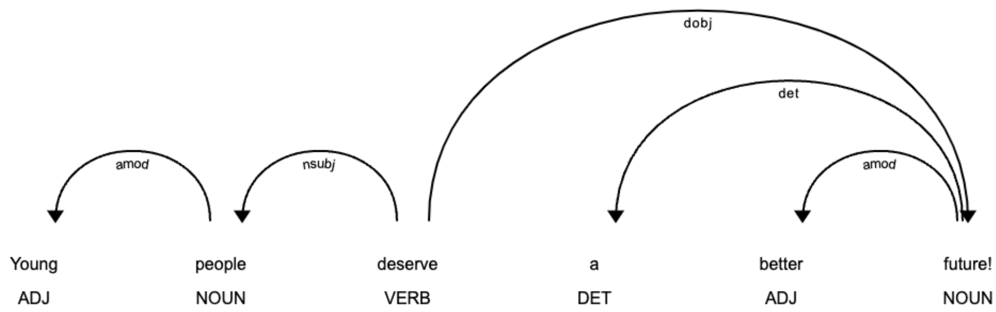

Notes: Visual representation of SpaCy output for parts-of-speech tagging and dependency parsing.

each tweet, which in this example are "young people" and "a better future."[3]

Step 2: Group-specific vectors. To establish a group-specific benchmark for determining whether the subjects/recipients identified in the previous step reflect a group appeal, we let n-grams emerge based on their frequency of co-occurrence (e.g., bi-grams like "young man" or tri-grams like "people of color"). We then train a word embedding model on the millions of tweets we collected, enabling the model to understand the relationships between words as used by candidates on Twitter.[4]

Next, we create group-specific lists of words by selecting the 50 words most similar to predefined seed words that represent the group of interest. For example, for the group "young people", we extract the 50 words semantically closest to the average embedding of `young_people` and `young`.The final list of words for each group is displayed in Section C.1.

Step 3: Semantic similarity. After extracting subjects/recipients and identifying group-specific n-grams, we need a decision rule to classify whether a tweet appeals to a group or not. Consistent with our conceptualization of group appeal, we classify a tweet as appealing to a group if its subject or recipient is semantically similar (i.e., akin to a synonym) to the group vectors.

To represent subjects/recipients and groups as embeddings, we utilize the sentence transformers framework to compute embeddings for both subjects/recipients and the group-specific words,

[3]In isolating subjects and direct objects, we retain the dependent words such as "young" and "better". This is done automatically in `SpaCy`. In Appendix B1 we explain further how `SpaCy` works, and we present other examples of how "young people" can be used in tweets as subjects or recipients.

[4]We use the `word2vec` model introduced by Mikolov et al. (2013).

Table 2: Examples of Group Appeals.

Party	Tweet	GAD Output	Group
Dem	To invest in **the next generation** of Americans, we need a President who believes in education and empowers our students to succeed. Free community college, expanding skilled trade programs, and early childhood education are a start. @JoeBiden knows that and will make it happen.	['the next generation', 0.60]	Young
Rep	**Black History** is good. Black Future is better! It's time to stop being victims and start being victorious! Establishing BLACK FUTURE MONTH	['Black History', 0.64]	Black
Dem	**Women** are used to fighting for our equality - and we're not going to stop till we have it. #ERANow #RatifyERA #EqualRightsAmendment	['Women', 0.78]	Women

Notes: Examples of group appeals detected by our measurement strategy. The first two columns report the party affiliation of the candidate posting the tweet and the text of the tweet. The GAD Output column reports the subject/beneficiary whose cosine similarity is above the .60 threshold and the respective group appeal assigned to the tweet in the last column.

estimating how semantically similar they are to each other.[5] Semantic similarity ranges from -1 to 1. We set an arbitrarily high threshold of 0.60, above which the subject/recipient in the tweet is classified as appealing to the specific group.[6]

In Table 2 we display three examples of tweets classified as group appeals, including the extracted subject/recipient and their cosine similarity to the group-specific embedding. These examples illustrate the advantages of our approach. For instance, by using embeddings, we can identify "the next generation" as a distinct but semantically relevant appeal to young people.[7]

Our group appeal detector offers two key advantages over alternative approaches: flexibility and transparency/replicability. By using embeddings to represent our target groups, we can capture various ways in which candidates appeal to these groups. Utilizing transformer models to generate group-specific embeddings allows us to retain the contextual richness and semantic properties of the group. Another advantage of this highly flexible strategy is that the same method can be used to extract candidates' appeals to any other groups by just defining a set of group-specific seed words. Finally, as we show in the following section, while achieving performance comparable to large language models, our approach is free and fully replicable. We provide further validation

[5]We use the pre-trained `all-MiniLM-L6-v2` model, publicly available `https://huggingface.co/`.

[6]In Figure D.1 in the Appendix, we validate the 0.60 threshold and demonstrate that different thresholds yield similar results.

[7]Note that our approach would not classify any reference to "next generation" as a young appeal (e.g., "The Next Generation EU is ..." is not classified as an appeal to the young).

in Section D of the Appendix, where we show high agreement with trained human coders and comparable performance to large language models.

Empirical Results

Descriptive Analysis of Candidates' Group Appeals

Candidates' communication on Twitter is highly heterogeneous, and the content of their tweets varies significantly. To characterize the volume of group appeals, we compare the supply of group appeals to the supply of other policy-based messages. Table 3 shows the percentage of tweets classified as group appeals by party, alongside the percentage of tweets classified into five major policy areas that are particularly salient in U.S. politics: abortion, environment, health care, gun control, and immigration.[8]

Table 3: Percentage of tweets with group appeals or policy content.

	Democrats	Republicans
Group Appeals		
Young	0.92	0.54
Black	0.81	0.22
Women	1.63	0.48
Any	3.36	1.25
Policy Content		
Abortion	0.38	0.50
Healthcare	4.31	1.82
Immigration	1.22	0.80
Gun Control	0.66	0.49
Environment	1.28	0.49

Notes: Percentage of tweets containing a group appeal compared to the percentage of tweets containing policy-specific contents.

The table shows that candidates' appeals to young voters, women, and black voters are as frequent as their mentions of key policy issues. Except for healthcare mentions, which are significantly more frequent than group appeals, the frequency of group appeals is remarkable if compared to how often candidates tweet about major policy topics.

[8] In Section E in the Appendix, we describe the simple approach we used to classify tweets into policy topics based on a combination of locally trained word embeddings and dictionary analysis.

Moreover, our data shows that the percentage of candidates' tweets containing a group appeal increased over time. From 2012 to 2021, the percentage of tweets containing a group appeal increased from 1.5% to 2.9%, peaking at 3.3% in 2019. Figure 2 illustrates this increasing trend. The first panel on the left shows the average number of tweets posted by candidates by party and over time. Each dot is a monthly party-specific average. The dashed lines represent a flexible approximation of the underlying trend. Democratic candidates post an average of 75 tweets per month, whereas Republican candidates post 52. This gap widens over time, especially from 2016, and decreases again in the early 2020s. The remaining three panels report the average monthly number of group appeals by Democratic and Republican candidates. While the supply of group appeals from Democrats and Republicans did not differ significantly between 2012 and 2015, the partisan gap widened in 2016 and peaked in 2020, with Democrats more likely to appeal to young, black, and female voters. Though less pronounced, Republicans also display an increasing supply of group appeals.

District and Candidate Correlates of Candidates' Group Appeals

We now turn to a more systematic test of our predictions about group appeals as a mobilization strategy. Our analysis dataset is at the candidate-district-election year level, where each observation represents the number of tweets containing appeals to group G posted by each candidate during the 100 days preceding the election. We relate this quantity to a battery of district and

Figure 2: Supply of Group Appeals by Party and Over Time.

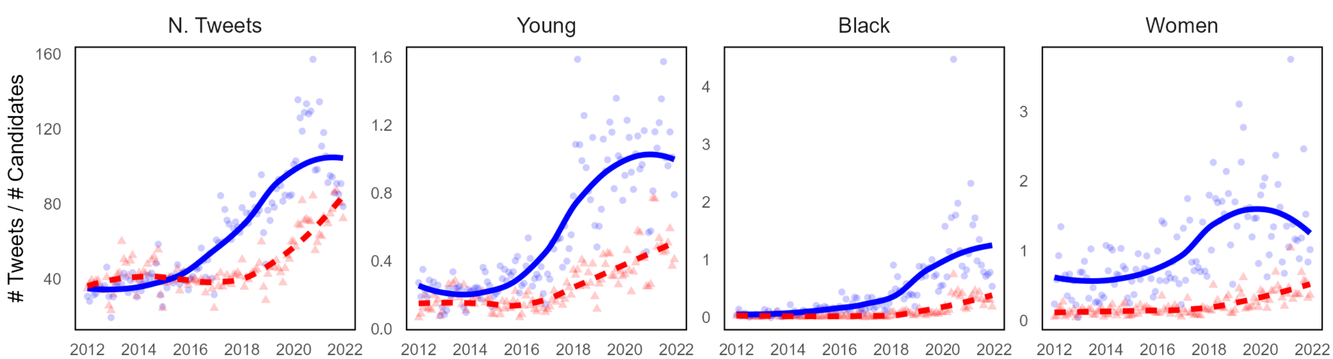

Notes: Average monthly number of tweets and group appeals posted by Republican and Democratic candidates over time.

candidate characteristics. To test *Prediction 1*, we create several variables of interest, namely the share of young (18-24), black, and female voters in the district. To test *Prediction 3*, we include as explanatory variables the candidate's age (continuous) and two dummy variables, one for black candidates and one for female candidates.

The unit of analysis is the total number of appeals to group G (logged) made during the 100 days before the election by candidate c, running in district d, and during election year t. We then estimate the following equation via OLS:

$$\#\text{Appeals}^{G}_{cdt} = \gamma_d + \delta_t + \beta^G \pi^G_{dt} + \zeta' x_{cdt} + \epsilon_{cdt} \tag{1}$$

where π^G is the percentage of individuals in the district belonging to group $G = \{$young, black, women$\}$ (e.g., % young individuals 18-24), x is a vector of candidate- and district-level time-changing covariates (number of tweets posted, party, incumbent status, gender, race, age of the candidate, employment rate, total population, median income, and Gini index). Notably, x includes the age, race, and gender of candidates, which we expect will be strong predictors of candidates' appeals to those groups.

We account for fixed characteristics of districts by estimating district fixed effects γ_d, and we partial out the effect of common shocks with election-year fixed effects δ_t. Because districts' characteristics may be correlated with candidates' decisions to run, district fixed effects enable us to account for the selection effect of candidates into districts. β estimates the percentage increase in the number of appeals resulting from a 1 percentage point increase in the share of group g in the district's population while holding constant same-size increases in the share of groups g' as well as a set of district-level socio-economic covariates and candidate-level characteristics. From the vector of coefficients ζ', we are interested in the effect of candidates' age, race, and gender on the supply of appeals to the group to which the candidate belongs.

