Research Article

American Politics Research
1–33

DOI: 10.1177/1532673X18800017
journals.sagepub.com/home/apr
SAGE

Etch-a-Sketching: Evaluating the Post-Primary Rhetorical Moderation Hypothesis

Brice D.L. Acree[1], Justin H. Gross[2], Noah A. Smith[3], Yanchuan Sim[4], and Amber E. Boydstun[5]

Abstract

Candidates have incentives to present themselves as strong partisans in primary elections, and then move "toward the center" upon advancing to the general election. Yet, candidates also face incentives not to flip-flop on their policy positions. These competing incentives suggest that candidates might use rhetoric to *seem* more partisan in the primary and more moderate in the general, even if their policy positions remain fixed. We test this idea by measuring ideological moderation in presidential campaign language. Using a supervised two-stage text analysis model, we find evidence that presidential candidates in 2008 and 2012 use more ideologically extreme language during primary campaigns, and then moderate their tone when shifting to the general election, with troubling implications for representation and accountability.

Keywords

political communication, ideology, presidential campaigning, text analysis

[1]The Ohio State University, Columbus, USA
[2]University of Massachusetts, Amherst, USA
[3]University of Washington, Seattle, USA
[4]Institute for Infocomm Research, Singapore
[5]University of California, Davis, USA

Corresponding Author:
Justin H. Gross, University of Massachusetts, Amherst, 200 Hicks Way, Thompson Hall 320, Amherst, MA 01003, USA.
Email: jhgross@umass.edu

> Fugelsang: It's fair to say that John McCain was a considerably more moderate candidate than the ones Governor Romney faces now. Is there concern that the pressure from Santorum and Gingrich might force the governor to move so far to the right it would hurt him with moderate voters in the general election?
>
> Fehrnstrom: Well, I think you hit a reset button for the fall campaign. Everything changes. It's almost like an Etch-A-Sketch. You can kind of shake it up and restart all over again.

—Eric Fehrnstrom, Spokesman for presidential candidate Mitt Romney, questioned by John Fugelsang on CNN during the 2012 presidential primaries.

Introduction

Conventional wisdom suggests that candidates should present themselves as strong partisans in primary elections, and then move "toward the center" upon advancing to the general election. Eric Fehrnstrom, in his candid comment above, offered after candidate Mitt Romney won the Illinois primary in 2012, simply provided a metaphor for a phenomenon politicos have long assumed to be true. Informal Google searches for "tack toward the center" or "flip-flop primary general" return thousands of news stories and blog posts about candidates, at all levels of government and spanning decades, undergoing similar metamorphoses: appealing to party diehards in primary elections before attempting to appear centrist when campaigning in the general election.

Yet this hypothesized phenomenon of post-primary moderation is directly at odds with the incentives candidates face not to flip-flop by changing their issue positions. Thus, on one hand, we should expect candidates to tack toward the center following the primary election, while on the other hand, we should expect candidates to maintain a set of stable positions. We argue that these expectations can be reconciled by considering that candidates can use rhetoric to *appeal* first to their base during the primary, then to more centrist voters during the general, all while maintaining essentially static policy views.

We refer colloquially to this expectation—that candidates shift from using more partisan *rhetoric* in the primary election to more centrist *rhetoric* in the general election—as the "Etch-a-Sketch"[1] hypothesis. If true, it has important implications, as electoral representation and accountability are crucial for democracy (Powell, 2000). If candidates do indeed choose language to appeal to different types of voters at each stage of the election, it suggests that during either the primary or the general election—or both—candidates are presenting a less-than-accurate version of themselves to voters; a wolf in

sheep's clothing, as it were. Thus, to the extent that voters' perceptions can be shaped by a politician's rhetoric—and research shows that they can be (Druckman & Holmes, 2004; Nelson, 2004)—the rhetoric candidates use to portray themselves as more or less partisan or more or less radical than they really are hinders voters' abilities to make informed choices about which candidate would best represent them and, moreover, which inferred promises candidates will be accountable for once elected.

To date, researchers have addressed the question of whether candidates "tack toward the center" by examining whether candidates change their *policy* stances, and find only modest degrees of moderation (Burden, 2001, 2004). These findings make sense in light of the notion that it should be electorally costly for candidates to "flip-flop" on policy positions, which are easy to compare over the course of a campaign. Yet such findings would appear to contradict the game theoretic expectation—supported by anecdotes such as Fehrnstrom's statement—that strategic candidates are those who find a way to start out more partisan or ideologically extreme during the primary, then become in some sense more moderate during the general election. Our study unifies the formal theoretic and empirical expectations of candidate behavior. We contend that scholars have overlooked perhaps the most fruitful and least risky manner in which candidates may signal moderation: changing rhetoric. Language is a powerful political tool. Through rhetoric, candidates can play to the particular audience at hand—the ideological base during the primary, the median voter during the general—without shifting positions.

Our *theory of rhetorical moderation* seeks to establish three points: (a) Candidates have an incentive to appear more moderate in general elections than in primary elections; (b) rhetoric affords a means to create and sustain a more moderate appearance without the risk of looking like a "flip-flopper"; and (c) candidates systematically change their rhetoric in transitioning from the primary to the general election. In the sections to follow, we explore these points and mobilize a large corpus of political rhetoric to illustrate our theory at work in recent presidential campaigns. To preview the findings, we show that presidential candidates do in fact use more ideologically extreme language in primary elections before moving rhetorically toward the center in the general elections.

Theory

Incentives for Moderation

Convergence theory starts from a basic premise, formalized by Hotelling (1929), Black (1948), and Downs (1957), concerning the median voter model

for majority decision making. We begin by imagining voters scattered along a single dimensional space which may represent "ideal points," broadly defined as political ideology, policy preference, or some other criterion. If voters possess single-peaked preferences, the utility each voter extracts from a potential outcome declines monotonically in the distance between the voter and the outcome. Downs argues that parties in a two-party system should thus maximize support by staking out positions near that of the median voter. If parties had no concern for maintaining unique identities, in fact, they would become indistinguishable. In practice, parties have genuine ideological and policy differences, and represent unique bases of party activists, and thus never truly converge. But the American example does show two parties more centrist in their orientations—at least historically—and more concerned with "independent" voters than parties in multiparty systems (Burden, 2001).

Applying the median voter theory to two-stage elections instead of single-cycle elections requires generalization beyond the simple theories of Black and Downs. The most naïve model asserts that candidates face two stages of median voter pressure: first, a candidate should seek the median primary voter to secure a party's nomination; if the candidate wins, he or she would then need to converge toward the general electorate median voter. As primaries tend to be dominated by party activists, who in turn tend to be more ideologically extreme (Brady, Han, & Pope, 2007), a full election cycle should see the candidate transitioning away from the general electorate median during the primary and then back toward the political center after winning the primary.

Cox (1990) defines these two conflicting sets of incentives as *centrifugal* and *centripetal* forces. Centrifugal forces push candidates outward, away from the political center. These include pressure from party activists and die-hard partisans, who tend to be ideologically consistent and demand ideological purity from candidates. Centripetal forces pull candidates toward the political center. This includes pressure from general election voters, who tend to be more moderate in their policy preferences and less ideologically constrained, and thus less likely to support "extreme" candidates.

This classic spatial model, though simple, provides ample intuition for post-primary moderation. During primary campaigns, centrifugal forces predominate as candidates seek activist support, campaign donations, and primary votes. After all, a candidate who cannot secure his or her party nomination will not earn the opportunity to pose a challenge in the general election, no matter how widely popular or centrist he or she appears. After winning a party nomination, though, centripetal forces kick in as candidates seek to win over the median constituent. Put simply, candidates will attempt to position themselves near the center of their current electorate.

Disincentives for Moderation

Post-primary moderation, broadly defined, occurs when a candidate changes his or her behavior to appeal more to centrist voters after securing a party nomination. In media coverage of elections, we frequently hear the term "run to the center," which invokes an image of candidates, erstwhile out toward the ideological poles during a primary election, rushing back to the middle ground to appeal to typical American voters, or at least reduce the risk of alienating them.

Commentators typically discuss post-primary moderation in terms of candidates shifting their political stances on one or more policy issues. This matches the canonical interpretation of the Hotelling, Black, and Downs models, which imply that party, candidate, or firm "positioning" represents a placement on a single issue, or on some scale of aggregated policy positions. Burden (2001) employs this definition in his examination of post-primary moderation, measuring aggregate policy positions in the primary and general election periods using Congressional DW-Nominate scores. Nominate scores rely upon roll call votes to measure legislators' latent aggregate policy preferences, and Burden applies the procedure to votes taken during primary and general election seasons. His analysis shows some modest evidence for post-primary moderation, though the change in legislator positions from the primary to the general election is small.

