

Are Voters' Preferences Being Ignored by Governor's Agendas?

Daniel M. Butler* Benjamin S. Noble[†]

February 24, 2025

Abstract

Co-partisan governors increasingly pursue similar agendas, and partisan control of government highly predicts government policy. Some attribute these trends to elite coordination at the expense of constituency preferences. We argue these accounts ignore the role of voters. Governors have incentives to represent constituency preferences; and partisan agendas may simply mirror governors' nationalized and polarized electorates. We measure the partisanship of nearly 2,500 State of the State Addresses (1962–2023) to evaluate whether governors' agendas are reflective of differences in constituents' preferences. We fine-tune a series of BERT models to predict the party of each held-out speech using other governors' addresses as training data. When a governor “sounds like” her co-partisans, the model more confidently predicts her party. We use this predicted probability as a measure of expressed partisanship and show that gubernatorial agendas have polarized, but this trend is increasingly correlated with constituency partisanship, not governor partisanship. Polarization in gubernatorial agendas has increased, but this shift reflects voter preferences rather than a departure from them.

*Professor, Department of Political Science, Washington University in St. Louis.

[†]Assistant Professor, Department of Political Science, University of California, San Diego.
benjamin noble.org b2noble@ucsd.edu.

In the spring of 2021, migrant crossings into the US at the southwest border increased dramatically. Prior to February 2021, Customs and Border Protection recorded about 70,000 encounters per month. In March of that year, encounters routinely hit or exceeded 170,000 per month.¹ In response, Republican governor Greg Abbot of Texas implemented a new policy that April in which his administration bused migrants to Democratic-led cities across the country. Soon after, Republican governor Ron DeSantis of Florida adopted a similar policy. The degree to which these governors, and others, adopted the policy raises a question about the degree to which partisanship and partisan bonds influence governors' policy agendas. As Butler and Sutherland (2022) show, governors agendas have nationalized, becoming increasingly similar since the 1970s. Similarly, Grumbach (2022, 51) argues that "the policy regime under which an individual lives is increasingly determined by her state of residence."

Why has policy similarity increased across states? Some argue that these trends are driven by elite-level co-partisan coordination. The strongest version of this claim holds that governors and state legislatures coordinate on co-partisan priorities irrespective of their voters' preferences (Grumbach 2022; Hertel-Fernandez 2019). Political elites—from governors, to interest groups, to activists—may leverage periods of unified state government and federal gridlock to implement their preferred policies, rather than the policy preferences of the public. Consistent with this argument, *POLITICO* reported that aides for Abbot and DeSantis communicated directly about implementing the busing policy in their respective states.² If elites pursue agendas that are driven primarily by their own partisan interests, a troubling implication is that governors may be pursuing policies their constituents *oppose*.

While elite-level partisanship and coordination surely drives some policy change in the states, we argue that gubernatorial representation is unlikely to be as dire as Grumbach (2022) and others fear. First, one consequence of nationalization is that gubernatorial

¹<https://www.cbp.gov/newsroom/stats/southwest-land-border-encounters>

²<https://www.politico.com/news/2023/07/19/desantis-abbott-immigration-00106902>

selection follows trends in presidential voting (Hopkins 2018), so governors typically represent co-partisan constituencies (Jacobson 2015). Republicans govern states in which a majority of voters are Republican, and Democrats govern states in which a majority of voters are Democrats. To the extent that party drives the state’s policy agenda, these policies are likely to be in-line with majority sentiment in the state. Consistent with this argument, a 2022 poll of Texans found that a narrow majority (52%) were supportive of Abbot’s busing policy.³ Second, governors, like other politicians, have incentives to pursue policies that align with voters’ preferences, at least some of the time (Besley and Case 2003; Tausanovitch and Warshaw 2013; Desmarais, Harden and Boehmke 2015). To the extent that governor’s private preferences are more ideologically extreme than their voters, electoral incentives can condition the degree to which they act against the will of their constituents (Canes-Wrone, Brady and Cogan 2002). Taken together, these propositions lead us to expect that a Republican governor in a state that is 85 percent Republican would pursue an agenda that is more “Republican” than a Republican governor representing a 55 percent Republican state. We are agnostic to whether these changes are driven by governors dynamically responding to voters or voters selecting congruent politicians, but, we do expect governors to reflect constituency preferences over time as the partisanship of the constituency changes. This expectation contrasts with that of Grumbach (2022) and others whose arguments suggest that agendas are invariant to the partisan composition of the state.

We evaluate gubernatorial representation by measuring and analyzing the partisanship of governor’s agendas as expressed in their State of the State Addresses (Butler and Sutherland 2022) between 1960 and 2023. These speeches, delivered near-annually across this period, provide a non-discretionary, consistent, and routine reflection of governors’ policy priorities, making them a reliable source for understanding elite partisanship across time. To measure a speech’s partisanship, we adopt a supervised learning ap-

³<https://texaspolitics.utexas.edu/blog/new-uttexas-politics-project-poll-abbott-maintains-45-40-lead-over-o%E2%80%99rourke-52-support-busing#Immigration%20and%20border%20issues>

proach and use classification accuracy as a substantive quantity of interest (Peterson and Spirling 2018). In short, we fine-tune a series of BERT models to predict the party label of a governor’s speech, using every other governor’s speech (in that, and the previous two years) as training data. Although we know the true label, this predicted probability provides information about “how Republican” a given speech is, as compared to other contemporaneous speeches. Intuitively, the more a governor “sounds like” other governors in her party, and the more governors within a party sound like one another, the more confident the classifier’s prediction of her partisanship. We conduct this classification task at the sentence-level and then aggregate these predictions to the speech level to explore the degree to which gubernatorial partisanship and constituency partisanship predict the partisan-lean of governors’ expressed agendas.

Our analysis reveals a positive correlation between voter preferences and the partisan-lean of gubernatorial agendas. This correlation holds when we control for a governor’s party, and is especially strong in the post-1994 period in which politics nationalized (Gentzkow, Shapiro and Taddy 2019; Hemmer 2022; Sinclair 2006) and partisanship became a more informative signal, especially for cross-state comparisons, of voter preferences. Although the governor’s own partisanship has always been positively associated with the partisan content of gubernatorial agendas, its influence has been constant across time, whereas the influence of constituency partisanship has *increased*, contrary to the dominant perspective in the literature. Our results suggest that the increase in the partisan polarization in state policy agendas has not occurred *in spite* of voters’ preferences, but could be a result of gubernatorial responsiveness to the preferences of well-sorted, polarized voters (Abramowitz 2018). Our findings raise normative concerns about mass-level nationalization and polarization, but should alleviate some concerns about whether voters’ preferences are ignored by their governors.

Nationalization, Partisanship, and State Policy Agendas

The content of the policy agenda is a major determinant of what governments enact (e.g., Cox and McCubbins 2005). Before governments can pass a policy, politicians must first propose, debate, and vote on a bill. From the perspective of the rank-and-file, the content of the bill is *fait accompli*. They cast an up-or-down vote, choosing between the status quo and an alternative bundle of changes. The same is not true, however, of coalition leaders—like presidents or governors—who often select from the set of all issues and structure subsequent alternatives for the rank-and-file (Arnold 1990). The agenda is crucial for shaping which policies are proposed, debated, and ultimately implemented (e.g., Jones and Baumgartner 2005; Light 1999). For this reason, lobbyists (Butler and Miller 2022), legislative leaders (Cox and McCubbins 2005; Smith 2007), and presidents (Canes-Wrone 2006) try to exercise positive and negative agenda control. Given this context, even if politicians only voted for bills with majority public support, partisan control of the agenda could still lead to partisan policy outcomes because Democratic and Republican leaders would propose different agendas. However, these agenda-setting decisions become even more important if politicians are willing to propose and enact partisan priorities that do not even receive support of the majority (Grumbach 2022).