We present the results in Table 4. We report two sets of coefficients, on district and candidate predictors. There is a good level of support for the predictions. Note the coefficients along the diagonal of both blocks of predictors. A 1 p.p. increase in the percentage of young people is associated with an increase in the supply of young appeals by 8.4 percent, and an increase in

Table 4: Correlates of Candidates' Supply of Group Appeals.

	DV: Group Appeals (Log #)		
	Young	Black	Women
	(1)	(2)	(3)
District Predictors			
% Young	0.084***	-0.120***	-0.019
	(0.032)	(0.033)	(0.033)
% Black	0.011	0.039***	0.011
	(0.010)	(0.008)	(0.008)
% Women	-0.078	-0.064	0.006
	(0.091)	(0.082)	(0.078)
Candidate Predictors			
Age Candidate	-0.005**	0.001	-0.002
	(0.002)	(0.002)	(0.002)
Black Candidate	0.008	0.376***	0.105
	(0.088)	(0.109)	(0.113)
Female Candidate	0.140***	0.080*	0.602***
	(0.049)	(0.047)	(0.061)
District Covariates	✓	✓	✓
Candidate Covariates	✓	✓	✓
Mean DV	1.6	1.3	2.7
R^2	0.558	0.649	0.686
Observations	2,108	2,108	2,108
District FE	✓	✓	✓
Year FE	✓	✓	✓

Notes: OLS estimates and SE clustered by district in parentheses. The outcome variable is the log-transformed number of tweets appealing to the group in the 100 days leading to the election. District-level covariates include employment rate, total population, median income, and Gini index. Candidate-level covariates include the number of tweets posted, incumbency status, and party. Significance levels: *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

the percentage of African Americans in the district is associated with an increase in the supply of black group appeals by 3.9 percent. Though positive, there is no statistically significant association between the share of women in a district and the candidates' strategic supply of appeal to women. While only changes in the share of young affect candidates' appeals to the group, we find some complementary effects for black appeals, which are particularly prominent in districts with large shares of older African Americans.

We also find a similar diagonal of precisely estimated coefficients for the candidate-level predictors. The age of candidates is negatively associated with appeals to the young; black candidates supply, on average, 37 percent more appeals to African Americans compared to white, Hispanic, and candidates of other races, and female candidates appeal to women 60 percent more than male candidates. However, consistent with the district-level results, we find that gender has a different role compared to age and race. Although the effect of gender is largest for the supply of appeals to women, we find that female candidates are more likely to appeal to any group, including young people and African Americans. This pattern is consistent with gender differences in appealing to traditionally underrepresented societal groups (Kaslovsky and Rogowski, 2022) and with recent work showing that women are more likely to have a local orientation in politics (Ban and Kaslovsky, 2024).

In the Appendix, we report several robustness tests, whose results are broadly consistent with those presented in Table 4. We show that the results are not driven by the log transformation of the dependent variable (Table G.6), are robust to the period analyzed: both before and after the elections (Table G.7), are robust to different operationalizations of the share of young individuals (Table G.9).

Candidates' Group Appeals in Response to Protest Events

We now turn to estimating how the dynamic importance of specific groups affects candidates' group appeals. According to *Proposition 2*, protest events with a clear group connection increase the perceived importance of certain demographic groups within a district, leading to an increase in candidates' appeals to the group. We empirically examine this prediction by studying the effect of two types of protests that occurred in several U.S. cities between 2018 and 2021 and were clearly associated with demographic groups: Fridays for Future climate protests and Black Lives Matter protests.

Fridays for Future. Fridays for Future (FFF) is a global climate strike movement that encourages demonstrations that inspired millions of young people worldwide to advocate for climate protection policies. Fridays for Future protests are registered through the movement's institutional website.[9] Each protest is recorded in the organization's database alongside three key variables: location, country, and number of demonstrators. We obtained the original records directly from

[9]See `https://fridaysforfuture.org/`.

the FFF coordinators. From the first protests in 2018 until the end of 2021, there were 2,717 protests in 1,138 unique cities mapped to 394 congressional districts.

Black Lives Matter. The murder of George Floyd in Minneapolis in May 2020 sparked a large wave of protests across U.S. cities, coordinated within the Black Lives Matter (BLM) movement. We collected data on the location and date of protests and the number of participants from the Armed Conflict Location & Event Data Project (ACLED), which includes information on the coordinates of the events, allowing precise matching of protests to congressional districts. We recovered information for 7,385 protests from January 2020 to December 2021, which in total affected 424 congressional districts.[10]

Even though most districts experienced both types of protests, not all districts experienced them simultaneously. We can leverage the quasi-random occurrence of a protest, which occurred over multiple waves, to estimate the effect of protest exposure to changes in the supply of group appeals using a stacked difference-in-differences design. Given that the vast majority of protests occur outside election periods, we focus on members of Congress, a subset of our candidates who are restricted to those in office on any given day.[11]

Our design treats each wave of protest (i.e., the different dates when protests occurred) as separate sub-experiments, around which we construct a difference-in-differences design using politicians representing districts with and without a protest. The unit of observation is candidate c representing district d over time (days t). We then "stack" all the wave-specific observations of treated and control observations before and after each protest wave to estimate the pooled effect of the protest across all protest waves.[12]

More precisely, let J be a vector of protest dates, and let k identify the days before and after the protest. We observe the same politician k times before and after the protest occurred in her district d. Because k is centered around each wave, negative (positive) values of k identify time periods leading to (following) the protest, and $k = 0$ denotes the day of the protest. Multiple protests can occur in the same district, so to allow for a clean pre-protest period, we focus on

[10]Section F in the Appendix shows the district-level correlates of the occurrence of both types of protests.

[11]Only 0.5% of FFF protests and less 2.3% of BLM protests occur in the month before election day.

[12]By making explicit the comparison group in each wave, this approach solves some of the issues of traditional two-way fixed effects estimators as recently discussed in the econometrics literature by Callaway and Sant'Anna (2021) and de Chaisemartin and D'Haultfœuille (2020).

protests that happen at least two months after the previous one. This results in 915 BLM protests and 1,411 FFF protests. BLM protests occurred over 271 different days (waves), and FFF protests occurred over 115 different days (waves).[13]

To assemble the stacked dataset, we then perform the following two steps for each wave: 1) we subset all districts that had a protest during wave j (treated districts) and those that either had not had a protest yet (not-yet treated) or never had a protest (never treated) as control districts. 2) We match politicians' tweets in treated and control districts. We observe politicians' tweets 15 days before and after the protest. Once we stack each wave-specific dataset, we obtain an analysis dataset of 449,346 observations at the wave-politician-district-day level for BLM protests and 308,432 observations for FFF protests.

For politician c, district d, wave j, and day k around the date of the protest, we estimate the following equation:

$$\#\text{Appeals}^{G}_{cdjk} = \gamma_{cj} + \delta_{kj} + \beta \text{Protest}_{dj} \times \text{post}_k + \epsilon_{cdjk} \tag{2}$$

where $\text{Protest}_{dj} = 1$ if district d experiences a protest in wave j, and 0 otherwise. The variable # Appeals is the (log) number of appeals to the young (for FFF protests) and to black voters (for BLM protests). The variable post_k is defined as $\text{post}_k = 1[k \geq 0]$, taking the value 1 in the post-protest period and 0 before. δ_{kj} are wave-specific time fixed effects, which capture common shocks and wave effects that apply to all districts. Recall that districts can serve in both the treatment and control groups. Consider, for instance, the first district of Alabama, which experienced two BLM protests, one on May 31, 2020, and another on August 29, 2020. This district is used in the control group for waves that occurred before May 31 and in the treatment group for the two waves when a protest occurred in that district. Therefore, we estimate politician (and implicitly district, since politicians only represent one district) fixed effects separately for each protest wave (γ_{cj}). By exploiting within-district and within-candidate variation in exposure to protests, we can control for fixed characteristics of candidates and their self-selection into specific districts as well as their baseline attention to environmental and racial issues. The parameter β captures the impact of

[13] Figure F.2 in the Appendix shows the number of protests in our sample over time.

experiencing a protest relative to control politicians that do not represent a district with a protest during the days near wave j.

For causal inference, we require that districts with and without the protest would have had common trends in the absence of the protest. The main concern for identification is that protests are not randomly assigned. Although the inclusion of candidate fixed effects alleviates this concern, time-varying differences across treated and control districts could still affect the supply of group appeals and bias the estimated effects. However, the high-frequency nature of our data, which tracks politicians' communication on Twitter daily, tempers this concern, allowing us to distinguish changes at the time of protest from slower-moving confounding factors.

Furthermore, as is often the case with difference-in-differences designs where treatment assignment follows geographic boundaries, estimated effects might be biased due to the presence of spill-overs to adjacent districts (see, e.g. Butts, 2023). Large protests in major cities may receive national media coverage and influence the rhetorical strategies of politicians representing other districts. Important for our purposes, ACLED reports the source of information used to code the protest event, which can be international, national, sub-national, regional, or local. This additional information for each protest allows us to restrict the sample of protests to those that did not receive national coverage, thus partially mitigating the threat to identification posed by the presence of spill-overs. For BLM protests alone, we report two sets of results, using the total and sub-national/local samples of protests, respectively.

Table 5 displays the results. Both types of protests had a small but positive effect on the supply of group appeals by candidates in districts affected by the protests. The effect sizes range between an increase of 0.4-0.5% for both the supply of black and young appeals. When looking at protests without national media coverage, the effects of BLM protests on candidates' supply of black appeals increases and gains precision.

As robustness tests, we demonstrate how the results change when the number of days before and after the protest included in the analysis is restricted or expanded (Figure H.4). For BLM protests, we obtain similar results to those reported in Table 5 for larger time windows (from 20 to 30 days), whereas the effects become noisier and shrink to zero when the analysis is limited to 5

Table 5: Protests and Group Appeals.

	DV: Group Appeals (Log #)		
Group Appeal:	Black		Young
Protest:	BLM (All) (1)	BLM (Local) (2)	FFF (3)
Protest × Post	0.0053* (0.0029)	0.0062** (0.0030)	0.0043** (0.0020)
Mean DV	0.044	0.045	0.026
R^2	0.1966	0.2001	0.0903
Observations	449,346	429,806	308,432
Candidate-Protest wave FE	✓	✓	✓
Day since protest-Protest wave FE	✓	✓	✓

Notes: OLS estimates of the effect of BLM and FFF protests on politicians' supply of black and young appeals. SE clustered by district-protest wave in parentheses. The outcome variable is the log-transformed number of tweets appealing to the group in the 15 days before and after the protest. Significance levels: *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

and 10 days before and after the protest. The opposite holds for FFF protests, with larger effects for short windows and smaller effects for larger windows. Finally, we show that the results are robust to using the non-transformed number of group appeals (count) and a dichotomous measure as outcome variables (see Table H.10).