This finding should not come as a surprise: moderation comes at a cost. Burden (2004) spells out three reasons candidates may not moderate their policy positions after securing their parties' nominations. First, candidates have their own preferences and beliefs, and may not be willing to "sell out" to gain an electoral advantage. Second, candidates may lose support if the public perceives them to be "waffling" or flip-flopping on the issues. And third, candidates have reputations that often prove difficult to change during a campaign.

For a quintessential example of the risk inherent in strategic shifts of policy position during the course of two-stage elections, consider John Kerry's presidential run. As a Senator in 2002, Kerry had voted to support the Iraq War. As a presidential candidate at the launch of the primary campaign in January 2004, Kerry said, "I don't believe the president took us to war as he should have." But then during the general campaign in August 2004, when asked if he would still have gone to war knowing Saddam Hussein did not possess weapons of mass destruction, he said, "Yes, I would have voted for the authority. I believe it was the right authority for a president to have." Unlike rhetoric that merely hints at ideology, in this case, Kerry effectively changed his policy position (and then changed it back again). And this changing of policy position did not go without notice or criticism. Video footage of

Kerry windsailing (which his campaign staffers had hoped would make him seem young and healthy) was co-opted for use in a negative ad about Kerry's flip-flopping, with the tag line "whichever way the wind blows." Bush further highlighted Kerry's shifting positions during the second presidential debate, on October 8, saying, "You know, for a while [Kerry] was a strong supporter of getting rid of Saddam Hussein. He saw the wisdom—until the Democrat primary came along and Howard Dean, the anti-war candidate, began to gain on him, and he changed positions." In short, Kerry's change in effective policy stance reflects some of the costs that policy moderation can incur. The problem in Kerry's statements, however, was not that they used loose rhetoric to signal partisan ideology in the primary and then more moderate ideology in the general. Rather, Kerry's statements cued a shift in *actual policy* that was easy to identify and attack.

Formal theories of two-stage elections have long sought to deal with these costs explicitly (Agranov, 2011; Hummel, 2010). In a pair of compelling articles on the subject, Tomz and Van Houweling (2010, 2014) employ experimental designs to show that voters strongly prefer candidates who maintain consistent platforms. Respondents exposed to a flip-flopping candidate were less likely to trust the candidate's new position; accordingly, even respondents sympathetic to the candidate's new stance tended to be less supportive of the candidate after the change in position. As the authors conclude, "candidates who emerge from primaries holding positions that diverge from the median voter in the general electorate will rarely find it in their electoral interest to shift to the center" (Tomz & Van Houweling, 2014, p. 16). Furthermore, they write, candidates do not typically see electoral benefits from reversing position unless their "new" stance is favored by at least 70% of the electorate.

These costs likely explain why Burden fails to find a strong moderation effect, and more broadly why post-primary moderation has proven so difficult to detect. Candidates face strong incentives not to flip-flop on policy, unless their new position is overwhelmingly popular. Furthermore, measuring convergence with roll call votes or similar data likely obfuscates moderation. Voters pay little attention to daily votes in Congress, and most roll call votes recorded would convey little about the general ideological position of candidates anyway. Policy moderation, then, gives candidates the worst of both worlds: little potential upside at the risk of incurring a high cost. Thus, it is unsurprising that we observe so little policy moderation. But what about *rhetorical* moderation? Although candidates rarely change their policy *positions* between the primary and general campaigns, we expect they tend to shift their ideological *language* in a manner that reduces cues aimed at more extreme ideological audiences and increases cues of moderation or broad appeal.

Rhetorical Moderation

> He hasn't changed his position on immigration. He's changed the words that he is saying. (Trump Campaign Spokesperson Katrina Pierson on CNN, August 25, 2016)

Ideology represents a constrained set of basic beliefs about politics (Converse, 1964). This characterization, most frequently invoked in studies of public opinion, defines ideology in terms of abstract beliefs on fairness, justice, ethics, and morality which organize individuals' political priorities and issue preferences (Freeden, 2003; Lane, 1962). In other words, ideology entails not just policy preferences but more importantly *why* citizens hold those preferences, and how they prioritize among them.

Policy positions may not adequately tap the intricacies of ideological thought. If ideology comprises the fundamental beliefs that organize political preferences, we should incorporate expressions of ideological beliefs into our measures. A candidate can take a position for any number of reasons, each of which may derive from a different ideological interpretation of the policy. Both liberals and libertarians tend to support less restrictive definitions of marriage, for instance, but the former support it in the name of protecting marginalized social groups while the latter are suspicious of government restricting individual liberty.

We therefore argue that candidates can "move toward the center" by changing (a) the language they use in explaining their positions—framing their policy preferences in terms less narrowly appealing to an ideological base—and (b) the relative attention they pay to various policy areas, de-emphasizing the most ideologically divisive issues. This leads to a notion of *attention-weighted* ideology—ideology defined not just on one's beliefs but on the priority given or attention paid to particular values or issue areas—and shapes how we define post-primary moderation theory. Recall that moderation serves to make a candidate more appealing to centrist voters and voters likely to be alienated by ideologically extreme language. It does not matter if a candidate actually *is* more centrist during the general election campaign, so long as the voters perceive him or her to be so. Instead of "becoming" more centrist by espousing a different set of policies, candidates can seek to make themselves *appear*—or, better put, *sound*—more centrist by changing the attention they pay to various issues, values, and beliefs, and by changing the language they use in discussing these.

Most citizens do not pay close attention to politics, and thus few boast sophisticated knowledge of candidates' policy positions. Hardcore partisans, ideologues, and activists serve as the exception, and also happen to comprise

a large share of the primary electorate. For this reason, Burden (2001) anticipates that primary voters will demand firm policy commitments in-line with majority party preferences. General election voters, who tend to be less ideological and less interested in politics, struggle to differentiate candidates based on policy (Bawn et al., 2012). This gives candidates some freedom to appear more moderate than their earlier policy positions might convey, as general election voters will not generally notice any mismatch.

Candidates can *appear* more moderate by employing more balanced rhetoric and avoiding the most divisive rhetorical cues, including language strongly associated with particular movements deemed outside the mainstream by most voters. Moderate voters, even if they cannot precisely identify candidates' ideological positions, may still know that they do not like political extremists. Candidates who frequently speak on fringe concerns (e.g., disbanding the Federal Reserve, banning genetically modified foods, allowing intelligent design to be taught alongside evolution in public schools) will appear outside of the mainstream. In a primary election, this may be advantageous, but in a general election, it is dangerous. Dangerous not simply because many voters oppose a particular stance, but also because attention to fringe ideas may be taken to imply risk of other—perhaps more distasteful—ones.

Moderate, less partisan general election voters typically pay scant attention to party primaries, though, which creates an opportunity for strategic candidates to reorient their focus and rhetorical cues when transitioning to the general election period. For example, by focusing heavily on immigration during a primary election—and doing so with heavy-handed language—a Republican contender may align themselves with the priorities of base Republican primary voters. Upon advancing to the general election, the candidate may find it advantageous to spend more time discussing bread-and-butter issues such as education, unemployment, and taxes, and either talking less about immigration or employing less punitive rhetoric. Moreover, references that may serve as shibboleths for members of an activist base—for example, "social justice" on the left or "natural law" on the right—may be jarring to a more moderate or less zealous political audience. Such a shift may serve the candidate's strategic purposes in shoring up moderate Republican support and appealing to some moderate Democrats without rendering them vulnerable to charges of flip-flopping.

Candidates can also change how they explain their political vision to voters. Bawn et al. (2012) argue that even sizable policy differences can appear negligible to many voters. A candidate can capitalize on this informational asymmetry: by employing strategic messaging, the candidate may be able to make his or her policy positions appear more centrist (Iyengar & Simon, 2000). In a Democratic primary, for example, a candidate may justify their

tax plan by appealing to social justice and aid to low-income workers, but this same candidate may switch the framing of their plan in the general election, preferring to speak about funding programs that support the middle class.

To summarize, a candidate's policy positions alone may not sufficiently reveal whether the candidate tailors his or her ideological messaging to the current electorate. Candidates might maintain fairly stable official policy stances across the two-stage election, but tailor how they explain, or how much they emphasize, their positions depending on the audience. This possibility opens up exciting research avenues, for it demands a more refined way of measuring ideological moderation than tallying explicit changes in policy positions. After all, candidates in political contests are engaging in a signaling game, attempting to claim common cause with members of the electorate. Examining political rhetoric can provide the necessary information about the signals candidates send, shedding light on the language they use to they express their ideological beliefs.

Expectations

In a primary campaign, candidates should use more clearly oriented ideological language—that is, a high proportion of rhetoric that we associate uniquely with liberals or conservatives (or more radical voices of the left or right) and with particular segments of their respective base (e.g., religious right, libertarians, and cultural populists for Republicans, and progressives, socialists, and various identity-oriented social movements[2] for Democrats). After winning a primary, though, the candidate should seek more balance in their rhetoric. Instead of trying to sound highly partisan, the candidate should use more diverse language, speaking to issues of importance to liberals and conservatives, and discussing his or her world view using language that will appeal to a broad swath of the public.