While there is evidence that Republicans and Democrats pursue ideologically-aligned policies when they control state government (Caughey, Warshaw and Xu 2017; Fowler and Kettler 2021; Leigh 2008; Warshaw 2019), it's not clear whether these policies go against voters' preferences. One argument is that partisanship operates primarily through the agenda-setter. In this view, the coalition leader (e.g., governor) is chosen by voters, but that is where constituency influence ends. Once in office, the leader faithfully enacts party orthodoxy irrespective of her electorate's preferences (Lee, Moretti and Butler 2004). Put differently, whether a Democratic governor takes office with 51 or 91 percent of the vote, policy outcomes will be similar and the governor will not moderate based on the underlying distribution of constituency preferences. This effect may be compounded when

government officials are more interested in learning about policy ideas that have been advocated by co-partisans, but not opposition partisans, in other states (Butler et al. 2017). Grumbach (2022) is one prominent advocate of this perspective, arguing that “rather than a sea change in public opinion, the major policy changes of the era of state resurgence are the result of increasingly coordinated national networks of activists and organizations that make up the modern Democratic and Republican parties.” Normatively, this lack of representation would be troubling to the extent that governors are subverting the preferences of their constituents and these outcomes would “limit [state parties’] capacity to represent those state and local interests that fit uneasily within national coalitions” (Hopkins, Schickler and Azizi 2022, 4).

However, we argue that one implication of nationalization (e.g., Hopkins 2018) is that the pursuit of partisan policy agendas does not necessarily imply that voters are being rolled. Governors are incentivized, just like any other politician, to seek re-election by responding to constituency demands. Even term-limited governors may cater to voters’ preferences either to benefit co-partisan successors or lay the groundwork for a federal-level run for office (e.g., for a Senate seat or the presidency). Further, national factors, such as attitudes toward the president or opinions on salient national issues, are more predictive of vote choice and policy attitudes across the federal system (Carson, Sievert and Williamson 2023). State politics has generally shifted away from localized issues and toward salient national debates (Butler and Sutherland 2022; Hopkins, Schickler and Azizi 2022), and voters’ own attitudes towards these issues are more constrained and polarized (Abramowitz 2018; Mason 2018). Voters increasingly value issue-based representation (Costa 2020; Lapinski et al. 2016) and are consuming more national, and less local, news (Darr, Hitt and Dunaway 2018; Hopkins 2018; Martin and McCrain 2019). Incumbents, then, are much more likely to represent constituencies composed primarily of their own partisans (Jacobson 2015). The governor may simply be acting as a conduit for public opinion (Erikson, Wright and McIver 1993; Tausanovitch and Warshaw 2013; Des-

Table 1: Trust in Different Levels of Government, 2022

Trust the State	Trust the Feds			
	A great deal	A fair amount	Not very much	None at all
A great deal	1,605	672	385	233
A fair amount	1,822	9,753	3,486	1,057
Not very much	842	6,188	10,953	2,378
None at all	767	2,093	2,776	5,801

marais, Harden and Boehmke 2015). By pursuing more partisan agendas than they have in the past, governors may simply be delivering what voters at the state level want in our polarized and nationalized error of politics.

In support of this contention, voters express more trust in their state government than the federal government. In its 2022 wave, the Cooperative Election Study (CES) asked respondents to separately report their level of trust in their state and federal government for handling the nation's problems. Table 1 gives the joint distribution of responses. Most respondents (61.9 percent) trusted both levels of government equally. However, there were more people who trusted their state government more than the federal government (24.2 percent) than the reverse (13.9 percent). There are, of course, many reasons that individuals may trust their state government more than the federal government.⁴ However, one reason may well be that state government produces policies people like. At the very least, this is not the pattern we'd expect to see if state governments were actively subverting voter preferences.

Although we expect both a governor's own partisanship and that of her state to be correlated with the partisan-lean of her agenda across our period of study, we also expect the relative influence of forces to vary. Nationalization and polarization have varied across U.S. history (Carson, Sievert and Williamson 2023), and several scholars have identified the 1990s as a key period of change due to the Southern Realignment, increasing com-

⁴At the state level, there were 0 states where more respondents trust the national government more than the state government. There was only 1 state (Iowa) where the number of respondents who trust the national government equals the number who trust the state government. In the other 49 states, there are more respondents who trust the state government more than they trust the national government.

petitiveness of congressional elections, Newt Gingrich, the rise of an ideological activist class, and changes in the media environment (Butler and Miller 2022; Gentzkow, Shapiro and Taddy 2019; Hopkins, Schickler and Azizi 2022; Lee 2016; Sinclair 2006; Theriault and Rohde 2011). While these forces were national in scope, “state parties and national-level officials were responding simultaneously to the same changes in the broader political environment” (Hopkins, Schickler and Azizi 2022, 17). As such, we expect the influence of elite and mass partisanship to vary around this break-point, which define concretely as the pre- and post-1994 era.

Existing scholarship leads to competing expectations about how changes in the 1990s affected trends in gubernatorial agenda-setting. If elite-level dynamics are driving policy agendas at the expense of voter preferences, one could hypothesize that constituency was more important prior to 1995, whereas party became a driving force since then. Before the era of intense elite polarization, governors may have tried to align their agendas with constituency sentiment, believing their own re-election would be based more on candidate-specific and local factors rather than national partisan dynamics. However, after 1994, governors may have leaned more heavily on partisanship with the expectation that their electoral prospects would be better served by appealing to their partisan base. However, a voter-centric perspective leads to an alternate expectation. To the extent that voters were poorly sorted and weakly polarized prior to the 1990s (e.g., Abramowitz 2018), a voter’s national partisan identity might convey little about their preferences over *state* policies. Thus, a governor observing strong support for the Republican presidential candidate would infer little about what kind of local policies those voters desired. However, as voters sorted and began to embrace a more constrained partisan worldview (e.g., Mason 2018; Levendusky 2009), then national partisanship conveyed much more information to governors about the policies their electorate wanted. National-level partisanship, in this context, serves as a strong and informative signal of voter preferences at the state level. While this perspective is agnostic to the effects of elite partisanship

over time, it hypothesizes that constituency is *more* strongly predictive of gubernatorial agendas in the post-1994 period.⁵

Our primary purpose is to investigate whether variation in the partisan-lean of gubernatorial agendas is correlated with changes in voter partisanship. We do not argue, that governors are only following voter preferences; multiple influences could simultaneously be shaping politicians' actions (Caughey, Warshaw and Xu 2017). We expect both elite and constituency partisanship to predict the partisan-content of governors' agendas, with the strength of that correlation possibly being different across time. And even if we observe the expected correlation, we cannot answer whether it is driven by direct responsiveness, selection of more aligned agenda-setters, or is simply incidental. Yet our question is extremely important for learning whether observed increase in state polarization means that voters are having their preferences trampled on as some researchers argue.