The Mobilization Effect of Candidates' Group Appeals

In a final step, we explore the effectiveness of group appeals in increasing turnout among group members. In keeping with our argument on strategic mobilization, we might expect group members to reward candidates who appeal to their group. Yet, thus far, we have provided no evidence that this is the case. Suppose young voters identifying with the Democratic Party must decide whether to turn out to vote for their party's candidate in the U.S. House elections. Are they more likely to do so if their party's candidate frequently appeals to the young?

Studying the mobilization effects of group appeals outside an experimental setting is challenging. Our high-frequency data, however, offers a unique opportunity to shed light on the effectiveness of group appeals in an observational and real-world context. Our approach matches candidates' daily Twitter posts with a large U.S. public opinion survey — Nationscape (Tausanovitch and

Vavreck, 2021) – and compares respondents' turnout intentions immediately – before and after a candidate makes a group appeal.

Nationscape is one of the largest surveys of the U.S. general public, encompassing approximately 500,000 interviews (on average, 575/day) conducted in nearly every mid-sized U.S. city in the lead-up to and months following the 2020 election. The survey spans from July 2019 to November 2020. The survey includes three categories of variables that enable us to assess the mobilization effect of group appeals. First, demographic variables allow us to identify respondents' group affiliations. Second, information on congressional districts and respondents' party identifications allows us to link each respondent to their respective candidate. Lastly, using the timestamp of each interview, we can map the day-specific supply of group appeals to respondents' turnout intentions.

With this data, we employ a high-frequency differences-in-differences strategy to examine whether respondents belonging to group G are more likely to report intentions to turn out and vote for their party's candidate to the U.S. House if interviewed in the days immediately following a group appeal posted by their candidate, compared to respondents in the same group interviewed in the days immediately prior to the appeal.

To build the dataset, we first extract the dates when candidates posted tweets containing appeals to group G, removing those dates too close to the previous one to allow for a clean pre-post period. We impose a 10-day period without additional group appeals before including a new date in the dataset. For example, during the period from the beginning of the survey until the 2020 election day, Republican candidate Amanda L. Adkins (KS-03) posted appeals to young people on four different days: July 11, August 25, August 27, and October 11, 2020. From this set of dates, we remove August 27 because it is less than 10 days apart from $t =$ August 25.

Consistent with the previous notation, let J be a set of unique dates/events when candidate c posted an appeal to group G. For every candidate-day of appeal pair (cj), we match respondents from the same district and party that answered the survey 3 days before and after the day of appeal ($k \in [-3, 3]$). The short time window allows us to assume that absent the candidate's appeal to group G, respondents belonging to group G would have answered similarly to respondents not belonging to group G (i.e., parallel trends assumption).[14] Within each district, we classify all respondents from group G (the group targeted by the appeal) as the treatment group, while all other respondents serve as the control group. This ensures that both treatment and control groups are exposed to the same district and candidate, isolating the effect of the appeal itself. Continuing with the example of Amanda L. Adkins, we match a total of 20 respondents who identify as Republicans in district KS-03 and were interviewed three days before and after the days when she posted one or more young appeals. We repeat these steps for each candidate-event cj combination and stack the datasets in order to observe individual turnout intention y_i for respondents belonging to group g, matched with candidate c, k days before and after the event j. Eventually, we obtain three stacked datasets, one for each of our three types of appeals to young, black, and female voters.

We then compare treated (those belonging to group G) to control respondents before and after the candidate posted an appeal to the same group G. To do so, we estimate the following event-study specification:

$$y_{igcjk} = \upsilon_{cgj} + \delta_{jk} + \sum_{\tau=-k}^{k} \beta_\tau D_{igc} \mathbf{1}[\tau = k] + \zeta' x_{icjk} + \epsilon_{igcjk} \tag{3}$$

where y_{igcjk} is the respondent's intention to turn out in support of the candidate of her party. Subscript g represents whether the individual belongs to the same group of the appeal made by candidate c on date j. Respondents are interviewed $k = 3$ days before and after j. υ_{cgj} are candidate-group-event fixed effects, aimed at controlling for the interaction between group-, candidate-, and district-specific factors that may differently influence respondents' turnout intentions. δ_{jk} are event-specific time effects, x_{icjk} is a set of covariates including respondents' age, gender, race, employment status, education level, household income, interest in politics, perception of the economy and the country's direction, as well as the number of tweets posted by the candidate in event-time jk. We also include two dichotomous variables equal to 1 if a FFF or BLM protest occurred in the district of the candidate c and time jk. D_{igc} is an indicator for treated individuals belonging to group $G = g$ (e.g., young respondents). β_τ estimates the mobilization

[14] Note that c captures at the same time the candidate, party, and district.

Figure 3: Mobilization Effects, High-Frequency Event-study Design.

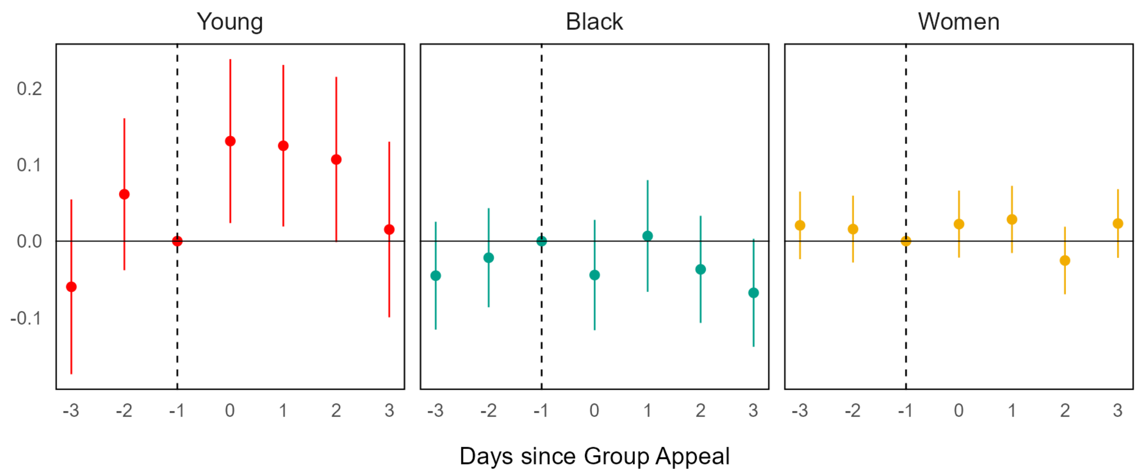

Notes: Event-study OLS estimates and 95% confidence intervals. The outcome is a dummy variable equal to one if respondents belonging to group $G = g$ (displayed as panel titles) report they intend to vote for the candidate of the party they identify with in the next U.S. House elections k days before and after the candidate makes an appeal to group g.

effect τ days relative to when the candidate makes an appeal. ϵ is the error term, and we cluster standard errors at the level of treatment assignment, specifically by candidate, group, and event.

Figure 3 displays the event-study results. Young respondents interviewed in the three days immediately following a young appeal by their party's candidate are more likely to express they intend to turn out to vote for that candidate compared to young respondents interviewed just before the appeal was posted. The effects are large and, averaged across the entire post-appeal period, account for +7.7 percentage points in the probability of turnout intentions. The figure shows a sizable and precisely estimated effect on the day of the appeal and the following day. The effect remains positive, albeit less precise, two days after the appeal, and it diminishes to 0 on the third day. The effect of young appeals is short-lived and – assuming turnout intentions correlate with actual turnout behavior – it is not clear whether group appeals that are not made in close proximity to the election day can have a meaningful impact on turnout.

Conversely, we find no significant effect for women and black voters. One reason why young respondents are more likely to react to group appeals might be the fact that younger people are more likely to be exposed to candidates' communication on social media (ANES data from 2020 shows that 60% of 20-year-olds have visited Twitter in the past year, whereas only 25% of 60-yearolds report doing so).[15] Similarly, young respondents might have less rooted opinions, and their propensity to change (self-reported) behavior might be higher than older people. Additionally, for older groups, including Black and female voters, voting habits may be more persistent over time. Decades of research have shown that voting in one election strongly predicts voting in subsequent elections (Brody and Sniderman, 1977) and that generational turnout patterns tend to persist once established in early adulthood (Plutzer, 2002; Franklin, 2004).[16]

Conclusion

The political mobilization of voters is a crucial aspect of politics and, at least in electoral democracies, the key to electoral success. Scholars and political observers have noted that political competition has increasingly shifted away from a sole focus on the traditional economic left-right to conflicts over identity. The rise of social movements – ranging from feminist and civil rights activism to climate justice efforts – has re-invigorated political conflict on second dimensions,reshaping the nature of political communication and campaign strategies. Mobilizing the voters based on their class belonging is now complemented by, and at times supplanted by, efforts to engage young voters, women, and racial minorities.This diversification of electoral appeals allows candidates to adopt multiple strategic approaches. We argue that candidates use group-based appeals to signal their commitment to specific constituencies, reinforcing their support among key electoral coalitions.

To study group appeals empirically, we introduce a new approach to detect group appeals and apply it to novel high-frequency data on U.S. candidates' campaign messaging. The richness of the data allows us to detect, on a daily basis, when candidates appeal to certain demographic groups, departing from existing approaches that focus on party-level campaign strategies. We find that candidates' appeals are systematically influenced by district demographics, political events, and candidate characteristics, with mobilization effects observed among young voters but not among

[15] See Figure I.7 in the Appendix.

[16] We report difference-in-differences estimates in Table I.11 in the Appendix. Additionally, we show similar results using alternative window sizes (Figure I.5) and we find that the effect of group appeals on turnout intentions is strongest (and precisely estimated) only when including respondents up to 23-25 years (see Figure I.6).

female or Black voters. By analyzing real-world campaign communication rather than relying on party manifestos or hypothetical scenarios typical of experimental settings, this study offers one of the most comprehensive attempts to document the strategic nature of group appeals in the U.S. context.

This study contributes to a broader literature on strategic communication and political targeting by providing novel data and methods. In addition to documenting candidates' responsiveness to structural, contextual, and personal factors, we examine whether group appeals are effective at mobilizing voters. Relying on panel data and a high-frequency design, we find that young voters are responsive to candidate appeals directed at the young, whereas we do not detect a mobilization effect for female and Black voters. Voters aged 18-24 are more likely to self-report that they plan to turn out as a result of the increased supply of appeals to the young by their party's candidate. While previous work has mostly relied on experimental manipulation, this paper is among the few attempts to study political mobilization in a real-world setting, with actual communication by actual candidates.

Finally, this study opens new interesting avenues for future research. First, empirically, our group appeal detector can be easily used by interested researchers to detect appeals to a wide array of demographic and political groups across different communication platforms. Second, theoretically, this study raises new questions about the determinants of heterogeneous mobilization effects – why some groups are more responsive than others to demographic appeals and how demographic group appeals interact as substitutes or complements to more traditional policy pledges? Candidates might strengthen associations between certain groups and specific policies, or alternatively, strategically use demographic appeals to obscure policy ambiguity.