Our rhetorical moderation hypothesis could theoretically hold in any two-stage election where a party nomination contest is followed by a general election. In many contexts, however, the electorates may not be distinct enough to create a true moderation incentive. As Brady et al. (2007) argue, many Congressional races involve primary and general election constituencies that overlap considerably, to the point where candidates face little incentive to make significant changes after securing the nomination. Furthermore, we should only expect transformations when both the primary and general elections are contested, and thus likely to attract broad voter attention (Agranov, 2011; Meirowitz, 2005). In low-salience or uncontested elections, the potential upsides of the moderation strategy would almost certainly be washed out by the risk of appearing disingenuous (Hummel, 2010).

Finally, we would only expect rhetorical moderation in cases where public attitudes toward candidates are fluid. As Burden (2001) argues, post-primary moderation only works if voters are persuadable with respect to the candidate's ideological position. A candidate with an established reputation as a strong liberal will have little reason to appear more moderate. In fact, that candidate may appear disingenuous in the eyes of a public that already believes him to be liberal (Burden, 2001).

For these reasons, we focus our energies on presidential candidates, and the 2008 and 2012 elections in particular (these being the two most recent presidential elections at the time of our data collection). In presidential elections, the stakes of party nominations are extremely high, attracting quality candidates and the investment of party activists and donors, who are likely to demand that nominees represent the preferences of party loyalists. Presidential campaigns also entail two distinct electorates: party primary voters and a more centrist general electorate. And presidential campaigns also vary in how established public attitudes toward candidates are. Early in 2008, for example, many voters had no opinion of either Barack Obama or John McCain. By 2012, however, most citizens had an established opinion about incumbent President Obama, even if they lacked a fixed opinion about Mitt Romney. We therefore expect to see strong rhetorical moderation from Barack Obama and John McCain in 2008, and Mitt Romney in 2012, but little or no moderation from Obama in 2012.

Design

Our theory demands a research design that can shed light on how presidential candidates communicate their ideological positions, and how their communication may change after winning a primary nomination contest. This naturally requires some form of text analysis. For many decades, scholars coded textual data with labor-intensive human coding. Yet, as King and Lowe (2003) point out, human coding places inordinate demands on researcher resources and may not produce reliable data. Given these revelations, and the burgeoning capacity to deal with large data sets, we employ an automated approach to studying presidential rhetoric.

Measuring ideology in language is not straightforward. We contend that ideological beliefs—or strategic representations thereof—will manifest through the written or spoken word. More specifically, we assume that political language contains ideological *signals*. These signals can take many forms: explicit policy positions, attention paid to particular issues, or even simple word choice. For example, conservative politicians may, on average, speak more on immigration than on Medicaid, and may be more likely to use phrases

such as "border security" or "illegal alien" than liberals, who may instead refer to "comprehensive immigration reform" and "undocumented immigrants."

Roderick Hart (2000), in a book on the word choice of presidential candidates, argues that campaigns serve to "teach" about politics, "preach" as a kind of performative ritual, "sensitize" us to political content at regular intervals, and "activate" the masses by increasing their sense of political efficacy to mobilize them to vote. He notes a drop in direct party-based references in the language of candidates since the 1950s, but warns of a potential for "stealth politics" in its place. He focuses primarily on style over ideology, noting that George McGovern—recalled by many as too radical to have a chance at beating Nixon in 1972—in fact presented a "midwestern, ministerial style [that] consistently overrode his philosophical biases" during the general election period. Nonetheless, Hart concedes that "the general campaign is not the primary, and more than one author has documented McGovern's slide to the middle after he secured the Democratic nomination" (chapter 4, endnote 57). Such documentation, however, remains impressionistic, as in the one citation provided by Hart (2000), Bob Greene's (1973) *Running: A Nixon-McGovern Campaign Journal*.

We treat ideological language quite generally. We seek to measure the ideological content of political language by tracking the signals that political candidates convey through their speech. To identify those signals, we employ a supervised approach—specifically, we seek to learn ideological signals from a corpus of political writings by (mostly) recognized authors of known ideological leanings. Then we use this information to observe whether, on average, the ideological signals employed by candidates change between the primary and general election.

Our research design unfolds as follows: using a large corpus of political texts, we identify the key phrases associated with particular ideological subcorpora. Then, using this vocabulary, we examine the campaign speeches of all major party presidential candidates in 2008 and 2012, paying particular attention to the three candidates who competed both in presidential primaries and in a general election: Barack Obama (2008), John McCain (2008), and Mitt Romney (2012). If the post-primary moderation theory bears out in the data, we would expect these three candidates to transition toward the ideological center after securing their parties' nominations.

Identifying Ideological Cues

To identify the key phrases that help us to infer a speaker's ideology, we construct a corpus of 170 books and magazines from authors of familiar ideological bents. We restrict our attention to books and issues of magazines

published that have nearly all been published over a narrow time-frame, to minimize temporal variation in political rhetoric as much as possible. Following the method described by Sim, Acree, Gross, and Smith (2013), we assign each text an ideological class from the set {Left, Center, Right} based on our understanding of the author's ideological beliefs, according to their reputations and self-identification. We further assign to each text an ideological *sub*class, again based on our judgment of the authors' predominant ideological beliefs. Liberals and others on the Left are assigned a subclass from the set {Socialist, Progressive, Center-Left, Religious Left}. Conservatives and others on the Right are assigned a subclass from the set {Populist Right, Religious Right, Far Right, Libertarian, Center-Right}. As outlined in Sim et al. (2013), we have informally validated the class labels with a panel of experts—five graduate students specializing in contemporary American politics—who confirmed that the phrases extracted from the model matched their intuition about the ideologies of the authors.[3]

Unlike many existing approaches to extracting ideological meaning from text (see Diermeier, Godbout, Yu, & Kaufmann, 2012; Laver, Benoit, & Garry, 2003; Lowe, Benoit, Mikhaylov, & Laver, 2011; Monroe, Colaresi, & Quinn, 2008; Monroe & Schrodt, 2008; Slapin & Proksch, 2008, inter alia), we do not explicitly seek to scale ideologues into a Euclidean space. Instead, we approach ideology as a discrete construct. This better matches our theoretical understanding of ideology as a hierarchical set of clustered ideas (Freeden, 2003). To maintain as much nuance as possible in the political texts, we avoid collapsing the unique ideas of libertarianism and religious conservatism, for example, into a single conservative class. Rather, we seek to capture the meaningful nuance of multiple ideological classes within broad left and right divisions.

To facilitate analysis of our textual data, we remove common words (stop words), such as "and," "the," and "or." We also stem words, meaning we combine words based on word stems so that words like "debate," "debated," and "debating" all register as versions of the same word, instead of counting as unique words. Finally, we combine individual words into bi-, tri-, and four-gram word phrases, or *tokens*. A token is simply a unique word or multi-word phrase that serves as our unit of analysis.[4]

Each document in the corpus is thus represented at this stage as a bag of multi-word phrases. Recall that each document also has a hand-labeled ideological class and subclass designation. To identify tokens that distinguish classes and subclasses from one another—for example, to identify language that occurs frequently in documents written by libertarian but not religious liberal authors—we fit a sparse additive generative model (SAGE) (Eisenstein, Ahmed, & Xing, 2011). With known document labels, this amounts to a hierarchical multinomial logistic regression with sparsity-inducing Laplace priors on

the ideological class and subclass effects.[5] In essence, we are estimating the added logodds of term appearance associated with particular ideologies, subideologies, topics, and authors.

Let W be a set of terms—in this context, bi-, tri-, and four-grams—that may possibly overlap multiple ideological classes. Furthermore, let ε_k be the additive effect for ideology k and let e_0 be the background effect. Then, for each term $w_i \in W$ and each effect e_0 and $e_k \in \varepsilon$, we generate $\theta_{e,w}$ from a zero-centered Laplace distribution, with hyperparameter λ_e^{-1} which varies over the effects.

Now, let D be the set of documents (books and magazines), and assume that each document $d \in D$ contains a bag of terms. We have observed the effects for each document, each given by $\varepsilon_d \subseteq \varepsilon$. For each term t in each document, the model generates $w_{d,t}$ from a multinomial distribution with probabilities given as follows:

$$p(w \mid \theta) = \frac{\exp \sum_{e \in \varepsilon_d} \theta_{e,w}}{\sum_{w' \in \mathcal{W}} \exp \sum_{e \in \varepsilon_d} \theta_{e,w}}. \quad (1)$$

This is the probability of observing a given term conditioned on the additive effect for the term associated with an ideological category. For more, see page 2 of Eisenstein et al.'s (2011) paper.