Measuring Partisanship in Gubernatorial Agendas

To test these competing expectations, we focus on the policy agendas of governors as they are increasingly important actors in American politics (Birkhead, Harden and Windett 2024). Polarization and gridlock at the federal level (Sinclair 2006; Lee 2016) have allowed state governments, particularly governors, to take a more prominent role in policy-making. As the chief executive, the governor is one of the most well-known state politicians. Further, their agenda can have an outsized influence on state politics as state legislators rarely have the resources to properly consider all bills introduced in a year and thus want to allocate time to those bills which stand the best chance of passing. Lorenz (2020) argues that organized interests can provide legislators with information about bill

⁵To be clear, we do not argue that the 1994 Republican Revolution *caused* a change in gubernatorial agenda setting, but rather, the changes in the political and media environment affecting national and state politics during this period made it more likely for governors to pursue polarized policy agendas.

viability through lobbying; the governor can play a similar role by focusing attention on particular issues. Further, because many state legislatures are under-resourced, there is an informational imbalance in the governor’s favor (Kousser and Phillips 2012), and in many states governors can directly influence policy through vetoes or unilateral action (Bolton and Thrower 2021). Thus, we focus our attention on partisanship in governors agendas given the influence of their office and platform in driving state policy outcomes.

In our empirical tests, we focus on governor’s expressed agendas as presented in their annual State of the State Addresses. The State of the State (SOTS) Address, similar to the president’s State of the Union Address, is a major policy speech given by the governor to the legislature either annually or biannually. Governors typically use this speech to recap their recent accomplishments and to outline their policy agenda for the coming year. As this speech is mandated by most state constitutions (and given regularly by tradition in those where it is not mandated), a governor’s choice to speak is not driven by features of the political environment. Thus the act of giving the SOTS can be treated as exogenous to the political environment. In our study, we leverage a corpus of SOTS Addresses collected by Butler and Sutherland (2022), beginning in 1960 and ending in 2016, which we extend through 2023. Our final analysis corpus contains 2,427 speeches given by 470 governors across all 50 states.⁶

Classification Accuracy as a Measure of Partisanship

We follow Peterson and Spirling (2018) in using classification accuracy as a substantive quantity of interest. As we detail below, we train a set of machine learning models to predict the party label of the speaker of each sentence in our corpus. This predicted

⁶This corpus does not include the theoretical maximum of 3,200 speeches because governors in some states (e.g., Texas) give bi-annual speeches, missingness in the data collection process, exclusions due to OCR processing challenges, and that we subset only to Republican and Democratic governors. Additionally, governors may give multiple speeches in a single year that could be considered a State of the State Address (e.g., a governor gives a budget and policy speech in the same year). When we encounter these, we include them as separate speeches in our data. In Figure A1, we plot a grid of state-years for which we have at least one speech.

probability provides a sentence-level measure of the model's confidence that the speaker is a Republican or Democrat. Probabilities closer to 1 or 0 indicate that the model is more certain about the speaker's party identification given the text of that sentence. Probabilities closer to 0.5 indicate the model is less certain about the speaker's partisan identity based on their spoken text. We choose the sentence as our unit of analysis given the token limitations of our machine learning models and our expectation that the degree of partisanship will vary within speeches. Of course, whether these predicted probabilities capture partisanship depends on the data used to train our machine learning models—i.e., the speeches given by other governors.

Many parts of a governors' SOTS will focus on topics that are not polarized. This can happen because they focus on topics where the two parties have similar positions, such as support for infrastructure spending. It can also happen when a governor talks about an issue specific to their state. For example, the government response to a natural disaster in the previous year. Finally, it may happen because the governor's sentence is non-substantive and commonly expressed (e.g., welcoming legislators). When the model encounters a new sentence about infrastructure spending, a local natural disaster, or a non-substantive and common expression, it will not be confident in its prediction, yielding a predicted probability close to 0.5. When we aggregate up from the sentence level to the speech level, we aggregate across sentences, many of which are not necessarily ideological or partisan. As such, the observed variation in predicted probabilities may be narrow and close to 0.5 on average.

However, as long as governors are at least somewhat polarized by party, the ideological portions of their expressed agendas should be similar to those of co-partisan governors and different from those of out-partisan governors (McCarty 2019). The larger this partisan divergence in what governors talk about or how governors talk, the more confident the model will be when predicting the party label of the speaker. Consider a hypothetical world in which Democratic governors only talk about abortion rights and Republicans

only talk about immigration. If the model encounters an unseen sentence about abortion rights, it will confidently predict that the sentence was spoken by a Democrat. Alternatively, suppose in an immigration context, Republicans only refer to “illegal aliens” while Democrats only use the phrase “undocumented immigrants.” If the model encounters a sentence about “illegal aliens,” it will confidently predict that the speaker is a Republican given the underlying distribution of lexical features.

A key implication, then, is that prediction errors (i.e., predicting Republican when the true speaker is a Democrat) are informative. For example, the more a Democratic Governor uses the phrase “illegal alien” despite its association with the Republican party, the more likely the model will be to classify this governor as a Republican. The more a governor speaks like their co-partisans, the more confidently the model will place them among their co-partisans. The more a given governor speaks like the other party while her co-partisans do not, the more likely the model will be to mis-classify their partisanship. To the extent that the training data is representative (more on this below), these predicted probabilities provide a continuous measure of the degree to which a governor’s expressed agenda is “Democratic” or “Republican.” Given the literature on nationalization and polarization (Butler and Sutherland 2022; Hopkins 2018; Gentzkow, Shapiro and Taddy 2019), we would expect the model to make more accurate predictions, indicative of increasing polarization, as we get farther from the 1960s and closer to the 2020s.

Predicting Partisanship with BERT Classification

We use a multi-step process to generate our predicted probabilities. We begin by selecting BERT-base, a transformer-based model, as our starting point.⁷ This model, developed and trained by Google in 2018, generates contextual embeddings by representing tokens (e.g., words or word sub-pieces) as dense vectors in a common, multi-dimensional

⁷We tested several different approaches and found that BERT maximized classification accuracy across our corpus and provided more qualitatively sensible classification at the sentence level.

space. Tokens that appear closer to each other in this space are more similar to one another. Unlike dictionary methods or static embeddings (e.g., word2vec), these embeddings are sensitive to the context in which the token appears. For example, the token for “immigrant” will be represented by a different embedding if it is preceded by the token “undocumented” or “illegal.” If we suspect that Democrats and Republicans discuss immigration (or any topic) using different contextual language (e.g., Rodriguez, Spirling and Stewart 2023), this model is well suited to leverage these differences when making its predictions.

The BERT-base model is trained on English Wikipedia articles and books. This pre-training process gives the model a starting point to build from—which is especially helpful given our small corpus of State of the State Addresses. While, the content of gubernatorial speeches differs significantly from the language we would find in Wikipedia and books, BERT is flexible and open-source, allowing researchers to fine-tune the model and improve performance for their specific context. We conduct unsupervised fine-tuning in which the model is exposed to distinctly political text [specifically speeches from the Congressional Record]](Gentzkow, Shapiro and Taddy 2019). This step enables the model to learn relationships between political concepts and words that may not have appeared in its original training data and allows it to better interpret gubernatorial speeches. For example, in this step, it might strengthen the association between the word “alien” and “immigration” rather than “Roswell, New Mexico.” Importantly, the model is *not* learning relationships between text and party identification at this stage, but rather, contextual patterns in political speech. Although this step is not strictly necessary, it helps the model better understand gubernatorial speeches in the subsequent step. As political language changes over time (and this is substantive point of interest in our research), we fine-tune a series of BERT base models, decade-by-decade, allowing for more flexibility and dynamism in modeling these contextual relationships.⁸ For example, we fine-tune a model

⁸The choice of decades is convenient but arbitrary. One could conceivably pre-train candidate models on more fine-grained or coarser time periods. Given the degree to which political language changes over

on sentences from speeches in the congressional record of greater than 30 words between 1960 and 1969.⁹ We do the same for all remaining decades between 1970 and 2010.