References

Adams, J. (2012). Causes and electoral consequences of party policy shifts in multiparty elections: Theoretical results and empirical evidence. *Annual Review of Political Science 15*, 401–419.

ANES (2021). Anes 2020 time series study full release [dataset and documentation].

Ansolabehere, S. and S. Iyengar (1994). Riding the wave and claiming ownership over issues: The joint effects of advertising and news coverage in campaigns. *Public Opinion Quarterly 58*(3), 335–357.

Ansolabehere, S., S. Iyengar, A. Simon, and N. Valentino (1994). Does attack advertising demobilize the electorate? *The American Political Science Review 88*(4), 829–838.

Baccini, L. and T. Sattler (2024). Austerity, economic vulnerability, and populism. *American Journal of Political Science n/a*(n/a).

Ban, P. and J. Kaslovsky (2024). Local orientation in the u.s. house of representatives. *American Journal of Political Science n/a*(n/a).

Banducci, S. A., T. Donovan, and J. A. Karp (2004). Minority representation, empowerment, and participation. *The Journal of Politics 66*(2), 534–556.

Barney, D. J. (2017). CCES 2016 Supplemental Data.

Barrie, C., A. Palmer, and A. Spirling (2024). Replication for language models: Problems, principles, and best practice for political science.

Bellodi, L., M. Morelli, A. Nicolò, and P. Roberti (2023). The shift to commitment politics and populism: Theory and evidence. *CEPR Discussion Paper (https://cepr.org/publications/dp18338)*(DP18338).

Bobo, L. and F. D. Gilliam (1990). Race, sociopolitical participation, and black empowerment. *American Political Science Review 84*(2), 377–393.

Bobo, L. and F. D. Gilliam Jr (1990). Race, sociopolitical participation, and black empowerment. *American Political Science Review 84*(2), 377–393.

Brody, R. A. and P. M. Sniderman (1977). From life space to polling place: The relevance of personal concerns for voting behavior. *British Journal of Political Science 7*(3), 337–360.

Burns, N., K. L. Schlozman, and S. Verba (1997). The public consequences of private inequality: Family life and citizen participation. *American political science review 91*(2), 373–389.

Butts, K. (2023). Difference-in-differences estimation with spatial spillovers. *arXiv preprint arXiv:2105.03737*.

Callaway, B. and P. H. Sant'Anna (2021). Difference-in-differences with multiple time periods. *Journal of Econometrics 225*(2), 200–230. Themed Issue: Treatment Effect 1.

Carey, J. M. and M. S. Shugart (1995). Incentives to cultivate a personal vote: A rank ordering of electoral formulas. *Electoral studies 14*(4), 417–439.

Cha, J., S. Kuriwaki, and J. M. Snyder (2021). Candidates in American General Elections. *Harvard Dataverse*.

Colantone, I. and P. Stanig (2018). The trade origins of economic nationalism: Import competition and voting behavior in western europe. *American Journal of Political Science 62*(4), 936–953.

Cox, G. (1999). Electoral rules and electoral coordination. *annual Review of Political science 2*(1), 145–161.

Cox, G. W. (2015). Electoral rules, mobilization, and turnout. *Annual Review of Political Science 18*(1), 49–68.

Dal Bo', E., F. Finan, O. Folke, T. Persson, and J. Rickne (2023). Economic and social outsiders but political insiders: Sweden's populist radical right. *The Review of Economic Studies 90*(2), 675–706.

Dalton, R. J. (2013). *Citizen politics: Public opinion and political parties in advanced industrial democracies*. Cq Press.

Dausgaard, C. H. and F. Hjorth (2025). Elite rhetoric and the running tally of party-group linkages. *(https://fghjorth.github.io/workingpapers/runningtally.pdf)*.

Davidson, C. and B. Grofman (1994). *Quiet Revolution in the South: The Impact of the Voting Rights Act, 1965–1990*. Princeton, NJ: Princeton University Press.

de Chaisemartin, C. and X. D'Haultfœuille (2020, September). Two-way fixed effects estimators with heterogeneous treatment effects. *American Economic Review 110*(9), 2964–96.

Dehdari, S. H. (2022). Economic distress and support for radical right parties—evidence from sweden. *Comparative Political Studies 55*(2), 191–221.

Deshpande, P. (2022a). CCES 2018 Supplemental Data.

Deshpande, P. (2022b). CCES 2020 Supplemental Data.

Dolan, K. (1998). Voting for women in the "year of the woman". *American Journal of Political Science 42*(1), 272–293.

Dolinsky, A. O. (2023). Parties' group appeals across time, countries, and communication channels—examining appeals to social groups via the parties' group appeals dataset. *Party Politics 29*(6), 1130–1146.

Druckman, J. N. (2001). On the limits of framing effects: Who can frame? *The Journal of Politics 63*(4), 1041–1066.

Eifert, B., E. Miguel, and D. N. Posner (2010). Political competition and ethnic identification in africa. *American journal of political science 54*(2), 494–510.

Fenno, R. F. et al. (1978). *Home style: House members in their districts*. Little, Brown Boston.

Finseraas, H., O. Heath, P. E. Langsæther, and K. Smeets (2025). How group appeals shape group-party linkages in two political systems. *The Journal of Politics, forthcoming*.

Fraga, B. L., E. G. Juenke, and P. Shah (2021). Candidate Characteristics Cooperative (C3) 2018 Data.

Franklin, M. N. (2004). *Voter turnout and the dynamics of electoral competition in established democracies since 1945*. Cambridge University Press.

Gallego, A. (2007). Inequality in political participation: Contemporary patterns in european countries.

Gause, L. (2022). Revealing issue salience via costly protest: How legislative behavior following protest advantages low-resource groups. *British Journal of Political Science 52*(1), 259–279.

Gillion, D. Q. (2012). Protest and congressional behavior: Assessing racial and ethnic minority protests in the district. *The Journal of Politics 74*(4), 950–962.

Green, D. P. and A. S. Gerber (2019). *Get out the vote: How to increase voter turnout*. Brookings Institution Press.

Haffert, L., T. Palmtag, and D. Schraff (2024). When group appeals backfire: explaining the asymmetric effects of place-based appeals. *British Journal of Political Science 54*(4), 1217–1238.

Hajnal, Z. L. (2006). *Changing white attitudes toward black political leadership*. Cambridge University Press.

Hersh, E. D. and B. F. Schaffner (2013). Targeted campaign appeals and the value of ambiguity. *The Journal of Politics 75*(2), 520–534.

Holbrook, T. M. and S. D. McClurg (2005). The mobilization of core supporters: Campaigns, turnout, and electoral composition in united states presidential elections. *American Journal of Political Science 49*(4), 689–703.

Holman, M. R., M. C. Schneider, and K. Pondel (2015). Gender targeting in political advertisements. *Political Research Quarterly 68*(4), 816–829.

Honnibal, M. and M. Johnson (2015). An improved non-monotonic transition system for dependency parsing. *Conference Proceedings - EMNLP 2015: Conference on Empirical Methods in Natural Language Processing* (September), 1373–1378.

Horn, A., A. Kevins, C. Jensen, and K. Van Kersbergen (2021). Political parties and social groups: New perspectives and data on group and policy appeals. *Party Politics 27*(5), 983–995.

Huber, J. D. and G. B. Powell Jr (1994). Congruence between citizens and policymakers in two visions of liberal democracy. *World Politics 46*(3), 291–326.

Huber, L. M. (2022). Beyond policy: The use of social group appeals in party communication. *Political Communication 39*(3), 293–310.

Huber, L. M., T. M. Meyer, and M. Wagner (2024). Social group appeals in party rhetoric: Effects on policy support and polarization. *The Journal of Politics 86*(4), 1304–1318.

Jacobson, G. C. (2015). How do campaigns matter? *Annual Review of Political Science 18*(1), 31–47.

Kahn, K. F. (1993). Gender differences in campaign messages: The political advertisements of men and women candidates for u. s. senate. *Political Research Quarterly 46*(3), 481–502.

Kaslovsky, J. and J. C. Rogowski (2022). Under the microscope: Gender and accountability in the us congress. *American Political Science Review 116*(2), 516–532.

Key, V. O. (1948). *Southern Politics in State and Nation*. Knoxville: University of Tennessee Press.

Kinder, D. R. (1998). Opinion and action in the realm of politics.

Kinder, D. R. and C. D. Kam (2010). *Us against them: Ethnocentric foundations of American opinion*. University of Chicago Press.

Kitanova, M. (2020). Youth political participation in the eu: evidence from a cross-national analysis. *Journal of Youth Studies 23*(7), 819–836.

Klar, S. (2013). The influence of competing identity primes on political preferences. *The Journal of Politics 75*(4), 1108–1124.

Kousser, J. M. (1974). *The Shaping of Southern Politics: Suffrage Restriction and the Establishment of the One-Party South, 1880–1910*. New Haven, CT: Yale University Press.

Kuklinski, J. H. and N. L. Hurley (1994). On hearing and interpreting political messages: A cautionary tale of citizen cue-taking. *The Journal of Politics 56*(3), 729–751.

la Cour Dabelko, K. and P. S. Herrnson (1997). Women's and men's campaigns for the u. s. house of representatives. *Political Research Quarterly 50*(1), 121–135.

Leighley, J. E. and J. Nagler (2013). *Who votes now?: Demographics, issues, inequality, and turnout in the United States*. Princeton University Press.

Licht, H. and R. Sczepanski (2023, June 20). Who are they talking about? detecting mentions of social groups in political texts with supervised learning. OSF Preprints.

Mazumder, S. (2018). The persistent effect of us civil rights protests on political attitudes. *American Journal of Political Science 62*(4), 922–935.

McGhee, E. and J. Sides (2011). Do campaigns drive partisan turnout? *Political Behavior 33*, 313–333.

Mikolov, T., K. Chen, G. Corrado, and J. Dean (2013). Efficient estimation of word representations in vector space. *1st International Conference on Learning Representations, ICLR 2013 - Workshop Track Proceedings*, 1–12.

Miller, A. H., C. Wlezien, and A. Hildreth (1991). A reference group theory of partisan coalitions. *The Journal of Politics 53*(4), 1134–1149.

Mycock, A. and J. Tonge (2012). The party politics of youth citizenship and democratic engagement. *Parliamentary affairs 65*(1), 138–161.

Nelson, T. E. and D. R. Kinder (1996). Issue frames and group-centrism in american public opinion. *The Journal of Politics 58*(4), 1055–1078.

Norris, P. (2004). Young people & political activism. *Harvard University, John F. Kennedy School of Government.(32p)*.

Paxton, P., S. Kunovich, and M. M. Hughes (2007). Gender in politics. *Annu. Rev. Sociol. 33*, 263–284.