To estimate the model parameters, we must first estimate the background effect. Fortunately, this is fairly straightforward; the background effect is fit using logged relative frequency of each term w in each document $d \in D$:

$$\theta_{e_0,w} = \log \frac{\text{freq}(w; \mathcal{D})}{|\mathcal{D}|}, \quad (2)$$

and the remaining ideological effects are estimated by solving the unconstrained convex problem via the expectation maximization (EM) algorithm:

$$\max_{\theta} \sum_{d \in \mathcal{D}} \sum_{t=1}^{|d|} \log p(w_{d,t} \mid \theta) + \sum_{e \in \varepsilon} \lambda_e \sum_{w \in \mathcal{W}} |\theta_{e,w}|. \quad (3)$$

The maximum a posteriori (MAP) estimates of the parameters give the added log-odds that each term appears in an ideological category. SAGE estimates, then, tell us how much more or less likely an author from a particular ideological class (and subclass) will use any of the words in our vocabulary.

Testing the Rhetorical Moderation Hypothesis

To examine whether presidential candidates exhibit rhetorical moderation following their party primaries, we analyze transcripts of presidential campaign

Table 1. Number of Speeches Available for Each Candidate, and the Average Number of Terms per Speech per Candidate.

Candidate	Party	Years of Speeches	Number of Speeches	Average Terms per Speech
Michele Bachmann	Republican	2011-2012	6	2,259
Joe Biden	Democrat	2007-2008	6	2,543
Herman Cain	Republican	2011	2	2,992
Hillary Clinton	Democrat	2007-2008	105	3,759
John Edwards	Democrat	2007-2008	27	3,072
Newt Gingrich	Republican	2011-2012	15	3,163
Rudy Giuliani	Republican	2007-2008	36	1,863
Mike Huckabee	Republican	2007-2008	14	2,556
Jon Huntsman	Republican	2011-2012	5	1,899
John McCain*	Republican	2006-2008	169	2,542
Barack Obama*	Democrat	2008-2012	258	3,450
Ron Paul	Republican	2012	7	2,439
Rick Perry	Republican	2011-2012	10	1,772
Mitt Romney*	Republican	2006-2012	59	2,551
Bill Richardson	Democrat	2007-2008	29	3,050
Rick Santorum	Republican	2011-2012	16	2,457
Fred Thompson	Republican	2007-2008	15	2,683

Note. Asterisks designate the candidates observed in both the primary and general election.

speeches from the 2008 and 2012 presidential primary and general elections. We focus on these two election cycles for practical reasons of data limitations. Digital copies of political texts did not become widely available until the mid-aughts. The earliest book in our Contemporary American Ideological Books (CAIB) Corpus[6] was published in 2007, and the newest was published in 2014. Because ideological signals might change over time, we limit ourselves to this time frame, namely, the 2008 and 2012 elections.

Collecting transcripts from the University of California Santa Barbara's American Presidency Project yields 779 speech transcripts from 17 candidates (five Democrats and 12 Republicans). Table 1 breaks down the number of speeches per candidate, and the average number of terms per speech. For the three candidates who advanced to the general election, we also denote each speech's "epoch," defined as whether the candidate delivered the speech during the primary election (before securing enough delegates to win his party's nomination) or the general election (after winning his party nomination).

Having built a "dictionary" of ideological phrases in the first stage of analysis, we proceed to examine the candidates' use of ideological language during

the presidential campaign. Following the innovative method introduced in Sim et al. (2013), we drop the bag-of-terms assumption and model each campaign speech as a nonexchangable sequence of cue phrases interdigitated by filler words. Similar to a hidden Markov model (HMM), the Cue-Lag for Ideological Proportions (CLIP) model allows inference of the ideology of a speaker based on the phrases the speaker uses *and* the order in which the speaker employs those phrases. The CLIP model additionally tracks the length of "lag periods"—the number of filler words an author uses between uttering ideologically important cue phrases—to illuminate how long the author has spent speaking from various ideological perspectives.

To see how the model works, consider starting at the beginning of a speech where we assume the originating ideological state can be treated as some generic background political state. Moving forward in time, we encounter the first ideologically separating term. To infer the ideological class and subclass from which that word was "emitted," we refer to the ideology-specific multinomial distributions over the terms in the vocabulary. We use the SAGE estimates from the first stage to specify priors on the the class- and subclass-specific term distributions.[7] As we move forward through the speech, we encounter perhaps more uninformative terms before coming to the second ideologically separating term.

At this point, we have two pieces of important information: the term distributions for the ideological classes and subclasses, and the inferred class and subclass from the previous term. The former is again informed by the SAGE weights. For the latter, the HMM allows us to model the speaker's transition from (sub)class to (sub)class. The tree structure (see Figure 1) counts the edges between subclasses, and a priori favors shorter transitions. For instance, if the model classifies w_1 as being emitted from the "religious left" subclass, the model places somewhat higher probability that cue term w_2 will be emitted from another left subclass rather than from a center or right subclass. This is effectively a substantive prior on the ideological structure of American politics: it smooths the potential switches between ideological classes and subclasses, but does not absolutely prevent the model from inferring switches when the text warrants it.

Our primary inferential task is to ascertain how long an ideologue spends in each ideological state. To do so, we tally the lag periods between ideologically interesting cue terms—that is, we count the filler terms between cues. We attribute half to the model's previous inferred ideological state, and half to the models currently inferred ideological state.[8]

The high-level generative story for a single speech with T cue-lag pairs follows the generative story of a HMM with Bayesian priors for the parameters:

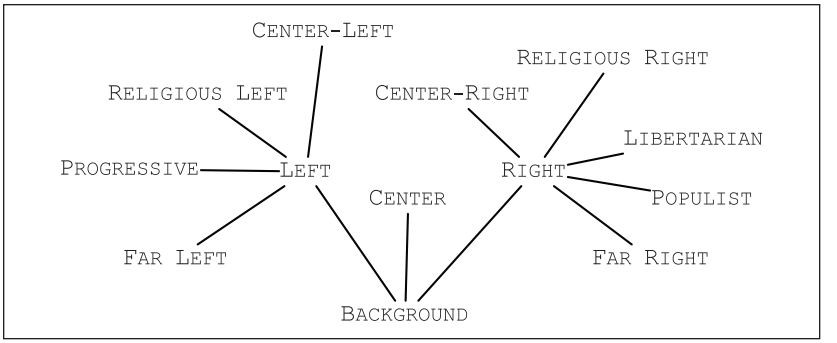

Figure 1. Ideology tree structure.

1. Parameters are drawn from conjugate priors.
2. Let the initial state be the Background state.
3. For $t \in \{1,2,\ldots,T\}$:[9]
   a. Transition to state S_t based on the transition distribution. This transition is conditioned on the previous state S_{t-1} and the lag at timestep $t-1$, denoted by L_{t-1}.
   b. Emit cue term W_t from the lexicon associated with state S_t, and emit the lag L_t.

What differentiates this model from a conventional HMM is that the transition parameters are defined in terms of the ideology tree (Figure 1), and the emission parameters are defined in terms of the lexicons and the lags, as discussed above.

A more complete technical model exposition can be found in the online appendix, or in Sim et al. (2013). Model parameters are estimated by Markov Chain Monte Carlo (MCMC) with EM steps, and with Gibbs sampling during the E-steps to sample the remaining latent variables $x_{d,i}$ and $r_{d,i}$, and direct optimization during M-step for ζ. λ can be set to the maximum likelihood estimate as it is essentially observed and independent.

This model provides certain advantages over other available methods.[10] First, we do not need to classify candidate speeches as deriving from a single ideological perspective. We recognize that candidates will necessarily adopt ideas from several ideological groups into their campaign appeals. The CLIP model allows for this by inferring ideological perspective on a token-by-token basis. Each time a candidate selects an ideological phrase (i.e., chooses to emit a cue), they have the option to transition from their current ideological perspective to a new one.

Second, the model leverages our hierarchical conception of ideology and the order in which phrases appear in a speech. Unlike bag-of-words methods, in which word order does not matter, the CLIP model reflects the simple truth that speakers are unlikely to skip around from ideological perspectives with abandon. Instead, the model uses the similarity between ideological classes and subclasses when estimating the perspective from which a speaker is emitting phrases. For example, if a candidate is sounding like left-leaning authors and then uses vocabulary associated with both Religious Left and Religious Right, these new terms will more likely be estimated as generated from the former than the latter.

Putting Together the Design Elements

To summarize, our research design involves two stages. In the first stage, we employ a large corpus of political texts to build a dictionary of ideological phrases. For each multiword phrase in this training corpus, we estimate an ideological and subideological "effect," which we use to infer how likely ideologues of a certain class–subclass profile are to employ each phrase a priori. After thus defining the dictionary, we proceed to the analysis of campaign speeches from the 2008 and 2012 presidential elections using only the selected terms. Using the CLIP model, we estimate the share of time each candidate spends emitting phrases from each ideological class and subclass. During inference, the probability that each token comes from each ideological class and subclass is estimated, given the context of other tokens, and also taking into account as a prior the associations uncovered by SAGE.