In the third and final step, we conduct a leave-one-out (LOO) supervised learning task (see also Das et al. 2022) in which the relevant fine-tuned model is further trained on gubernatorial speeches in a supervised context—associating sentences with party labels—and used to predict the party label of each sentence in a held out speech.¹⁰ Specifically, to predict the party label of sentences in speech i given by governor g in year t , we begin with the relevant fine-tuned decade model.¹¹ We then train this model to associate party labels with sentences of all speeches given by governors $\neg g$ in year t and all speeches given by all governors in years $t - 1$ and $t - 2$.¹² Finally, we use this model to predict the party label of each sentence in speech i given by governor g in year t , which was not included in the training process. As a concrete example, to predict the party label of each sentence spoken by Governor Chris Christie (R-NJ) in 2015, we begin with the 2010s fine-tuned model. We then further train this model on all speeches given in 2013 and 2014 as well as all speeches in 2015, excluding the speech given by Christie in 2015 (which we are going to predict). We then use this newly trained model to predict the party label of each sentence spoken by Christie in 2015. We discard this model and begin with a clean version of the 2010s model to predict labels for the next governor’s sentences. We repeat

time versus the computational time and resources needed to fine-tune these models, decades strike a fair balance.

⁹We exclude speeches of fewer than 30 words as they are primarily parliamentary or non-substantive.

¹⁰We exclude all non-major party governors from training and prediction. We also exclude a small number of speeches sourced from YouTube transcripts that were lacking punctuation, which is important for our sentence-splitting algorithm. We conduct minimal pre-processing. The major steps we take are lowercasing and replacing state, city, and major place names with a placeholder [place] to prevent the model from inferring partisanship from location names.

¹¹Our corpus of gubernatorial speeches ends in 2023, but our Congressional Record corpus ends in 2016. Thus, the 2010 model is trained only on floor speeches given through 2016 and is used as the base model for all gubernatorial speeches given after 2009.

¹²The Butler and Sutherland (2022) dataset represents an important effort to digitize the complete set of gubernatorial State of the State Addresses. While we identified some transcription errors during our analysis, such errors are an inherent challenge in projects of this scale. Correcting these errors is beyond the scope of this article; however, we conduct robustness tests in Appendix A.2 and confirm that these errors have limited implications for our results.

this process for all speeches in all years between 1962 and 2023.¹³

Of course, governors are limited in how much they can say in an address and may select or avoid specific issues. Thus, we do not have access to some measure of governors’ “true” agendas or partisan proclivities that appear in other forums (e.g., social media) or in their hearts. What we do have is the full set of State of the State Addresses in most years, and thus, our training data represents the universe of governors’ mandated, expressed agendas in a given year. Intuitively, our LOO process asks, for each governor, “how much do you sound like peers in your own party versus the other party” conditional on what is said in governors’ State of the State Addresses in this year and the two previous years. In this sense, our training data is representative of the diversity of how governors express themselves in this particular speech. To the extent that parties speak differently, and governors speak more like their co-partisans, the model’s predictions should get increasingly accurate about the governors’ partisanship.

Gubernatorial Partisanship Across States and Over Time

At the end of this process, we have a matrix of sentence-prediction pairs for all sentences spoken by all governors in our corpus between 1962 and 2023. These predictions can theoretically range from 0 (most Democratic) to 1 (most Republican), but empirically, they range from 0.055 to 0.962. In Table 2, we present sentences and sentence excerpts from 2023 State of the State Addresses along with the probability the speech was given by a Republican. In the top portion of the table, the five sentences least likely to be given by Republicans concern conventional Democratic owned issues: climate change and programs to address poverty and homelessness. In the bottom portion of the table, the five sentences most likely to be given by Republicans include parental rights in education and law enforcement. Sentences with 0.5 probability are generally non-partisan, non-

¹³Although our corpus includes speeches beginning in 1960, our analysis uses only speeches given in 1962 and after as we need two previous years of speeches to train our models.

Table 2: Most Republican and Democratic Sentences, 2023

State	Sentence	Pr(Rep.)
VT (R)	this also means keeping our communities and families in mind as we address climate change.	0.11
NM (D)	we know that these fires were intensified by extreme weather brought on by climate change, and since the beginning of my administration, we have taken that threat seriously—embracing science in order to mitigate climate change’s devastating impact, protect our environment and ramp up clean energy production.	0.15
WI (D)	we created the state’s first-ever plan to respond to and mitigate the effects of climate change.	0.15
DE (D)	my budget next week will build on these investments in affordable housing and community development.	0.15
WA (D)	i’ say it again: until we fix our housing crisis, thousands of people will remain homeless..	0.16
SD (R)	i promised to bring the next big industry to [place].	0.50
CA (D)	now, the fourth largest economy in the world.	0.50
WI (D)	i’m tony evers, and i’m proud to be standing here tonight as the 46th governor of the great state of [place] to deliver my fifth state of the state address.	0.50
MO (R)	this year, we are again funding the program with an additional \$32 million dollars to continue the state’s part and benefit more [place] teachers.	0.82
NV (R)	for the first time, parents will have an advocate inside government promoting the expansion of school choice in [place].	0.82
NE (R)	trooper sutton, please stand and be recognized.	0.82
MI (R)	through the parents’ bill of rights, we will reaffirm that in [place], it is the state who answers to parents and not vice versa.	0.83
MI (R)	this parents’ bill of rights would further cement that when it comes to the usage of names, pronouns, or health matters, schools will adhere to the will of parents.	0.84

substantive, or are not indicative of party without additional context. Table 2 provides some face validity to our measure of the partisan content of governors’ speeches.

Although sentences provide rich data to understand micro-variation within and across speeches, our key quantity of interest is speech-level partisanship. We expect governors to sometimes speak like members of the other party, but for the most part, speak like their co-partisans. So, we ask: across everything a governor says in their State of the State Address, how Republican or Democratic do they sound? To construct this measure of the

partisan-lean of their speech, we simply aggregate sentences to speech-level by taking the mean of the distribution of sentence-level predictions. Thus, for each governor, we have a measure of the average prediction, which we use as a measure of the partisan-lean of their expressed agenda. Those whose mean prediction is at or above 0.5 are classified as Republican whereas those below are classified as Democrat. The most basic feature of interest is classification accuracy: how well does the model identify true Republicans and true Democrats? Note: our goal is not to *accurately* classify each speech. Rather, classification accuracy provides insight into the degree to which governors sound like members of their party and tell us something about the annual level of polarization—where higher accuracy indicates greater partisan polarization in speech.