Plutzer, E. (2002). Becoming a habitual voter: Inertia, resources, and growth in young adulthood. *American political science review 96*(1), 41–56.

Porter, R. and S. A. Treul (2024). Evaluating (in)experience in congressional elections. *American Journal of Political Science n/a*(n/a).

Posner, D. N. (2004). The political salience of cultural difference: Why chewas and tumbukas are allies in zambia and adversaries in malawi. *American political science review 98*(4), 529–545.

Powell, G. B. and G. B. Powell Jr (2000). *Elections as instruments of democracy: Majoritarian and proportional visions*. Yale University Press.

Robison, J., R. Stubager, M. Thau, and J. Tilley (2021). Does class-based campaigning work? how working class appeals attract and polarize voters. *Comparative Political Studies 54*(5), 723–752.

Rosenstone, S. J. and J. M. Hansen (2003). *Mobilization, Participation and Democracy in America*. Longman Press.

Schaffner, B. F. (2005). Priming gender: Campaigning on women's issues in u.s. senate elections. *American Journal of Political Science 49*(4), 803–817.

Tausanovitch, C. and L. Vavreck (2021). Democracy Fund + UCLA Nationscape Project, October 10-17, 2019 (version 20211215).

Tertytchnaya, K. and T. Lankina (2020). Electoral protests and political attitudes under electoral authoritarianism. *The Journal of Politics 82*(1), 285–299.

Thau, M. (2018). The demobilization of class politics in denmark: The social democratic party's group-based appeals 1961–2004. *World Political Science 14*(2), 169–188.

Thau, M. (2019). How political parties use group-based appeals: Evidence from britain 1964–2015. *Political Studies 67*(1), 63–82.

Thau, M. (2021). The social divisions of politics: How parties' group-based appeals influence social group differences in vote choice. *The Journal of Politics 83*(2), 675–688.

Thau, M. (2024). The group appeal strategy: Beyond the policy perspective on party electoral success. *Political Studies 72*(4), 1653–1674.

Valentim, A. (2023). Repeated exposure and protest outcomes: How fridays for future protests influenced voters. *Working Paper: London School of Economics*.

Wasow, O. (2020). Agenda seeding: How 1960s black protests moved elites, public opinion and voting. *American Political Science Review 114*(3), 638–659.

Weiss, J. (2020). What is youth political participation? literature review on youth political participation and political attitudes. *Frontiers in Political Science 2*, 1.

Wouters, R. and S. Walgrave (2017). Demonstrating power: How protest persuades political representatives. *American Sociological Review 82*(2), 361–383.

Zhuravskaya, E., M. Petrova, and R. Enikolopov (2020). Political effects of the internet and social media. *Annual review of economics 12*(1), 415–438.

Appendix

A Obtaining Twitter Handles of Candidates

We collect the Twitter handle(s) of candidates from several sources:

- **Ballotpedia**, an online repository on U.S. elections and candidates. For instance, for the 2020 elections, we accessed candidates' directory at the following link: `https://ballotpedia.org/List_of_congressional_candidates_in_the_2020_elections`. Then, for each candidate, we extract the link to their social media accounts.

- the **House of Representatives Press Gallery**: Twitter handle for members of the 117th Congress available at `https://pressgallery.house.gov/member-data/members-official-twitter-`

- **Poliwoops**, an organization that tracks deleted tweets by public officials and maintains a list of active Twitter accounts: Dataset available at `https://www.propublica.org/datastore/dataset/politicians-tracked-by-politwoops`.

- For the candidates not matched in these sources, we program a scraper to directly obtain the Twitter account through the search engine of the official Twitter website and we complete the data collection through several manual searches.

B Measurement Strategy

B1 Dependency Parsing Examples

We consider the following dependency relations as subjects or recipients:

1. Young people as a nominal subject (*nsubj*): Young people deserve a better future. (see Figure 1)

2. Young people as a passive nominal subject (*nsubjpass*): Young people are treated without respect.

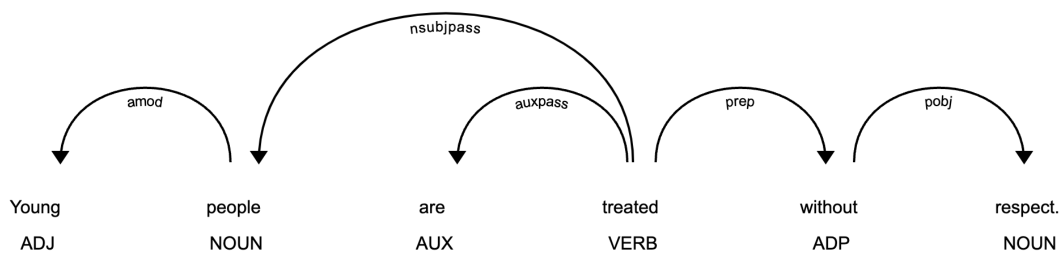

3. Young people as a direct object (*dobj*): I will protect young people.

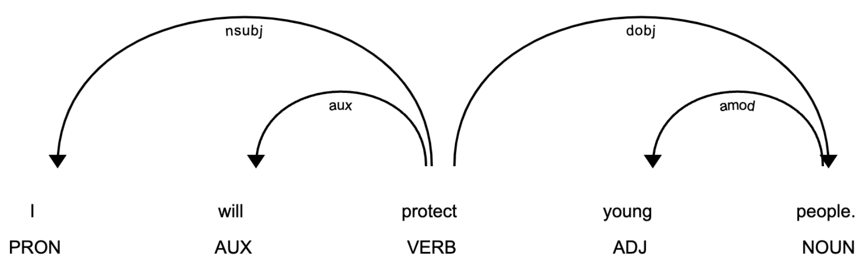

4. Young people as an object of preposition (*pobj*): I will fight for young people.

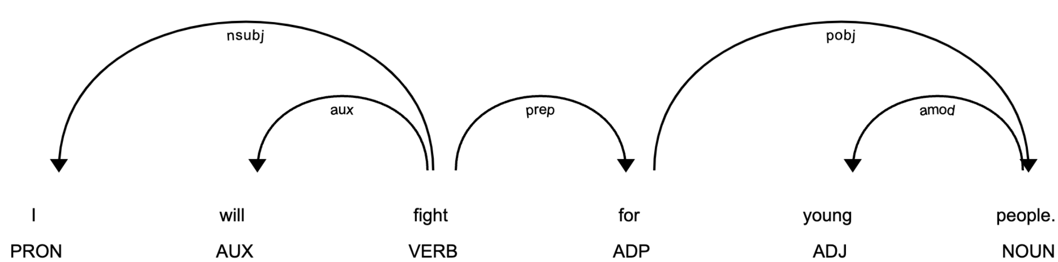

C Group-specific Lists of Words

We create group-specific lists of words by selecting the 50 words most similar to predefined seed words that represent the group of interest. For the group "young people", we extract the 50 words semantically closest to the average embedding of `young_people` and `young`. For women, we use `women` and `female`, and for Black voters, we use `black` and `africanamerican`. We remove false positives, i.e., words that do not accurately represent the group (e.g., removing `educator` but keeping `student`).

Table C.1: Group-specific Lists of Words.

Group	Words
Young	['young_people', 'young_woman', 'student', 'young_adult', 'younger_generation', 'young_leader', 'college_student', 'teen', 'young_girl', 'millennials', 'young_men', 'teenager', 'kid', 'talented_student', 'next_generation', 'high_schoolers', 'young_artist', 'yearolds', 'child', 'talented_young', 'young_man']
Women	['woman', 'female', 'black_woman', 'woman_of_color', 'latina', 'woc', 'young_woman', 'latinas', 'woman_and_girl', 'wmn', 'womengirls', 'womenofcolor', 'indigenous_woman', 'girl', 'womenvets', 'woman_in_uniform', 'servicewomen']
Black	['black', 'africanamerican', 'african_american', 'blk', 'black_and_latino', 'black_and_brown', 'afam', 'africanamericans', 'blackbrown', 'nonwhite', 'poc', 'black_men', 'black_woman', 'people_of_color', 'africanamerican_woman']

Notes: 50 n-grams with the largest cosine similarity between the group-specific seed words. False positives – i.e., n-grams with large cosine similarity without clear connection with the demographic group – have been removed. The average of the vector representation of these words constitutes the group embedding.

D Measurement Validation

To validate our group appeal detector, we build a validation set and compare the outcomes of our proposed method with those produced by a trained human annotator, which we consider to be the most reliable way to identify group appeals for small sets of data. Furthermore, to benchmark the performance of our method to alternative techniques, we assess how large language models perform, comparing the answers given by GPT with those of the human annotator.

Given that group appeals represent a small fraction of our sample of tweets, we build a stratified dataset that artificially increases the difficulty of our measurement strategy in detecting group appeals. To do so, for each of the three groups we study, we extract 500 tweets (1,500 tweets in total) matching the following criteria: 150 random tweets mentioning a group-related keyword (e.g., tweets mentioning the word "young"), 150 random tweets mentioning a group-related keyword but *not* classified as group appeals, 150 random tweets classified as group appeals, and 50 fully random tweets.

We begin by providing detailed instructions to the human coder, including a clear definition of group appeals with examples. The coder is then asked to determine if a tweet contains an appeal to each one of the three groups. The same instructions and examples are prompted to OpenAI's model via API (use use GPT 4o model). For each group, we compare the human coder to both our group appeal detector and GPT.

Table D.2 shows the performance metrics. As per standard practice, we present accuracy, precision, recall, and F1 from comparing the human coder to both the group appeal detector (columns from 1 to 4) and GPT (columns from 5 to 8). We observe a very high level of agreement between our proposed measure and the human coder, with accuracy and F1-score above 0.85. Additionally, there are minimal differences between the performance metrics of our method and those of GPT. Furthermore, Table D.3, we show that the group appeal detector performs well across strata, and Figure D.1 in the Appendix, we validate the 0.60 threshold and demonstrate that different thresholds yield similar results.

Table D.2: Validation Test: Performance Metrics.

	Group Appeal Detector				GPT 4o			
Group	Accuracy	Precision	Recall	F1	Accuracy	Precision	Recall	F1
Young	0.862	0.895	0.862	0.876	0.853	0.932	0.853	0.877
Black	0.884	0.907	0.884	0.893	0.814	0.926	0.814	0.848
Women	0.853	0.883	0.853	0.864	0.788	0.917	0.788	0.823

Notes: Performance metrics of validation test. Columns 2-4 display accuracy (share of agreement), precision (positive predictive value), recall (true positive rate), and F1 score (the harmonic mean of precision and recall) of the proposed group appeal detector and GPT, using the human coder's classifications as a benchmark.

Table D.3: Performance Metrics. Accuracy by stratum.