The post-primary moderation theory leads us to predict that McCain and Romney should employ rhetoric consistent with the more right-wing ideologies during the Republican primaries, before moving toward the center in the general election; conversely, Obama should employ more liberal language in the Democratic primary and more balanced rhetoric in the general election in 2008. During 2012, we would also anticipate that Obama should largely abandon a balanced rhetoric, as his 4 years in office had largely cemented the view of Obama as a liberal.

Results

To assess the veracity of the rhetorical moderation hypothesis, we examine the estimated proportion of "time" each candidate spends speaking from liberal and conservative perspectives. Even though the CLIP model does not scale candidates into a Euclidean space, we can represent the degree of ideological "balance" by the proportion of time candidates spend using left- and right-leaning ideological language. If a candidate used roughly equal parts

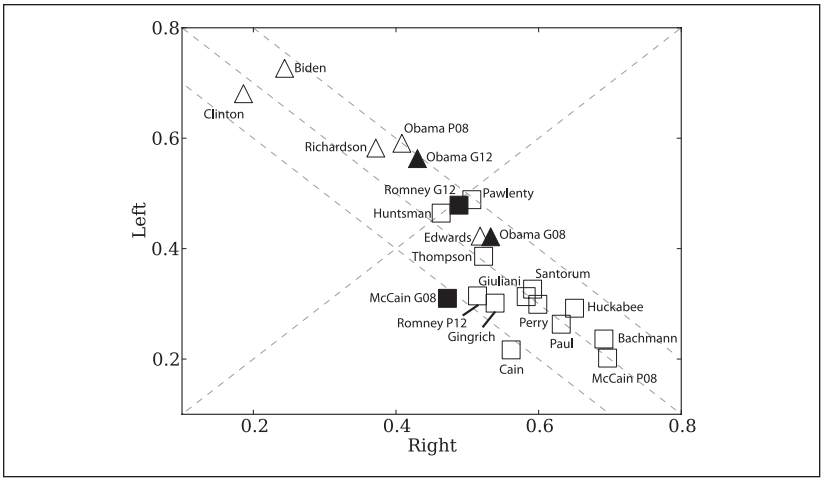

Figure 2. Proportion of "time" spent using phrases from the liberal (*y* axis) and conservative (*x* axis) ideological vocabularies.

Note. Hollow figures indicate the primary speeches and solid figures indicate general elections; squares represent Republican candidates and triangles represent Democrats. The intersection of the main diagonal lines marks a 50-50 balance between time spent sampling from the left and right ideological vocabularies. Because the "centrist" category is omitted from the figure, proportions may not sum to unity.

liberal and conservative language, we could think of the candidate as sounding moderate in their ideological expression.[11] Put another way, the balanced use of "liberal" and "conservative" language does not necessarily mean a candidate is appealing to each side at equal rates; rather a lack of clear signaling is likely to confound model estimation in a manner analogous to how it confounds voters attempting to pigeonhole the candidate. Table 2 provides the estimated proportional speech time spent drawing on phrases from each basic ideological orientation and subideological class.

Figure 2 shows aggregate left-right proportions of all candidates from the 2008 and 2012 presidential elections. To compute these estimates, we simply add up the proportion of time each candidate spends speaking from *any* of the liberal versus *any* conservative vocabularies, as detailed in Table 3. Accordingly, a candidate at the intersection of the two diagonals would represent campaign rhetoric perfectly balanced between left and right. Candidates in the lower right quadrant represent a high proportion of conservative rhetoric and low levels of liberal rhetoric, while the top-left quadrant represents the inverse.

As the figure shows, our basic intuition about American party identities is borne out in the data. Republican candidates tend to occupy the conservative section of the figure, while the Democratic candidates tend to occupy the

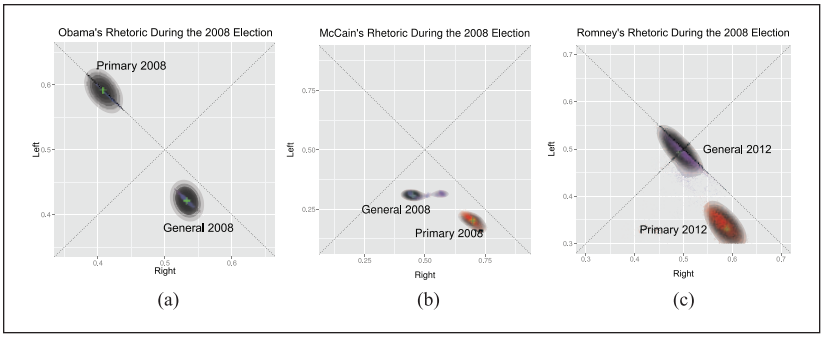

Figure 3. Evidence for the rhetorical moderation hypothesis from three presidential candidates in two election cycles: (a) Barack Obama in 2008, (b) John McCain in 2008, and (c) Mitt Romney in 2012.

Note. Shaded regions represent high posterior density regions, smoothed using a bivariate normal kernel density smoother. Crosshairs ("+") mark the posterior medians for each marked electoral stage. The intersection of the main diagonal lines marks a 50-50 balance between time spent sampling from the left and right ideological vocabularies. Because the centrist category is omitted from the figure, proportions may not sum to unity.

liberal section. Candidates who occupy positions further from the center also tend to match our intuition: the model estimates Michele Bachmann, Herman Cain, Ron Paul, Mike Huckabee, Rick Santorum, and Rick Perry to use higher proportions of conservative language, while moderates Jon Huntsman and Tim Pawlenty are located nearer the center. We also draw reader attention to evidence supporting one of our hypotheses: setting aside whether he moderated in 2008, we see that President Obama's 2012 general election rhetoric was quite liberal. In fact, his rhetoric nearly matched the ideological balance of his *primary election* rhetoric from 4 years earlier. This matches our expectations based on Burden (2001): because Obama was well-established in the public mind by 2012, he had little to gain from attempting to appear more centrist.

Figure 2 also offers evidence supporting the rhetorical moderation hypothesis, although the abundance of other candidates makes it difficult to decipher. Referring instead to Figure 3, we can see how Obama, McCain, and Romney all employed more balanced rhetoric in the general election. To indicate measurement uncertainty, we include the highest Bayesian posterior density regions around these point estimates. Darker areas of the plot signify higher posterior probabilities, which we can interpret as the most likely ideological mixture proportion[12] of the candidate at each election stage. In each case, we see little to no overlap in the primary and general election posterior

distributions, providing strong evidence that the candidates shifted their rhetoric in the expected direction. In fact, the posterior probability of the expected shifts (e.g., Obama sounding more liberal during the Democratic primary and more conservative during the general election) are at least 0.99 for each candidate, indicating negligible chances of observing these patterns by chance alone.[13]

Given the aggregate evidence supporting the rhetorical moderation theory, we turn now to more refined measures of ideological language. Here, the noisy nature of text returns slightly more complicated results. As Table 2 shows, we estimate that candidates spend relatively little time speaking from the more refined ideological perspectives. We believe this to most likely occur for both substantive and methodological reasons. Substantively, we believe that candidates must appeal not only to their unique ideological sympathizers but to the party base as a whole. As a result, the candidates' language will borrow some phrases from specialized subideologies, but generally will represent a broader class of ideological language. And on the methodological side, the ideological subclasses both overlap in their vocabularies and contain less information than we can learn from the aggregate top-level class labels. It thus becomes easier to classify phrases as deriving from the broader ideological class than from the more refined subclasses, yielding much higher proportions for the top-level "left" and "right" proportions than for the refined subclasses.

Tables 4 and 5 contain the top distinguishing phrases of left and right ideologies, respectively, based on documents from the book and magazine corpus, as learned in the first stage of the CLIP model estimation process. Table 6 does the same for texts designated as centrist. Some of these terms are obviously meaningful, as for example in the populist right's fixation on immigration, trade, race, and culture–with certain phrases such as "illegal aliens" and "open border" distinctive to their particular perspective. The same can be said for the religious right, with "Christian nation," "anti-Christian," "pro-life," and "San Francisco"–a modern version of Sodom and Gomorrah in such writings–among top cues. Others serve as shibboleths likely unfamiliar to outsiders. References to "raw milk" by libertarians abound, as laws requiring pasteurization serve as symbols of government regulations run amok in a supposed nanny-state. "Property rights," "economic freedom," and mentions of various constitutional rights and protections have more obvious connections to libertarian ideology.

Centrist (often bipartisan) vocabulary includes many technocratic concerns, references to visions of middle America–"average American," "American dream," "country music"–and the word "Centrist" itself, as well as names of right-wing talk radio hosts, ostensibly as examples of

Table 2. Estimated Percentage of Time Each Candidate Spent Speaking From Each Ideological Perspective.