To provide some insight at the speech level, we present the rank ordered probability a governor is a Republican in 2023 in Table 3. In line with what we would expect about gubernatorial polarization in 2023, we see that our measure almost perfectly clusters Republican and Democratic governors by party. Democratic governors tend to have predictions below 0.5 while Republicans have predictions above 0.5. The overall accuracy in 2023 is 86%. Deviations are also informative. For instance, Phil Scott (R-VT), a Republican, is actually the seventh most Democratic-sounding governor in 2023—consistent with Vermont’s democratic lean. Other governors representing out-party states, such as Beshear (D-KY) and Youngkin (R-VA) are right at the threshold. On the other hand, governors representing states that lean heavily toward their party such as Hochul (D-NY) and Bergum (R-ND) gave speeches where the content was most strongly associated with their party. Overall, these rankings accord with our general understanding of party polarization in 2023, suggesting that our approach performs as designed. Further, the degree to which governors representing heavily partisan states are predicted to be more partisan than those representing out-party states on average suggests some role for constituency influence.

Across six decades of gubernatorial addresses, our model accurately classifies 70% of

Table 3: Ordered Probability that Governor is Republican, 2023


State	Governor	Party	R. Vote	Pr(R)	Pr(R) (Std. Dev.)	State	Governor	Party	R. Vote	Pr(R)	Pr(R) (Std. Dev.)
NY	Hochul	D	38.30	0.41	-1.68	AR	Sanders	R	64.20	0.49	0.16
MN	Walz	D	46.37	0.43	-1.37	VA	Youngkin	R	44.85	0.50	0.21
CT	Lamont	D	39.80	0.43	-1.24	OH	Dewine	R	54.11	0.50	0.25
WI	Evers	D	49.64	0.43	-1.18	MD	Moore	D	32.99	0.50	0.34
MA	Healey	D	32.86	0.44	-1.10	AK	Dunleavy	R	55.23	0.51	0.47
PA	Shapiro	D	49.39	0.44	-1.08	GA	Kemp	R	49.85	0.51	0.58
VT	Scott	R	31.71	0.44	-1.06	NH	Sununu	R	46.28	0.51	0.58
HI	Green	D	35.00	0.44	-1.01	WV	Justice	R	69.79	0.51	0.66
NM	Grisham	D	44.48	0.44	-0.95	MS	Reeves	R	58.36	0.52	0.73
RI	Mckee	D	39.39	0.45	-0.93	ID	Little	R	65.84	0.52	0.83
MI	Whitmer	D	48.58	0.45	-0.85	OK	Stitt	R	66.94	0.52	0.85
NC	Cooper	D	50.66	0.45	-0.75	TN	Lee	R	61.81	0.52	0.87
WA	Inslee	D	40.08	0.45	-0.73	IN	Holcomb	R	58.16	0.53	0.89
AZ	Hobbs	D	49.85	0.46	-0.52	TX	Abbott	R	52.84	0.53	0.92
CA	Newsom	D	35.07	0.47	-0.43	IA	Reynolds	R	54.18	0.53	1.01
IL	Pritzker	D	41.39	0.47	-0.43	NV	Lombardo	R	48.77	0.53	1.05
CO	Polis	D	43.06	0.47	-0.43	SD	Noem	R	63.45	0.53	1.08
NJ	Murphy	D	41.95	0.47	-0.39	NE	Pillen	R	59.75	0.54	1.21
LA	Edwards	D	59.45	0.47	-0.36	AL	Ivey	R	62.88	0.54	1.32
DE	Carney	D	40.41	0.47	-0.35	FL	Desantis	R	51.66	0.55	1.42
KS	Kelly	D	57.46	0.47	-0.29	MO	Parson	R	57.84	0.55	1.52
ME	Mills	D	45.31	0.48	-0.17	SC	Mcmaster	R	55.94	0.56	1.62
UT	Cox	R	60.65	0.49	0.06	WY	Gordon	R	72.44	0.56	1.65
KY	Beshear	D	63.17	0.49	0.08	ND	Burgum	R	67.18	0.57	1.83
MT	Gianforte	R	58.36	0.49	0.14						

speeches. Overall, this classification accuracy is quite high given the limited training data, suggesting that both our models perform well and that governors are fairly polarized in their speech across time. Interestingly, the model is better at classifying Democrats than Republicans, suggesting that Democrats use more homogeneous speech features.¹⁴ There is, as expected, substantial heterogeneity across years.

In Figure 1, we plot classification accuracy (the y -axis) at the year level (the x -axis) to visualize this temporal heterogeneity. The size of the points indicates the total number of State of the State Addresses we have in our corpus that year. We see that classification accuracy is variable over time. It is higher earlier in the time series—in the 1960s and 1970s—then declines to its minimum during the 1980s. Significantly, the classification accuracy begins increasing in the 1990s, in line with existing literature on polarization occurring in the 1990s. The increase continues through the 2000s and 2010s—where it

¹⁴This pattern persists across different BERT specifications and even alternative modeling strategies, like random forest classification using a document-term matrix. This result suggests that there is something heterogeneous about Republican speeches, and this is not simply a bias of the model or training process.


Figure 1: Proportion of SOTS Correctly Predicted by Year



reaches its maximum.

Although accuracy is one way to understand polarization, a visual alternative in the literature is to assess whether the scores conform to a unimodal or bimodal distribution separated by party (McCarty 2019). In Figure 2, we plot the density of speech-level predictions by year. On the x -axis, we plot the mean-centered probability a speech was given by a Republican—where higher values indicate that the speech is more likely to be delivered by a Republican governor. We shade these density curves by party of the actual speaker, Democrats in blue and Republicans in red. We plot these curves by year, beginning with 1962 at the top of the figure and ending in 2023 at the bottom. Early in our time series, the distributions are either unimodal or contain substantial heterogeneity within party. However, speeches begin to polarize in the late 1990s. This polarization increases through the 2010s and into the 2020s, at which point, the distribution becomes increasingly bimodal and party-separated—indicative of strong rhetorical polarization.

Figure 2: Polarization in SOTS Speeches Over Time



Empirical Approach

For our main regression analysis the dependent variable is the Republican-lean of the speech (which is the aggregated, average predicted probability that each sentence in the speech was spoken by a Republican, which we then standardize). We standardize this variable to facilitate interpretation. In our empirical tests we look at how the dependent variable correlates with the governor's own partisanship and the state's presidential two-party Republican vote share (from Amlani and Algara 2021). If, as some argue (e.g., Grumbach 2022), elite-level partisanship influences gubernatorial agendas, we should observe a positive correlation between the predicted probability a governor is Republican and whether the governor is actually a Republican. If, as we have argued, governors' policy agendas are aligned with voters' interests, increases in the Republican two-party vote share should be correlated with increases in the probability as speech was given by a Republican. We use presidential vote share as a proxy for district partisanship because it is easily measured and is not endogenous to the governor's own electoral results (e.g., Grimmer 2013).

We also include state and year fixed effects. State fixed effects help account for time-invariant, state-specific features (e.g., region) that may lead governors to speak more or less like their fellow partisans. Year fixed effects isolate changes in gubernatorial partisanship from state-invariant, year-specific features of the political environment that could induce governors to speak more like their fellow partisans (e.g., the occurrence of a presidential election). We also account for state-varying and time-varying variables including whether it is an election year in the state, whether the governor is running for reelection and whether state government is unified (i.e., the same party controls the governor's mansion and both state legislative chambers).¹⁵

¹⁵These data come from (Klarner 2013) through 2011 and are updated by the authors thereafter.