	Accuracy			
Stratum	Young	Black	Women	Sample Description
Human Coder - GAD				
1	0.980	0.993	0.973	Fully random.
2	0.884	0.896	0.860	Group-specific words mentioned.
3	0.729	0.784	0.753	Group-specific words mentioned and appeal detected.
4	0.933	0.936	0.904	Group-specific words mentioned but no appeal detected.
Human Coder - GPT				
1	0.980	0.960	0.967	Fully random.
2	0.838	0.818	0.762	Group-specific words mentioned.
3	0.858	0.800	0.780	Group-specific words mentioned and appeal detected.
4	0.820	0.776	0.762	Group-specific words mentioned but no appeal detected.

Notes: Performance metrics of the Group Appeal Detector and GPT 4o.

Figure D.1: Validation Test: Performance Metrics with Different Thresholds.

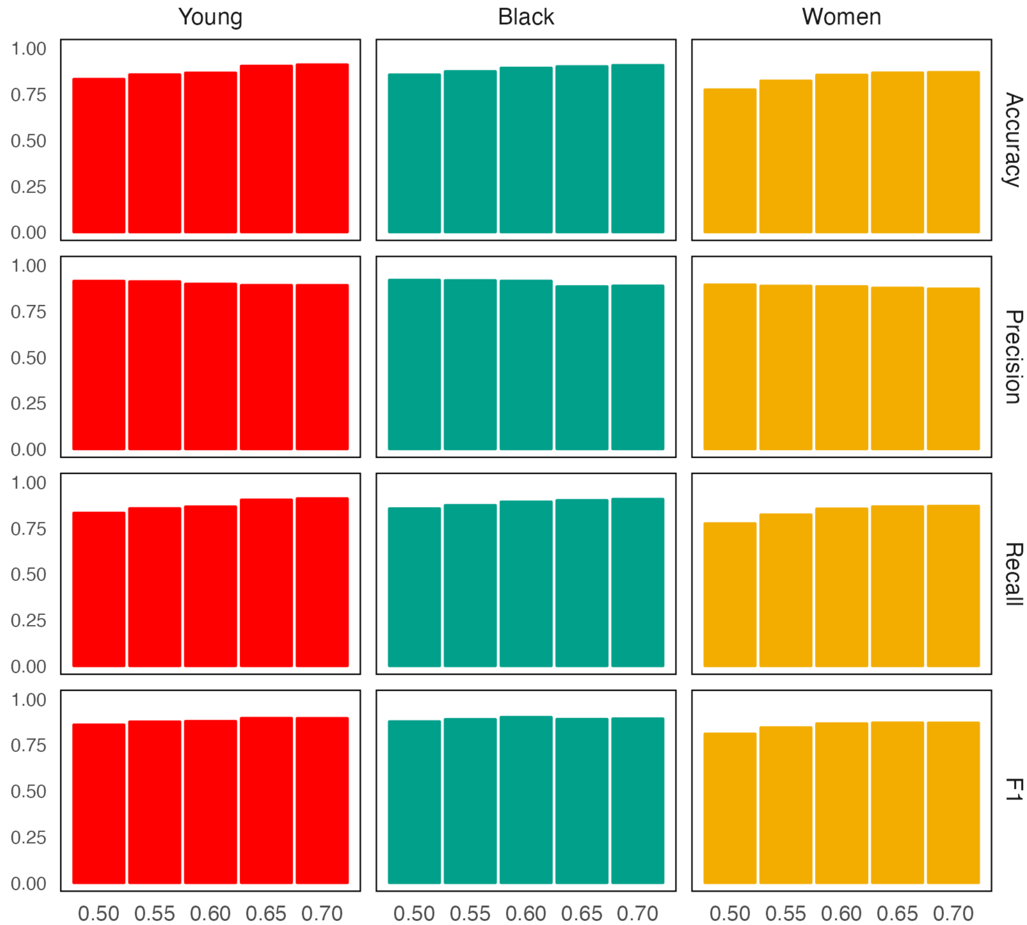

Notes: Performance metrics of validation test using different cutoff thresholds to classify a tweet as group appeal. Each bar represents the accuracy (share of agreement) and F1 score (the harmonic mean of the true positive and true negative rates) between the proposed group appeal detector (GAD) using different thresholds and the GPT models, using as a benchmark the classification choices of the human coder.

E Comparing the Frequency of Group Appeals and Policy Content

In this section, we provide comparative evidence on the frequency of group appeals and policy content on Twitter. We select five salient policy areas in U.S. politics, namely abortion, environmental protection, gun control, health care, and immigration. We combine a dictionary approach with word embeddings to classify tweets into topics.

First, we build topic-specific dictionaries. We extract the 20 words whose vector representation is most similar to the topic label (e.g., immigration) with `word2vec` (Mikolov et al., 2013), an unsupervised algorithm that learns fixed-length feature representations from how often words co-occur with one another, with the assumption the meaning of a word is given by the company it keeps.

We pre-process the 5.9 million tweets posted by Congressional candidates between 2012 and 2021. We lowercase the text, we remove the Twitter accounts tagged in the text (e.g., @NAME), we keep hashtags (#) because they generally convey important information, we remove punctuation, and we convert all nouns to their singular form (e.g., houses → house). We then allow bi-grams to emerge based on how often two words occur next to one another, imposing a minimum count of 200. This step allows for words like ”gun_control” or ”health_care” to be considered single words when training the model. We finally train the `word2vec` model on the pre-processed corpus of tweets with the `gensim` Python library, estimating 200-dimensional vectors, excluding words appearing less than 10 times, and setting a window size (where to compute word co-occurrences) to 4.

Once we have a word embedding for each word used at least ten times in the corpus, we extract the 20 words with the largest cosine similarity to the topic label. We manually remove false positives to ensure that each word is used almost exclusively in the context of the topic (e.g., we remove the word “government” from the military-specific dictionary, for it can be used in many different contexts without referring to military issues) and we assemble the eight topic-specific dictionaries reported in Table E.4 below. We consider tweet i to be about topic j, with $J = j_1, ..., j_5$, if tweet i contains at least one of the words in the dictionary of topic j.

Table E.4: Topic specific dictionaries.

Topic	Dictionary
Gun Control	gun, firearm, handgun, assault_weapon, weapon_of_war, bump_stock, rifle, weapon, assault_rifle, domestic_abuser, silencer, shotgun, pistol, assaultweapons, weapon_ban, gunfreezones, semiautomatic, concealed_carry
Immigration	immigration, immigrant, asylum_seeker, undocumented_immigrant, refugee, asylumseekers, immigrant_and_refugee, legal_immigration, undocumented, immig, immigration_policy, legal_immigrant, illegal_immigration, migrant, deportation, imm, family_reunification, familybased, assimilation, illegalimmigration, sanctuarycities, chain_migration
Abortion	abortion, lateterm_abortion, infanticide, birth_control, family_planning, abortion_clinic, abortionist, legal_abortion, abortion_provider, unborn_baby, abortion_ban, contraceptive, planned_parenthood, reproductive_health, contraception, unborn_child, title_x, born_alive
Environment	environment, our_environment, ecosystem, our_planet, enviro, air_and_water, env, clean_air, airwater, climate, quality_of_life, cleanair, natural_resource, planet, biodiversity, wildlife, energy_sector, waterway, landscape
Health Care	healthcare, health_care, affordable_healthcare, hc, access_to_healthcare, quality_healthcare, heathcare, healthcare_coverage, health_coverage, health, health_insurance, reproductive_healthcare, insurance_coverage, retirement_security, hcare, child_care, universal_healthcare, preventative_care, quality_care, medicare

Notes: Topic specific dictionaries derived from the `word2vec` model based on the largest cosine similarity between words and the topic label.

F District-level Correlates of Protest Occurrence

Figure F.2: Number of protests over time.

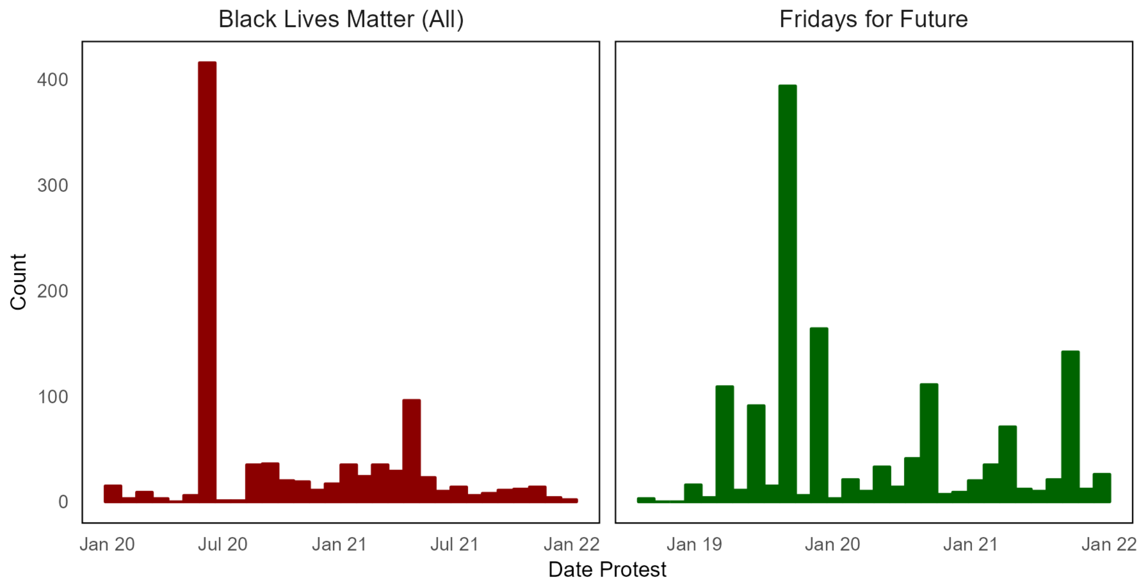

Notes: Number of protests by protest date. The sample includes protests occurring at 60-day intervals, selected from each district.

Figure F.3 shows the relationship between district characteristics and the number of protests occurring over time for FFF (in yellow) and BLM protests (in gray). The figure shows the estimated coefficients of a regression relating the number of protests in any given district-year observation (mean BLM = 17; mean FFF = 6.0) to a set of district characteristics and a battery of state and year dummies.[17] Holding constant differences across states and controlling for common shocks, we find that the number of FFF protests is correlated with income inequalities in the district. BLM protests are more likely to occur in younger districts with a larger share of African Americans, greater employment rates, and higher economic inequalities. Table F.5 shows full regression results using both the number of protests and the logged number of protests as outcome variables.

[17]Our datasets comprise the 435 districts $\times$ 2 years (2020-21) for BLM and 4 years (2018-2021) for FFF.

Figure F.3: District Correlates of Number of Protests.