	Left	Socialist	Center-Left	Progressive	Religious-Left	Right	Far Right	Libertarian	Center-Right	Populist	Religious-Right
Obama P08	16	6	19	5	12	4	5	5	9	12	5
Obama G08	15	7	9	8	3	6	1	19	12	6	9
Obama G12	50	0	1	5	0	38	0	1	0	0	4
Romney P08	25	1	0	2	20	10	1	6	17	4	13
Romney 08-11	28	2	5	12	0	28	14	7	3	0	0
Romney P12	28	0	1	1	2	34	8	14	0	1	2
Romney G12	29	13	2	1	4	29	11	5	1	2	0
McCain P08	8	4	2	2	4	18	2	4	39	6	1
McCain G08	20	4	2	4	1	21	2	4	12	5	2
Bachmann	14	2	2	3	3	38	8	12	4	2	5
Biden	56	3	2	8	3	19	2	1	1	1	1
Cain	11	1	4	2	4	33	5	2	5	2	8
Clinton	48	1	1	15	2	13	0	2	2	1	1
Edwards	40	2	14	11	1	21	1	1	3	1	2
Giuliani	20	2	4	3	3	31	5	5	6	9	3
Huckabee	23	0	2	1	2	33	3	7	6	4	12
Huntsman	33	1	2	1	10	31	4	2	5	2	2
Paul	21	1	2	2	1	41	3	7	3	1	7
Pawlenty	33	5	4	3	8	16	3	6	7	1	7
Perry	21	2	2	3	2	33	5	7	9	2	4
Richardson	29	10	1	13	5	21	1	1	3	6	4
Santorum	25	1	1	2	1	36	3	4	5	2	3

Note. Point estimates are given as posterior means. "Left" and "Right" refer to the top-level ideological classes, and other designations represent ideological subclasses. Rounding may result in estimates not summing to unity.

Table 3. Estimated Percentage of Time Each Candidate Spent Speaking From the Summed Left and Right Perspectives, and From Centrist Perspective.

	Σ_{left}	Σ_{right}	Centrist
Obama (Primary 2008)	59	41	0
Obama (General 2008)	42	53	5
Obama (General 2012)	56	43	1
Romney (Primary 2008)	49	51	0
Romney (2008-2011)	46	52	2
Romney (Primary 2012)	33	59	8
Romney (General 2012)	48	49	3
McCain (Primary 2008)	20	70	10
McCain (General 2008)	31	47	22
Bachmann	24	69	7
Biden	73	24	3
Cain	22	56	22
Clinton	68	19	13
Edwards	68	28	4
Giuliani	31	58	10
Huckabee	29	65	6
Huntsman	46	46	7
Paul	26	63	11
Pawlenty	53	39	7
Perry	30	60	10
Richardson	58	37	5
Santorum	30	54	16

Note. Point estimates are given as posterior means. Rounding may result in estimates not summing to unity.

people contributing to polarization. Similar distinctions can be identified by examination of the distinguishing phrases on the left. Terms such as "monopoly capitalism," "class struggle," "ruling class," and "corporate state" seem to be distinctively socialist. Other top terms of the socialist left, such as "social justice," are also associated with other left-leaning writers, especially progressives, but do not appear in the top handful of their terms shown in Table 4.

To facilitate examination of this research by other scholars, we have provided full results from the model online via an interactive module.[14] An interesting feature of the module is that it allows for easy perusal of the actual ideological tokens used by each candidate and the modal posterior ideological classification estimated by the model. These results shed some light on how exactly the candidates altered their ideological language as they transitioned from the primary to the general election.

Table 4. Phrases That Most Clearly Distinguish Each Ideological Subclass From the Others, Based on the Class- and Subclass-Effects Estimated Using the First-Stage SAGE Model.

Left Ideologies			
Progressive L	Socialist L	Center L	Religious L
United States	United States	United States	biological family
American Prospect	monopoly capitalism	modern art	progressive religion
Abu Ghraib	class struggle	young woman	nuclear family
executive director	occupy movement	eighteenth century	bad theology
public info	political economy	al Jazeera	religious issue
State Department	capitalist system	Mitt Romney	early church
public policy	trade union	twentieth century	religious community
Vice President	labor movement	great deal	American creed
interest rate	ruling class	free market	Willow Creek
Head Start	developing country	tax code	strict father
United Nations	public option	free trade	early Christians
mental illness	working people	common law	Mother Theresa
John Kerry	Dr. King	supply siders	family values
San Francisco	American worker	South Africa	God's love
Iraq war	economic crisis	living conditions	tax collector
child care	social justice	K Street	NUM_Matthew
civil war	working class	interest rate	NUM_Mark
political power	fossil fuel	Grover Norquist	NUM_Luke
community college	corporate state	private sector	NUM_John

Note. Many of the terms match the expectations of area experts, with the exception the rather anemic centrist category. The difficulty in extracting meaningful "centrist" words largely stems, we believe, from the relatively few truly centrist authors of political tracts. SAGE = sparse additive generative model.

Take, for example, Barack Obama's usage of the phrase "middle class." During the 2008 Democratic primary, Obama referred to the middle class relatively infrequently—only 18 times—compared to frequent use (134 times) in the general election. In the primary, Obama also focused much more attention on universal health care (41 mentions in primary speeches vs. four in the general election), the death penalty (11 vs. 0), and mandated fuel efficiency standards (11 vs. 3), while increasing attention on capital gains taxes on small businesses (1 time in the primary vs. 121 times in the general), small businesses more generally (8 vs. 114), tax credits (26 vs. 76), and tax relief more generally (0 vs. 48).

Romney altered his rhetorical attention similarly, sharply cutting his references to "economic freedom" (25 times in the primary vs. 8 in the general

Table 5. Phrases That Most Clearly Distinguish Each Ideological Subclass From the Others, Based on the Class- and Subclass-Effects Estimated Using the First-Stage SAGE Model.

Right ideologies				
Center R	Populist R	Religious R	Far R	Libertarian R
Governor Bush	illegal aliens	Christian nation	illegal immig.	raw milk
Ronald Reagan	illegal immigrants	human beings	North Korea	property rights
human being	border patrol	Jesus Christ	flat tax	chief justice
foreign policy	immig. reform	anti-Christian	political correct	natural right
Middle East	birthrate	New York	big gov't	founding fathers
John McCain	national media	God's word	affirmative action	eminent domain
radio program	civil war	Sarah Palin	Jewish state	Supreme Court
Saddam Hussein	Los Angeles	fed. gov't	human life	constitutional right
Repub. Party	white America	San Francisco	Saudi Arabia	TSA agent
mass destruction	African American	private property	Arab Spring	Ninth Circuit
Karl Rove	open border	elementary school	radio show	bear arms
Cold War	trade deficit	Holy Spirit	American exception	free speech
Sam's Club	elected officials	year old	left wing	Habeas Corpus
look back	corporate America	pro life	mainstream media	Fourth Amendment
Social Security	free trade	Joseph Smith	popular culture	Fourteenth Amendment
Miss America	culture war	public schools	global warming	judicial activism
Iraq war	special interest	Planned Parenthood	hard work	economic freedom
mind-set	national interest	Judeo-Christian	Muslim Brotherhood	Federal Reserve
class voter	working men	daily saint	Ron Paul	medical marijuana
health care	border security	Jim Wallis	foreign aid	defense attorney

Note. Many of the terms match the expectations of area experts, with some confusion between Far Right and Populist Right classes. SAGE = sparse additive generative model.

election), accusing Democrats of wanting to raise taxes (13 vs. 4), and completely dropping self-references as a conservative (7 vs. 0). Meanwhile,

Table 6. Phrases That Most Clearly Distinguish Centrists From Subideologies of Left or Right, Based on the Class- and Subclass-Effects Estimated Using the First-Stage SAGE Model.

Centrist ideology
Centrist
Long Beach
debt limit
stock option
country music
average American
corporate America
original intent
tax increase
Mark Levin
American dream
Wal Mart
super PAC
George Washington
loan office
Republican Party
Glenn Beck
Washington, DC
Alexander Hamilton
debt ceiling

Note. Many of the terms match the expectations of area experts, with the exception the rather anemic centrist category. The difficulty in extracting meaningful "centrist" words largely stems, we believe, from the relatively few truly centrist authors of political tracts. SAGE = sparse additive generative model.

Romney increased his references to the middle class (8 to 23 times), the African American community (0 to 9 times) and education reform (0 to 14 times) after winning the Republican nomination. And McCain, for his part, cut his references to the Bush tax cuts in half despite the fact that McCain used a far greater number of ideological phrases in the general than in the primary.

In examining these examples from the source material, we do not find candidates changing or adopting new policy positions. Rather, the candidates merely use their language to draw attention to different issues, or to explain those positions in a different manner. Of course not all changes are quite so stark, which speaks to why our methodology serves us well in this exercise. By leveraging a supervised but automated approach to analyzing political rhetoric, we balance the need for substantive area expertise with the ability

for statistical models to detect subtle but prevalent patterns of rhetorical change. Our results show that candidates strategically craft their rhetoric to their electorate, and most importantly, that candidates systematically engage in post-primary ideological moderation.