Results

We present our results in Table 4, where our dependent variable is the probability that a speech is given by a Republican governor, which have been standardized. In column 1, we present coefficients for a naive party model—which includes information about a governor’s partisanship but excludes information about the state’s partisanship. As anticipated, the predicted probability that a speech was given by a Republican is higher when that governor is a Republican: 0.78 standard deviations higher as compared to a Democrat. Looking at the data from 2023 (see Table 3), that is the difference between the speech given by Andy Beshear (D-KY) and Brian Kemp (R-GA). This result is statistically and substantively significant.

In column 2, we test a naive voter model, in which we regress predicted probability on the state’s two party Republican vote share in the previous election (excluding governor partisanship). Here, we see a positive association between the probability a speech was given by a Republican and the state’s Republican vote. A 10-point change in Republican vote share (about one standard deviation) is associated with quarter of a standard deviation increase in the probability a speech was given by a Republican. In column three, we test a combined model, including both party and constituency as independent variables. In this model, the coefficients’ size and statistical significance do not change substantively and both maintain their statistical significance. Goodness of fit statistics, including adjusted R^2 , AIC and BIC all improve marginally over the naive party model, but these differences are not particularly substantive. However, either the naive party model or combined model seem to out-perform the naive voter model. Together, these results imply that party explains much of the variation in the partisanship of governor’s agendas between 1962 and 2023, however, models that exclude voter partisanship fit the data slightly worse.

In columns 4–6, we replicate these models, interacting our two key independent vari-

Table 4: Relationship Between Party, State Republican Vote, and Predicted Probability a Speech is Given by a Republican Governor

	(1)	(2)	(3)	(4)	(5)	(6)
Republican	0.777*** (0.083)		0.753*** (0.079)	0.645*** (0.114)		0.615*** (0.107)
Republican x Post-1994				0.255+ (0.129)		0.189 (0.134)
Two-Party Republican Vote Share (10s)		0.250*** (0.071)	0.153* (0.058)		-0.020 (0.052)	-0.004 (0.043)
Two-Party Rep. Vote (10s) x Post-1994					0.425*** (0.080)	0.248*** (0.066)
Election Year	0.066 (0.062)	0.073 (0.061)	0.075 (0.062)	0.067 (0.064)	0.077 (0.060)	0.078 (0.063)
Governor Running for Re-Election	-0.054 (0.077)	-0.063 (0.081)	-0.068 (0.078)	-0.053 (0.077)	-0.085 (0.080)	-0.079 (0.078)
Unified State Government	0.036 (0.063)	-0.073 (0.100)	0.026 (0.064)	0.020 (0.060)	-0.102 (0.100)	-0.008 (0.064)
Fixed Effects						
State	✓	✓	✓	✓	✓	✓
Year	✓	✓	✓	✓	✓	✓
Num.Obs.	2427	2427	2427	2427	2427	2427
R2 Adj.	0.334	0.217	0.339	0.337	0.245	0.351
R2 Within Adj.	0.165	0.020	0.172	0.169	0.055	0.187
AIC	6014.8	6404.8	5995.0	6004.5	6317.9	5954.8
BIC	6681.2	7071.1	6667.2	6676.7	6990.1	6638.5

+ $p < 0.1$, * $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

Note: Coefficients come from an ordinary least squares model where the dependent variable is the probability a speech is given by a Republican governor. Standard errors are clustered at the state level.

ables with an indicator for pre- and post-1994.¹⁶ We have competing expectations about how these coefficients could vary across time. Recall, theories of elite partisanship imply an increasing relationship between party and gubernatorial partisanship over time—which is what we see in model 4. Here, gubernatorial partisanship is always influential, but it is even more predictive of speech partisanship after 1994 (at $p < 0.1$). These same theories would predict a decreasing influence of voter partisanship over time as elites polarize. However, we suspect the influence of voter partisanship could *increase* over time given that national-level partisanship could provide a clearer signal of voter preferences given polarization and sorting. Our hypothesis receives support in column 5. An


¹⁶As we include year fixed effects, we necessarily exclude the constitutive term for post-1994.

increase in Republican vote share is not associated with an increase in Republican predicted probability before 1994 (and the coefficient is incorrectly signed). However, after 1994, it is a strong, and statistically significant predictor of Republican speech. Finally, we re-test the combined model with the interaction in column 6. These results are more in line with theories of voter influence than elite partisanship. The coefficient on gubernatorial partisanship in the post-1994 period loses statistical significance and shrinks in size. Meanwhile, the influence of voter partisanship continues to be strong and correctly signed in the post-1994 period. By all three goodness-of-fit metrics, this model best explains the variation in our data. It maximizes adjusted R^2 and minimizes AIC and BIC across the models we've fit. From these results, we conclude that theories excluding voter partisanship are missing part of the story.

To what extent is this post-1994 result simply proxying for some sort of secular increase over time versus a sharper change that occurred in the 1990s? To assess this possibility, we re-specify the model in column 6, interacting party and Republican vote share with a decade indicator (see Appendix B1). There, we see that across 1960–1990, a state's Republican vote share has no statistical influence on the partisanship of a governor's speech. However, beginning in 2000, Republican vote share has a positive and statistically significant effect as compared to 1960. This result is more consistent with a structural change around the mid-1990s than it is of a secular increase over time. In the appendix, we also split these results by party and show that what governors of both parties say is increasingly reflective of constituency partisanship in the post-1994 period.

Given the complicated interaction in column 6, interpreting effects directly from the table is not straightforward. As such, we compute the marginal effect of gubernatorial partisanship and voter partisanship (holding the other at its mean) pre- and post-1994 to interpret effect sizes and statistical significance in Figure 3. When we do this, we find that gubernatorial partisanship has a statistically significant relationship with the predicted probability in both periods: increasing it by 0.62 standard deviations before 1995 and

Figure 3: Marginal Effects of Governor Party and Constituency Partisanship on Probability Speech is Given By a Republican.



Note: Although party exerts a substantive influence on the party prediction of the speech, this effect does not vary over time. However, state partisanship is associated with predicted probability only in the post-1994 period. Marginal effects are generated from the model in column 6 of Table 4.

0.80 after. However, these two marginal effects are not statistically different from one another. While gubernatorial partisanship *does* correlate with the predicted probability, this effect does not vary over time. A governor's partisanship is no more influential on their expressed priorities throughout our time series. When we repeat the same exercise with voter partisanship, we find that Republican vote share does not correlate with speech partisanship before 1995 but does after 1994, an increase of about half a standard deviation shift per 20-point vote shift toward the Republican presidential candidate. This 20 point difference is close to the presidential vote share difference between the average Republican and Democratic governor in our time series. Taken together, these results show that both gubernatorial partisanship and voter partisanship explain variation in governor's expressed agendas. These results suggest that governors may not be rolling their voters so much as representing their polarized preferences.

Conclusion

Recent research on state politics has bemoaned the influence of elite-level partisanship on state policy outcomes, often arguing that it is the core determinant of the kind of policy regime under which voters live (Grumbach 2022; Caughey, Warshaw and Xu 2017). We agree that partisan officials enact partisan agendas, and we suspect that some of that is because elites coordinate on their policy agendas. However, we argue that voters also play a role. Voters select elites and can incentivize the actions that politicians take with the potential threat of electoral accountability. Most importantly, we argue that the observed polarized partisan outcomes at the state level are not necessarily inconsistent with what voters themselves want from their state officials. When governors enact partisan agendas, they may be aligning themselves with the partisan (polarized) majority of the voters in their state. Thus we expect that voter partisanship—even after controlling for the partisanship of the governor—is predictive of the partisan content of the governors’ expressed policy agendas.