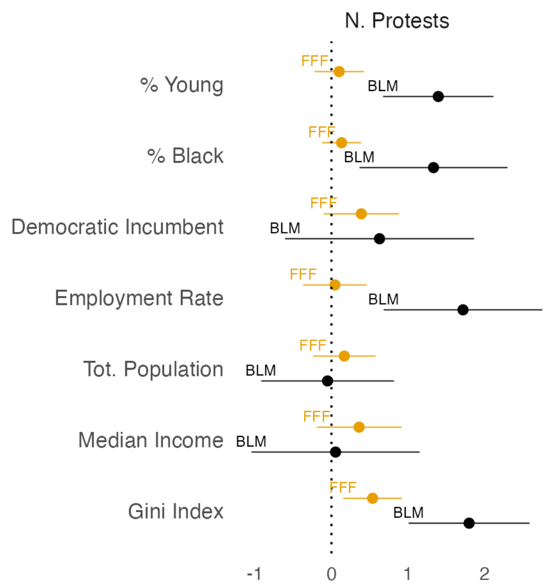

Notes: OLS estimates and 95% confidence intervals of the effect of district characteristics on the number of protests in the district. Estimation includes state and year fixed effects with SE clustered by state.

Table F.5: District Characteristics and Protest Events.

	DV: Number of Protests			
Protest:	FFF		BLM	
	N	Log(1+N)	N	Log(1+N)
	(1)	(2)	(3)	(4)
% Young	0.10	0.03	1.39***	0.14***
	(0.16)	(0.03)	(0.37)	(0.04)
% Black	0.13	0.02	1.33***	0.16***
	(0.13)	(0.03)	(0.42)	(0.05)
Democratic Incumbent	0.39	0.10**	0.63	0.00
	(0.25)	(0.05)	(0.64)	(0.09)
Employment Rate	0.05	0.00	1.72***	0.22***
	(0.21)	(0.04)	(0.49)	(0.06)
Total Population	0.17	0.07	-0.05	-0.03
	(0.21)	(0.05)	(0.39)	(0.05)
Median Income	0.36	0.05	0.05	-0.02
	(0.28)	(0.03)	(0.41)	(0.04)
Gini Index	0.54***	0.11***	1.80***	0.15**
	(0.20)	(0.03)	(0.41)	(0.06)
R^2	0.39	0.63	0.52	0.67
Observations	1,740	1,740	870	870
State FE	✓	✓	✓	✓
Year FE	✓	✓	✓	✓

Notes: OLS estimates of the effect of district characteristics on the number of protests and the average number of demonstrators. SE clustered by state in parentheses. For FFF protests, the dataset is a panel of district-year observations for the years 2018-2021. For BLM protests, the dataset is a panel of district-year observations for the years 2020-2021. All predictors are standardized. Significance levels: *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

G Robustness Tests

Table G.6: Correlates of Candidates' Supply of Group Appeals: Alternative Transformation of Outcome Variable.

	DV: Group Appeals					
	Dichotomous			Count		
	Young	Black	Women	Young	Black	Women
	(1)	(2)	(3)	(4)	(5)	(6)
District Predictors						
% Young	0.066***	-0.049***	0.027	0.123	-0.652***	-0.343
	(0.023)	(0.018)	(0.021)	(0.154)	(0.232)	(0.272)
% Black	0.015**	0.011**	-0.010*	0.001	0.255***	0.132**
	(0.007)	(0.005)	(0.006)	(0.038)	(0.056)	(0.059)
% Women	-0.150**	0.056	0.087*	0.038	-0.804	-0.821
	(0.064)	(0.046)	(0.051)	(0.364)	(0.570)	(0.713)
Candidate Predictors						
Age Candidate	-0.004***	0.000	-0.002	-0.018**	0.001	0.001
	(0.001)	(0.001)	(0.001)	(0.008)	(0.011)	(0.014)
Black Candidate	-0.042	0.212***	0.061	0.247	1.954**	0.154
	(0.060)	(0.054)	(0.072)	(0.425)	(0.839)	(0.702)
Female Candidate	0.064	0.049	0.195***	0.456**	0.402	3.039***
	(0.039)	(0.031)	(0.037)	(0.204)	(0.323)	(0.397)
District Covariates	✓	✓	✓	✓	✓	✓
Candidate Covariates	✓	✓	✓	✓	✓	✓
Mean DV	0.50	0.30	0.60	1.60	1.30	2.70
R^2	0.417	0.538	0.502	0.512	0.553	0.622
Observations	2,108	2,108	2,108	2,108	2,108	2,108
District FE	✓	✓	✓	✓	✓	✓
Year FE	✓	✓	✓	✓	✓	✓

Notes: Columns 1-3 report OLS estimates using as outcome variable a dummy equal to 1 if the number of tweets appealing to the group in the 100 days leading to the election is greater than 0, and 0 otherwise. Columns 4-6 report OLS estimates using as outcome variable the number of tweets appealing to the group in the 100 days leading to the election. SE clustered by district in parentheses. Percentages of young, black, and women are standardized. District-level covariates include: employment rate, total population, median income, and Gini index. Candidate-level covariates include the number of tweets posted, incumbency status, and party. Significance levels: *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

Table G.7: Correlates of Candidates' Supply of Group Appeals: Alternative Time Periods.

	DV: Group Appeals (Log #)					
Time Period:	Before Election			After Election		
	Young	Black	Women	Young	Black	Women
	(1)	(2)	(3)	(4)	(5)	(6)
District Predictors						
% Young	0.073**	-0.131***	-0.021	0.025	-0.080**	-0.016
	(0.034)	(0.045)	(0.038)	(0.030)	(0.039)	(0.027)
% Black	0.009	0.042***	-0.006	0.009	0.027**	0.005
	(0.011)	(0.011)	(0.011)	(0.009)	(0.011)	(0.008)
% Women	-0.201*	-0.047	0.035	-0.126	-0.050	-0.039
	(0.116)	(0.100)	(0.102)	(0.104)	(0.070)	(0.096)
Candidate Predictors						
Age Candidate	-0.006***	0.002	0.000	-0.006***	-0.001	0.000
	(0.002)	(0.002)	(0.002)	(0.001)	(0.002)	(0.002)
Black Candidate	0.060	0.630***	0.273**	0.095	0.407***	0.056
	(0.090)	(0.112)	(0.120)	(0.082)	(0.118)	(0.087)
Female Candidate	0.162**	0.095	0.787***	0.158***	0.121**	0.551***
	(0.067)	(0.060)	(0.070)	(0.045)	(0.051)	(0.053)
District Covariates	✓	✓	✓	✓	✓	✓
Candidate Covariates	✓	✓	✓	✓	✓	✓
Mean DV	4.1	3.6	6.5	2.0	2.3	3.0
R^2	0.722	0.766	0.787	0.768	0.685	0.783
Observations	2,124	2,124	2,124	4,014	4,014	4,014
District FE	✓	✓	✓	✓	✓	✓
Year FE	✓	✓	✓	✓	✓	✓

Notes: OLS estimates and SE clustered by district in parentheses. The outcome variable is the log-transformed number of tweets appealing to the group in the year of the election (Columns 1-3) and in the period following the elections (from the day immediately following the election until the end of the following year (e.g., for the 2012 election, the period analyzed is from November 7th, 2012 until the end of 2013)). Percentages of young, black, and women are standardized. District-level covariates include: employment rate, total population, median income, and Gini index. Candidate-level covariates include the number of tweets posted, incumbency status, and party. Significance levels: *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

Table G.8: Correlates of Candidates' Supply of Group Appeals: Omitting Candidate Predictors.

	DV: Group Appeals (Log #)		
	Young	Black	Women
	(1)	(2)	(3)
District Predictors			
% Young	0.083***	-0.102***	-0.011
	(0.032)	(0.030)	(0.038)
% Black	0.010	0.037***	0.002
	(0.009)	(0.009)	(0.009)
% Women	-0.024	-0.037	-0.005
	(0.082)	(0.074)	(0.089)
District Covariates	✓	✓	✓
Candidate Covariates	✓	✓	✓
Mean DV	1.5	1.1	2.7
R^2	0.490	0.553	0.556
Observations	2,829	2,829	2,829
District FE	✓	✓	✓
Year FE	✓	✓	✓

Notes: OLS estimates and SE clustered by district in parentheses. The outcome variable is the log-transformed number of tweets appealing to the group in the 100 days leading to the election. Candidate predictors omitted (to keep challenger candidates with no available demographic data). Percentages of young, black, and women are standardized. District-level covariates include: employment rate, total population, median income, and Gini index. Candidate-level covariates include the number of tweets posted and party. Significance levels: *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

Table G.9: Correlates of Candidates' Supply of Group Appeals: Alternative Characterization of Young Voters (18-34).

	DV: Group Appeals (Log #)		
	Young	Black	Women
	(1)	(2)	(3)
District Predictors			
% Young	0.084***	-0.120***	-0.019
	(0.032)	(0.033)	(0.033)
% Black	0.011	0.039***	0.011
	(0.010)	(0.008)	(0.008)
% Women	-0.078	-0.064	0.006
	(0.091)	(0.082)	(0.078)
Candidate Predictors			
Age Candidate	-0.005**	0.001	-0.002
	(0.002)	(0.002)	(0.002)
Black Candidate	0.008	0.376***	0.105
	(0.088)	(0.109)	(0.113)
Female Candidate	0.140***	0.080*	0.602***
	(0.049)	(0.047)	(0.061)
District Covariates	✓	✓	✓
Candidate Covariates	✓	✓	✓
Mean DV	1.6	1.3	2.7
R^2	0.558	0.649	0.686
Observations	2,108	2,108	2,108
District FE	✓	✓	✓
Year FE	✓	✓	✓

Notes: OLS estimates and SE clustered by district in parentheses. The outcome variable is the log-transformed number of tweets appealing to the group in the 100 days leading to the election. Percentages of young (18-34), black, and women are standardized. District-level covariates include: employment rate, total population, median income, and Gini index. Candidate-level covariates include the number of tweets posted, incumbency status, and party. Significance levels: *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

H Protest Event Study: Robustness Tests

Figure H.4: Protests and Group Appeals: Alternative Window Size.