Discussion

In this article, we have presented a theory of rhetorical moderation. Journalistic accounts of presidential campaigns suggest that strong partisans drag electoral candidates toward the ideological extremes during primary contests. Candidates who secure their party nominations must then shift back toward the median general electorate voter or else risk voters perceiving them as extremists. This tendency to move toward the center features prominently in campaign narratives, especially of presidential elections, though the phenomenon has seen almost no empirical support.

To test the idea that candidates are indeed using rhetoric to make it appear as though they are more partisan during the primary but more moderate during the general, we employ a novel approach to computational text analysis of campaign speeches. Using a two-stage model for classifying the ideological composition of political rhetoric, we show that all three major party nominees from the 2008 and 2012 election cycles underwent similar ideological transformations after winning their party primaries. In both quantifiable and substantively meaningful ways, Obama, McCain, and Romney all spoke as more ideologically extreme candidates during the primary season, before changing their language to sound more ideologically balanced in the general election.

We find the evidence we have presented compelling, and encourage the reader to consider three points. First, the moderating pattern occurs in two election cycles and with three candidates, and in all three cases, the rhetorical moderation was pronounced, and in the predicted direction. This is unlikely to occur by chance (a conservative estimate would be a one-in-eight [12.5%] chance, the same as getting all heads when flipping a fair coin 3 times). Second, the methods we employ to measure moderation are clearly picking up on an ideological component to political language, as evidenced by the top terms associated with ideological classes and the reasonable liberal-conservative positioning of the 17 candidates under consideration. Third, the pattern we find is robust to model specification.

Can we go further in assessing the nature of the linguistic shifts observed? Namely, do they seem purely rhetorical, or do they reflect actual changes in policy position? To investigate this question, we considered the speeches by the three general election candidates in 2008 and 2012 on five policy areas:

health care, immigration, taxes, abortion, and environmental policy. We chose these as a hard test for the hypothesis that, although their rhetoric is shifting, their policy positions are not. On each of these issues, the two party bases hold more extreme preferences than the general public. These areas are thus ripe for policy flip-flopping should it occur. Thus, while absence of evidence is not evidence of absence, we look closely where the most pronounced evidence of policy shift should occur if it exists.

Our first observation is simple: we find no evidence for policy shifts between the primary and general elections. In fact, the speeches in our sample contain rather few outright policy statements in general. When such statements were made, they were either (a) consistent across the primary and general election, or (b) they were only discussed in one stage or the other. For examples of (a), across the primary and general elections of 2008, Barack Obama consistently advocated for a US$4,000 tax credit to make college affordable, for expanding the Child Tax Credit, for eliminating insurance discrimination for preexisting health conditions, and for increasing mandatory fuel standards for automobiles. And across the primary and general elections of 2012, Mitt Romney consistently called for repealing the Affordable Care Act, for deregulating the coal mining industry, for increasing limits and incentives for high-skilled immigration, and for school vouchers.

For examples of (b), we can consider John McCain's "gas tax holiday," which he first proposed a month after securing the Republican nomination as a solution to rising gas prices. McCain also proposed a tax credit for individuals who purchase zero emission vehicles in the general election, but had spoken of no such proposal during the primary election. Obama similarly proposed raising the cap on the Social Security tax during the 2008 general election, but had not voiced the proposal during the primary. Both Romney and McCain endorsed a temporary worker visa program during their respective general election campaigns, but did not address this policy during the Republican primaries.

Thus, to the extent that candidates indicate policy shifts in response to a changing electorate, it occurs only in the proposal of new policy ideas, not changing existing commitments. Yet in our evaluation of the corpus, we neither find this to be prevalent—happening only a few times between the three candidates—nor does it consistently point to moderation. Proposing temporary worker programs appears more centrist than conservative; neither gas tax holidays nor Obama's Social Security tax appears more moderate than liberal.

Our second observation concerns the nature of moderation in language. We note a consistent shift in how issues are discussed by candidates between primary and general campaigns, and some evidence of shifting issue

attention. Mitt Romney, for instance, discussed abortion many times, and with conviction, during the 2012 Republican primary: advocating the Hyde Amendment, the overturning of *Roe v. Wade*, and cutting funding to Planned Parenthood. During the general election, Romney was largely silent on this score, making only two statements on abortion: once praising his vice presidential nominee Paul Ryan for his position on the "sanctity of every human life," and once more vowing to "protect the sanctity of life" during his address to the Republican National Convention, an audience composed heavily of members of the partisan base.

John McCain shows a similar pattern when speaking on immigration during the 2008 cycle. During the primary, McCain prefaced immigration statements by asserting that America needed to "secure our borders first," and only a couple of times did he refer to his previous bipartisan comprehensive immigration reform proposals.[15] In the general election, McCain maintained his position on border security, even explicitly telling CNN anchor Wolf Blitzer that he has "not changed [his] position" on immigration. Yet during the general election, McCain's explanation did shift toward sympathy, with McCain making new references to immigrants as "God's children" seeking a better life in the United States.

This pattern, in fact, holds across many issues and all three candidates. Romney spoke often of taxes during the 2012 election, but mostly of "cutting taxes" during the primary, and of not "raising taxes" during the general. McCain frequently vowed to make the "Bush tax cuts permanent" while campaigning during the Republican primary, but never used the phrase during his general election speeches, preferring instead to refer only generically to keeping tax cuts in place. Obama maintained a pro-union stance during the 2008 election, and he spoke extensively of his support of unions during the primary. During the general election, however, Obama only vaguely referred to unions, repeating the line that Americans should not scapegoat groups for America's problems: "not welfare recipients, not corporations, not unions, not immigrants, not gays."

To summarize, then, our deeper qualitative exploration of presidential campaign rhetoric affirms what we found in using sophisticated computational methods. We find strong evidence for post-primary moderation in presidential campaigns. Unlike the traditional definition, however, we find no evidence of policy flip-flopping. Instead, we find evidence of *rhetorical* moderation. Candidates, presented with different electorates at different stages of the campaign, may be constrained in their ability to change policy positions, but they can, and do, adapt the language they use to explain those policies to voters. That shift appears to happen in two ways: by changing focus away from "red meat" issues for party bases, by and reframing positions to make them more palatable to a general electorate.

As we noted at the outset, these findings are normatively troubling. In a democracy, voters are charged with selecting the candidate who will best represent their ideological (as well as other) interests. And one significant way that voters learn about candidates' ideology is through candidate rhetoric. Our evidence that candidates retain relatively fixed policy positions but morph their language, chameleon-like, from appearing more extreme to appearing more moderate means that voters are, at least at some point during the campaign, receiving misleading cues. Are these cues strong enough to influence vote choice? Probably not, in the case of most voters. But they matter, nonetheless, if we are led to believe that a candidate's judgment is similar to our own and subsequently surprised to learn it is not. And in the modern era, where it seems candidates are increasingly able to "get away" with communicating less substance and more generalities via outlets such as Twitter, we contend that our findings in support of the rhetorical moderation theory matter now more than ever.

This research project moves American political research forward in two major ways. First, we provide the first systematic evidence for the academic and popular intuition behind post-primary moderation. By employing text analytic tools, we gain leverage on a claim that—while obvious according to folk wisdom—has evaded systematic empirical measurement for decades. Second, and perhaps more excitingly, we open up a new and fruitful area for future study. It remains an open question how well voters may, or may not, perceive subtle rhetorical differences. As with any campaign message, the fact that candidates seem to carefully craft their ideological moderation may (or may not) mean the strategy has a significant effect on voters.

Declaration of Conflicting Interests

The author(s) declared no potential conflicts of interest with respect to the research, authorship, and/or publication of this article.

Funding

The author(s) disclosed receipt of the following financial support for the research, authorship, and/or publication of this article: This research was supported by National Science Foundation grants #1211201, #1211266, and #1211277, under the Social-Computational Systems program (SoCS).

Supplemental Material

Supplemental material for this article is available online.