Empirically, we find that *voter* partisanship is an important predictor of the partisan nature of the agendas that governors pursue. Across the time series, both gubernatorial and constituency partisanship are predictive of the kinds of agendas governors promote, and in the post-1994 period, constituency partisanship has a stronger correlation with governors’ agendas than it does in the pre-period. The correlation with governor partisanship is statistically static across these two time periods. This finding does not contradict the prior research about the importance of elite coordination and partisan policy diffusion—those things matter. Instead our main point is those effects are happening alongside an elite-level responsiveness to voters’ partisan leanings. Republicans like Charlie Baker (R) in Massachusetts, Larry Hogan (R) in Maryland articulated different agendas than what their Republican co-partisans in red-states did. And the same is true, on average, of governors who lead states with smaller partisan differences.

In this paper, we assume gubernatorial agendas are an intermediate step between

voter preferences and state policy outcomes. However, future research could more directly consider the links between what governors express in their State of the State Addresses and subsequent activity in state legislatures. Given that many studies focus on outputs of state government, it is theoretically possible (although we are skeptical) that while governors take actions that are responsive to public opinion, legislatures may not do so, instead producing policies based on pure partisanship. Further, future work could do more to disentangle the causal effects of state partisanship on gubernatorial agendas, and the degree to which these changes are driven by selection, responsiveness, or something else.

These findings point back to the important changes we have observed among voters more broadly. Voters have ideologically sorted into the party that more closely reflects their preferences (Levendusky 2009) and have turned their attention primarily to national issues and personalities at the expense of distinctly state or local issues (Hopkins 2018). Split-ticket voting has declined and the correlation between presidential vote and down ballot offices has increased (Carson, Sievert and Williamson 2023; Jacobson 2015). Thus, politicians tend to represent states where their own party is dominant and, as a consequence, promote agendas that are consistent with both their own, and their state's, partisan leaning. As party platforms nationalize (Hopkins, Schickler and Azizi 2022) and voters select politicians that share their party identity, policy agendas—and subsequent policies enacted at the state level—should look increasingly partisan. In other words, politicians pursuing more partisan agendas is not necessarily a sign that state government are acting against the will of the public. Instead, at least some of the divergence between red and blue state policy programs appears to simply be that voters in these states want different things.

References

- Abramowitz, Alan I. 2018. *The Great Alignment: Race, Party Transformation, and the Rise of Donald Trump*. New Haven: Yale University Press.
- Amlani, Sharif and Carlos Algara. 2021. "Partisanship nationalization in American elections: Evidence from presidential, senatorial, gubernatorial elections in the U.S. Counties, 1872–2020." *Electoral Studies* 73:102387.
- Arnold, R. Douglas. 1990. *The Logic of Congressional Action*. Yale University Press.
- Besley, Timothy and Anne Case. 2003. "Political Institutions and Policy Choices: Evidence from the United States." *Journal of Economic Literature* p. 67.
- Birkhead, Nathaniel A., Jeffrey J. Harden and Jason H. Windett. 2024. "Executive-Legislative Policymaking Under Crisis." *The Journal of Politics* .
- Bolton, Alexander and Sharece Thrower. 2021. *Checks in the Balance: Legislative Capacity and the Dynamics of Executive Power*. Princeton: Princeton University Press.
- Butler, Daniel M., Craig Volden, Adam M. Dynes and Boris Shor. 2017. "Ideology, Learning, and Policy Diffusion: Experimental Evidence." *American Journal of Political Science* 61(1):37–49.
- Butler, Daniel M. and David R. Miller. 2022. "Does Lobbying Affect Bill Advancement? Evidence from Three State Legislatures." *Political Research Quarterly* 75(3):547–561.
- Butler, Daniel M. and Joseph L. Sutherland. 2022. "Have State Policy Agendas Become More Nationalized?" *The Journal of Politics* p. 000–000.
- Canes-Wrone, Brandice. 2006. *Who Leads Whom?: Presidents, Policy, and the Public*. Chicago: University of Chicago Press.
- Canes-Wrone, Brandice, David W. Brady and John F. Cogan. 2002. "Out of Step, Out of Office: Electoral Accountability and House Members' Voting." *American Political Science Review* 96(01):127–140.
- Carson, Jamie L., Joel Sievert and Ryan D. Williamson. 2023. *Nationalized Politics: Evaluating Electoral Politics Across Time*. Oxford University Press.
- Caughey, Devin, Christopher Warshaw and Yiqing Xu. 2017. "Incremental Democracy: The Policy Effects of Partisan Control of State Government." *The Journal of Politics* 79(4):1342–1358.
- Costa, Mia. 2020. "Ideology, Not Affect: What Americans Want from Political Representation." *American Journal of Political Science* .
- Cox, Gary W. and Matthew D. McCubbins. 2005. *Setting the Agenda: Responsible Party Government in the U.S. House of Representatives*. Cambridge: Cambridge University Press.

- Darr, Joshua P, Matthew P Hitt and Johanna L Dunaway. 2018. "Newspaper Closures Polarize Voting Behavior." *Journal of Communication* 68(6):1007–1028.
- Das, Sanmay, Betsy Sinclair, Steven W. Webster and Hao Yan. 2022. "All (Mayoral) Politics Is Local?" *The Journal of Politics* 84(2):1021–1034.
- Desmarais, Bruce A., Jeffrey J. Harden and Frederick J. Boehmke. 2015. "Persistent Policy Pathways: Inferring Diffusion Networks in the American States." *American Political Science Review* 109(2):392–406.
- Erikson, Robert S., Gerald C. Wright and John P. McIver. 1993. *Statehouse Democracy: Public Opinion and Policy in the American States*. Cambridge University Press.
- Fowler, Luke and Jaclyn J. Kettler. 2021. "Are Republicans Bad for the Environment?" *State Politics Policy Quarterly* 21(2):195–219.
- Gentzkow, Matthew, Jesse M. Shapiro and Matt Taddy. 2019. "Measuring Group Differences in High-Dimensional Choices: Method and Application to Congressional Speech." *Econometrica* 87(4):1307–1340.
- Grimmer, Justin. 2013. "Appropriators not Position Takers: The Distorting Effects of Electoral Incentives on Congressional Representation." *American Journal of Political Science* 57(3):624–642.
- Grumbach, Jacob. 2022. *Laboratories Against Democracy: How National Parties Transformed State Politics*. Princeton: Princeton University Press.
- Hemmer, Nicole. 2022. *Partisans: The Conservative Revolutionaries Who Remade American Politics in the 1990s*. Basic Books.
- Hertel-Fernandez, Alexander. 2019. *State Capture: How Conservative Activists, Big Businesses, and Wealthy Donors Reshaped the American States—and the Nation*. Oxford University Press.
- Hopkins, Daniel J. 2018. *The Increasingly United States: How and Why American Political Behavior Nationalized*. Chicago: University of Chicago Press.
- Hopkins, Daniel J., Eric Schickler and David L. Azizi. 2022. "From Many Divides, One? The Polarization and Nationalization of American State Party Platforms, 1918–2017." *Studies in American Political Development* 36(1):1–20.
- Jacobson, Gary C. 2015. "It's Nothing Personal: The Decline of the Incumbency Advantage in US House Elections." *The Journal of Politics* 77(3):861–873.
URL: <https://www.journals.uchicago.edu/doi/10.1086/681670>
- Jones, Bryan D. and Frank R. Baumgartner. 2005. *The Politics of Attention: How Government Prioritizes Problems*. University of Chicago Press.
- Klarner, Carl. 2013. "State Partisan Balance Data, 1937 - 2011."
URL: <https://dataverse.harvard.edu/dataset.xhtml?persistentId=doi:10.7910/DVN/LZHMG3>