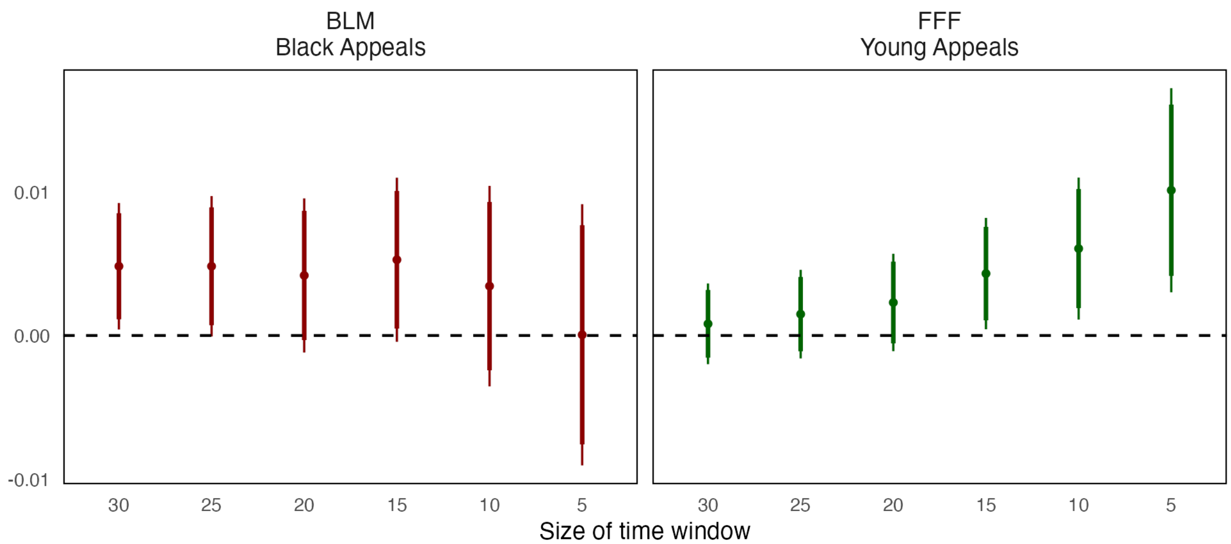

Notes: OLS estimates with 90% and 95% confidence intervals of the effect of BLM and FFF protests on politicians' supply of black and young appeals. SE clustered by district-protest. The horizontal axis shows the number of days before and after used for estimation (e.g., 15 days before and after the protest). The outcome variable is the log-transformed number of tweets appealing to the group.

Table H.10: Protests and Group Appeals: Poisson Estimator.

	DV: Group Appeals			
	Count		Dummy	
Protest:	BLM	FFF	BLM	FFF
Group Appeal:	Black	Young	Black	Young
	(1)	(2)	(3)	(4)
Protest × Post	0.009*	0.007**	0.006*	0.005**
	(0.005)	(0.003)	(0.003)	(0.003)
Mean DV	0.044	0.026	0.036	0.024
R^2	0.188	0.090	0.190	0.089
Observations	449,346	308,432	449,346	308,432
Candidate-Protest wave FE	✓	✓	✓	✓
Day since protest-Protest wave FE	✓	✓	✓	✓

Notes: OLS estimates of the effect of BLM and FFF protests on politicians' supply of black and young appeals. SE clustered by district-protest wave in parentheses. The outcome variable is the log-transformed number of tweets appealing to the group in the 15 days before and after the protest. Significance levels: *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

I Mobilization Effects: Robustness Tests

Table I.11: Candidates' Appeals and Voters' Mobilization, Difference-in-Differences Results.

	DV: Turnout for Own Party's Caniddate					
	Young		Black		Women	
	(1)	(2)	(3)	(4)	(5)	(6)
Post × Young	0.060**	0.090***				
	(0.030)	(0.030)				
Post × Black			-0.017	-0.014		
			(0.020)	(0.021)		
Post × Women					-0.006	0.000
					(0.012)	(0.012)
Covariates		✓		✓		✓
Mean DV	0.89	0.88	0.90	0.89	0.89	0.87
R^2	0.425	0.474	0.432	0.501	0.399	0.454
Observations	20,856	19,176	15,248	14,059	21,374	19,701
Candidate-Date of appeal-Group FE	✓	✓	✓	✓	✓	✓
Date of appeal-Days from appeal FE	✓	✓	✓	✓	✓	✓

Notes: Difference-in-differences estimates of the effect of being exposed to a group appeal on same-group intentions to turn out for their own candidate. SE clustered by candidate-group-date of appeal in parentheses. The outcome variable is a dichotomous variable equal to one for respondents' turnout intentions for their own candidate. Columns 1, 3, and 5 report estimates of simple specification without covariates. Columns 2, 4, and 6 report results including covariates (race, age, gender, religion, employment status, household income, education, the approval rate of the president, news consumption from Facebook, feeling the economy is doing worse, feeling the country is on the wrong track, interest in politics, the occurrence of a Friday for Future or Black Lives Matter protest). Significance levels: *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

Figure I.5: Mobilization Effects, Difference-in-Differences Results with Alternative Window Sizes.

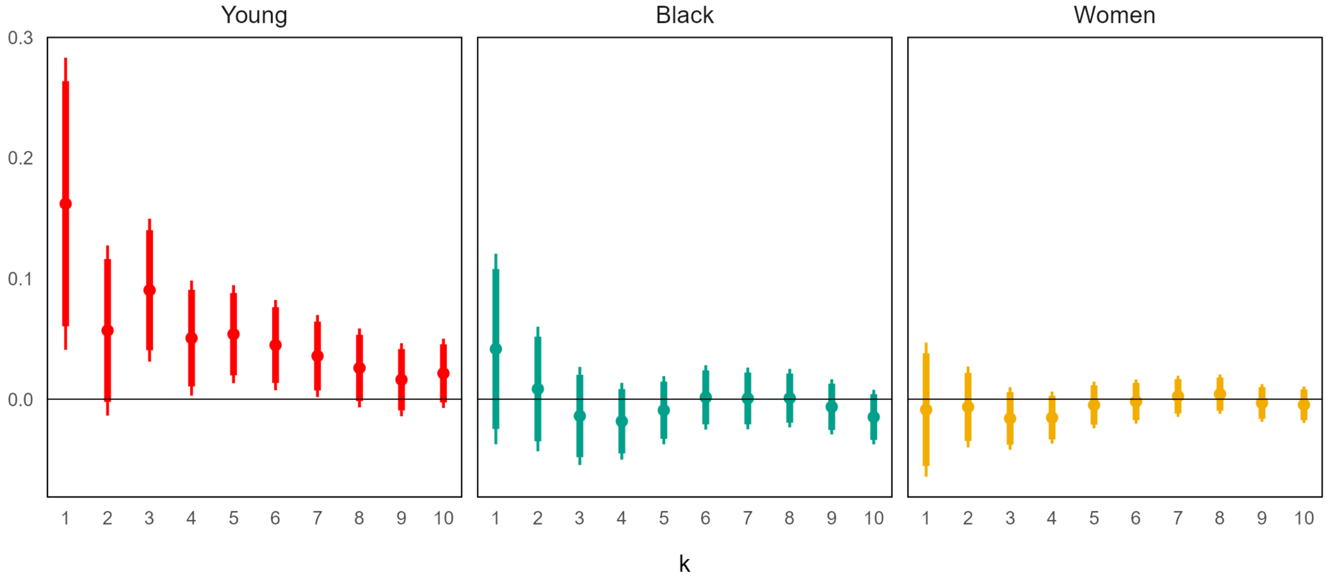

Notes: Difference-in-differences estimates with 95% and 90% confidence intervals. The outcome is a dummy variable equal to one if respondents belonging to group $G = g$ (displayed as panel titles) report they intend to vote for the candidate of the party they identify with in the next U.S. House elections in the k days after the candidate makes an appeal to group g. Covariates included (see Table I.11).

Figure I.6: Mobilization Effects, Difference-in-Differences Results with Alternative Age Bands for the Young.

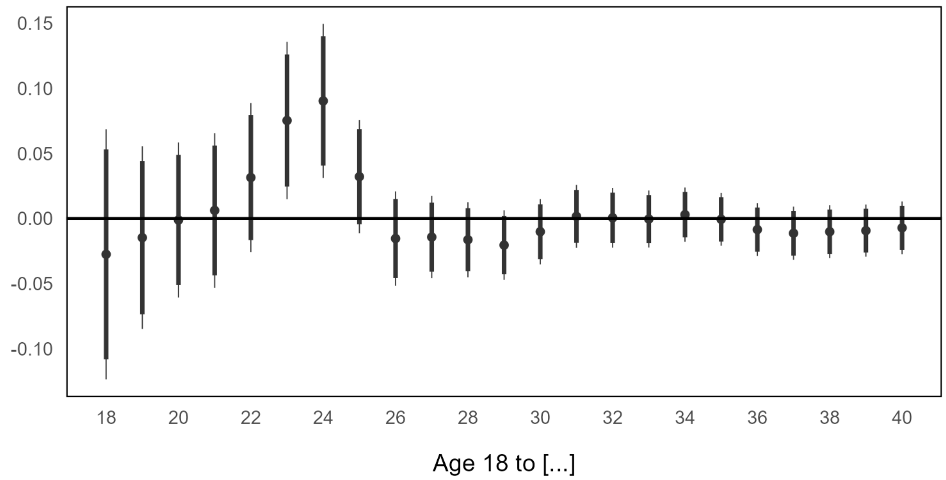

Notes: Difference-in-differences estimates with 95% and 90% confidence intervals. The outcome is a dummy variable equal to one if respondents report they intend to vote for the candidate of the party they identify with in the next U.S. House elections. The estimates are coefficients for the young × post group appeal interaction, with young respondents identified as those aged from 18 to k (displayed on the horizontal axis). Covariates included (see Table I.11).

Figure I.7: Twitter Usage by Age.

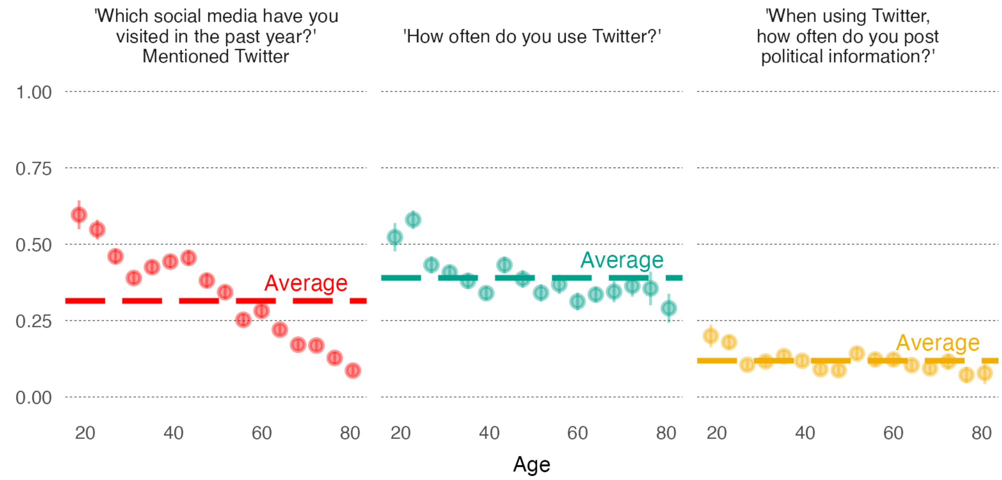

Notes: Share of respondents using Twitter by age. Data is from the 2020 ANES wave. The thick horizontal dashed line represents the sample average, and each dot represents the binned share of respondents answering affirmatively to the survey question reported in the panel title.