Notes

1. The Etch-a-Sketch is a classic American toy, introduced in 1960, that allows users to "draw" pictures by turning two knobs that control vertical and horizontal movement of a stylus. The stylus works by scraping off aluminum shavings that coat the inside of the toy's screen. To erase a drawing and begin again, one turns the toy over and shakes it to recoat the screen.
2. Consider attention in 2016 Democratic primaries to which candidates were willing to say that "Black lives matter."
3. The most frequently confused were Religious Left and Religious Right, drawing as they did on a number of terms common to both and uncommon elsewhere.
4. These data management choices serve two purposes. First, they help us to reduce noise in the data by combining words with common stems (so that, for example, plurals and singulars are counted as the same rather than separate words), and by removing stop words that play important grammatical but little semantic roles in language, and thus do not aid us in measuring ideological content in speech. Second, we recognize the important meaning conveyed in multiword expressions that we would lose if we considered words in isolation. If we ignored phrases, we might break up expressions like "illegal immigration" into the unigrams "illegal" and "immigration." To do so would ignore the inferential power gained from knowing that an author routinely writes not just of illegality or of immigration, but specifically of illegal immigration.
5. One can also cast this as an constrained optimization problem with ℓ_1 regularization.
6. In earlier work, the current and other authors have referred to this simply as the Ideological Books Corpus.
7. We use the sparse additive generative model (SAGE) weights as priors, and add some smoothing so (a) any word has some nonzero probability of being emitted from all ideological states, but (b) words with nonzero SAGE weights for a particular ideological class or subclass will have a higher probability a priori of being emitted from that (sub)class.
8. In practice, this works well except when a speaker goes for a long time without using an ideological cue term. In such cases, the lag period appears large, and the time-splitting assumption inflates the importance of the previous and current ideological states. That is, long lags might mean that the speaker has simply switched to talking about nonideological topics, rather than that they have split their time between the previous and current state. To address this, we also introduce a reset parameter, which allows the model to reset the ideological state back to the background (i.e., nonideological) state. The probability of resetting increases with lag length, and the reset indicator is sampled at each (cue or lag) term emission.
9. The length of the sequence is assumed to be exogenous, so that no stop state needs to be defined.
10. To ensure that our results are not simply driven by our model selection, we also estimated candidates' use of ideological language with a simple bag-of-words

approach. The results from the finite mixture model tell the same substantive story as the ones we present in the "Results" section, though face validity suffers in this approach.

11. In our data, we have only a few self-described centrists and we estimate very few tokens that distinguish the Centrist class from the others. In the data, what seems to distinguish the Centrist class is not a unique rhetorical style, but rather a tendency to borrow language from both the left and right classes.
12. Recall that these are not "positions" but rather a visual representation classification proportions.
13. As Figure 3 shows most clearly, John McCain makes for an interesting example of the rhetorical moderation hypothesis in practice. Not only did McCain make a sizable shift toward the center, but he also increased his use of centrist phrases (posterior probability ≈ 0.98), which is why his liberal and conservative proportions clearly do not sum to one.
14. See http://www.cs.cmu.edu/~ark/CLIP/
15. However, McCain did go so far as to say, in an interview with Tom Brokaw, that he had "stood up against my party, not just President Bush, but others; and I've got the scars to prove it, including taking up, with Ted Kennedy, immigration reform, knowing full well that that was going to hurt my chances in the primaries."

References

Agranov, M. (2011). *Flip-flopping, intense primaries and the selection of candidates*. Retrieved from https://papers.ssrn.com/sol3/papers.cfm?abstract_id=2412252

Bawn, K., Cohen, M., Karol, D., Masket, S., Noel, H., & Zaller, J. (2012). A theory of political parties: Groups, policy demands and nominations in American politics. *Perspectives on Politics*, *10*, 571-597.

Black, D. (1948). On the rationale of group decision-making. *The Journal of Political Economy*, *56*, 23-34.

Brady, D. W., Han, H., & Pope, J. C. (2007). Primary elections and candidate ideology: Out of step with the primary electorate? *Legislative Studies Quarterly*, *32*, 79-105.

Burden, B. C. (2001). The polarizing effects of congressional primaries. In P.F. Galderisi, M. Ezra, and M. Lyons (Eds.) (pp. 95-115). *Congressional primaries and the politics of representation*. New York: Rowman & Littlefield Publishers.

Burden, B. C. (2004). Candidate positioning in us congressional elections. *British Journal of Political Science*, *34*, 211-227.

Converse, P. (1964). The nature of belief systems in mass publics. *Critical Review*, *18*(1-3), 1-74.

Cox, G. W. (1990). Centripetal and centrifugal incentives in electoral systems. *American Journal of Political Science*, *34*, 903-935.

Diermeier, D., Godbout, J.-F., Yu, B., & Kaufmann, S. (2012). Language and ideology in congress. *British Journal of Political Science*, *42*, 31-55.

Downs, A. (1957). *An economic theory of democracy*. New York, NY: Harper & Row.

Druckman, J. N., & Holmes, J. W. (2004). Does presidential rhetoric matter? Priming and presidential approval. *Presidential Studies Quarterly*, *34*, 755-788.

Eisenstein, J., Ahmed, A., & Xing, E. P. (2011). Sparse additive generative models of text. In L. Getoor and T. Scheffer (Eds.), ICML'11 *Proceedings of the 28th International Conference on Machine Learning*, Bellevue, Washington, USA. (pp. 1041-1048).

Freeden, M. (2003). *Ideology: A very short introduction*. Oxford, UK: Oxford University Press.

Greene, B. (1973). *Running: A Nixon-McGovern campaign journal*. Chicago, IL: Henry Regnery.

Hart, R. P. (2000). *Campaign talk: Why elections are good for us*. Princeton, NJ: Princeton University Press.

Hotelling, H. (1929). Stability in competition. *The Economic Journal*, *39*(153), 41-57.

Hummel, P. (2010). Flip-flopping from primaries to general elections. *Journal of Public Economics*, *94*, 1020-1027.

Iyengar, S., & Simon, A. F. (2000). New perspectives and evidence on political communication and campaign effects. *Annual Review of Psychology*, *51*, 149-169.

King, G., & Lowe, W. (2003). An automated information extraction tool for international conflict data with performance as good as human coders: A rare events evaluation design. *International Organization*, *57*, 617-642.

Lane, R. E. (1962). Political ideology: Why the American common man believes what he does. New York: The Free Press.

Laver, M., Benoit, K., & Garry, J. (2003). Extracting policy positions from political texts using words as data. *American Political Science Review*, *97*, 311-331.

Lowe, W., Benoit, K., Mikhaylov, S., & Laver, M. (2011). Scaling policy preferences from coded political texts. *Legislative Studies Quarterly*, *36*, 123-155.

Meirowitz, A. (2005). Informational party primaries and strategic ambiguity. *Journal of Theoretical Politics*, *17*, 107-136.

Monroe, B. L., Colaresi, M. P., & Quinn, K. M. (2008). Fightin'words: Lexical feature selection and evaluation for identifying the content of political conflict. *Political Analysis*, *16*, 372-403.

Monroe, B. L., & Schrodt, P. A. (2008). Introduction to the special issue: The statistical analysis of political text. *Political Analysis*, *16*, 351-355.

Nelson, T. E. (2004). Policy goals, public rhetoric, and political attitudes. *Journal of Politics*, *66*, 581-605.

Powell, G. B. (2000). *Elections as instruments of democracy: Majoritarian and proportional visions* (Vol. 1). New Haven, CT: Yale University Press.

Sim, Y., Acree, B., Gross, J. H., & Smith, N. A. (2013). Measuring ideological proportions in political speeches. In *Proceedings of the 2013 Conference on Empirical Methods in Natural Language Processing*, Seattle, Washington, USA. The Association for Computational Linguistics.

Slapin, J. B., & Proksch, S.O. (2008). A scaling model for estimating time-series party positions from texts. *American Journal of Political Science*, *52*, 705-722.

Tomz, M., & Van Houweling, R. P. (2010). *Candidate repositioning*. Unpublished Manuscript.

Tomz, M., & Van Houweling, R. P. (2016). Political repositioning: A conjoint analysis. Unpublished Manuscript.

Author Biographies

Brice D.L. Acree is an assistant professor of Political Science and Translational Data Analytics at the Ohio State University. He studies computational methods for natural language processing.

Justin H. Gross is an assistant professor of Political Science and Core Faculty of the Computational Social Science Institute at UMass, Amherst. His research interests include political communication in mass and social media, ideology as pop philosophy, and Latino politics, as well as problems of measurement, networks, and text analysis. He is the author of several articles, appearing in such journals as the *American Journal of Political Science, British Journal of Political Science*, and *Social Networks*, among others.

Noah A. Smith is a professor of Computer Science & Engineering at the University of Washington and Senior Research Manager at the Allen Institute for Artificial Intelligence. He designs algorithms for natural language processing and its many applications. He is the author of *Linguistic Structure Prediction* (Morgan & Claypool), co-mentor of the student team that won Amazon's inaugural Alexa Prize competition, and winner of several conference paper awards.

Yanchuan Sim earned his PhD in Language and Information Technologies at Carnegie Mellon University. He currently researches novel techniques for improving state-of-the-art NLP systems at the Institute for Infocomm Research.

Amber E. Boydtun is an associate professor of Political Science and a Chancellor's Fellow at the University of California, Davis. She uses lab experiments, large-scale media studies, and manual and computational text analysis to study how issues make the news, the dynamics of "media storms," and how media attention shapes public opinion. She is author of *Making the News* (University of Chicago Press) and co-author of *The Decline of the Death Penalty and the Discovery of Innocence* (Cambridge University Press), as well as many journal articles.