- Kousser, Thad and Justin H. Phillips. 2012. *The Power of American Governors: Winning on Budgets and Losing on Policy*. Cambridge University Press.
- Lapinski, John, Matt Levendusky, Ken Winneg and Kathleen Hall Jamieson. 2016. "What Do Citizens Want from Their Member of Congress?" *Political Research Quarterly* 69(3):535–545.
- Lee, David S., Enrico Moretti and Matthew J. Butler. 2004. "Do Voters Affect or Elect Policies? Evidence from the U. S. House." *The Quarterly Journal of Economics* 119(3):807–859.
- Lee, Frances E. 2016. *Insecure Majorities: Congress and the Perpetual Campaign*. Chicago: University of Chicago Press.
- Leigh, Andrew. 2008. "Estimating the impact of gubernatorial partisanship on policy settings and economic outcomes: A regression discontinuity approach." *European Journal of Political Economy* 24(1):256–268.
- Levendusky, Matthew. 2009. *The Partisan Sort: How Liberals Became Democrats and Conservatives Became Republicans*. University of Chicago Press.
- Light, Paul. 1999. *The President's Agenda: Domestic Policy Choice from Kennedy to Clinton*. The John Hopkins University Press.
- Lorenz, Geoffrey Miles. 2020. "Prioritized Interests: Diverse Lobbying Coalitions and Congressional Committee Agenda Setting." *The Journal of Politics* 82(1):225–240.
- Martin, Gregory J. and Joshua McCrain. 2019. "Local News and National Politics." *American Political Science Review* 113(2):372–384.
- Mason, Lilliana. 2018. *Uncivil Agreement: How Politics Became Our Identity*. Chicago, Illinois: University of Chicago Press.
- McCarty, Nolan. 2019. *Polarization: What Everyone Needs to Know*. New York, NY: Oxford University Press.
- Peterson, Andrew and Arthur Spirling. 2018. "Classification Accuracy as a Substantive Quantity of Interest: Measuring Polarization in Westminster Systems." *Political Analysis* 26(1):120–128.
- Rodriguez, Pedro L., Arthur Spirling and Brandon M. Stewart. 2023. "Embedding Regression: Models for Context-Specific Description and Inference." *American Political Science Review* 117(4):1255–1274.
- Sinclair, Barbara. 2006. *Party Wars: Polarization and the Politics of National Policy Making*. Norman: University of Oklahoma Press.
- Smith, Steven S. 2007. *Party Influence in Congress*. Cambridge: Cambridge University Press.

- Tausanovitch, Chris and Christopher Warshaw. 2013. "Measuring Constituent Policy Preferences in Congress, State Legislatures, and Cities." *The Journal of Politics* 75(2):330–342.
- Theriault, Sean M. and David W. Rohde. 2011. "The Gingrich Senators and Party Polarization in the U.S. Senate." *The Journal of Politics* 73(4):1011–1024.
- Warshaw, Christopher. 2019. "Local Elections and Representation in the United States." *Annual Review of Political Science* 22(1):461–479.

Supplementary Information

Paper Title

Contents


A	Data and Methods	1
A.1	Speech Availability and Missingness	1
A.2	Transcription Errors and Effects on Conclusions	2
B	Additional Models	3

A Data and Methods

A.1 Speech Availability and Missingness

Figure A1 visualizes missingness in the corpus (1962–2023). We use speeches from the Butler and Sutherland (2022) before 2017, and we extend the data through 2023. Our database covers 78% of all potential speeches, with gaps due to unavailable records as well as states (like Texas) where speeches may not be given in a particular year.

Figure A1: SOTS by Year in Data



A.2 Transcription Errors and Effects on Conclusions

In pre-processing the Butler and Sutherland (2022) speech corpus, we identified several speeches in which the OCR is, understandably, imperfect. Given the scope of the project, the diversity of formatting across state journals, and the difficulty of processing older scans, we identified several instances in which non-speech content is included (e.g., the Speaker introducing the governor) as well as instances where sentences may be spliced together due to irregular formatting.

In particular, the most worrying concern is that these transcription errors seem to be correlated with time. That is, older journals have more problematic formatting and lower quality scans than more recent journal scans or text scraped directly from websites on which speeches were hosted. This could impact our model to the extent that our unit of analysis is the sentence, and it is possible, then, that these transcription errors in the form of spliced sentences could be consequential. Given our descriptive conclusion that classification accuracy, and thus partisanship, has increased over time, is it possible that this conclusion is driven entirely by transcription error?

Although correcting these errors is beyond the scope of our project, we acknowledge that these transcription errors could be consequential for the conclusions we draw. To probe the sensitivity of our results, we use our manually collected and validated 2021–2023 set of State of the State addresses, which we are confident contain few, if any, transcription errors. In our analysis, we intentionally degrade the quality of the text by randomly splicing sentences in a random subset of speeches. Then, we run the same prediction exercise as the main text, using the now degraded 2021–2023 speeches to predict the label of all sentences in 2023 through the LOO process.

From our qualitative examination, these errors tend to be persistent within certain states (given that their journals have consistent formatting across a period of time) and they appear to affect a low-to-moderate number of sentences within speeches. To probe the sensitivity of our conclusions we take a grid-based approach in which we independently vary the proportion of states affected and the proportion of sentences affected within speeches (both, 0 to 0.3 by 0.05 intervals). Each time we iterate through a combination of values, we conduct an entire LOO prediction loop on all 2023 speeches. We then compute the predicted probabilities and assess the accuracy of predictions for all 49 possible combinations.

To determine the robustness of our original results, note that the true accuracy of 2023 speeches is 0.86. The interquartile range of our permutation tests is [0.88, 0.90]. As accuracy changes by just 0.02 across this range, we argue that these transcription errors have minimal impact on our results. The full range of accuracy values is [0.8, 0.96]. We

also note that the standard deviation of accuracies across these permutation tests is 0.03, whereas the standard deviation of accuracies across the original data (i.e., all speeches 1962–2023) is 0.09. To the extent that transcription errors degrade accuracy, they account for about one-third of that variation on average. And, interestingly, introducing errors can actually improve classification—likely due to introducing artifacts and regularizing text in a way that is informative. That leaves two-thirds of the variation to be explained by speech-specific factors. Together, this test provides evidence that our results are robust to these transcription errors.

B Additional Models

In the main text, we interact our key independent variables with a binary post-1995 indicator variable. However, the significant interaction terms could be due to a secular increase over time in how these coefficients relate to the outcome. Here, we re-specify the model, interacting an indicator for each decade with our key independent variables. If our binary indicator is picking up a secular increase, we should see that appear across the decade-vote share interactions. Instead, we see that 1970, 1980, and 1990, the effect of Republican vote share is no different than it was in 1960 (the baseline). However, in 2000, 2010, and 2020, the effect of Republican vote share becomes a positive predictor. This result is consistent with a structural change in how governors responded to their constituency in beginning in the 2000s. That 1990 is not significant is likely due to (1) the change not occurring until 1995 and (2) a lag in gubernatorial responsiveness to this new political environment.

Table B1: Interaction with Decade	
	(1)
Republican	0.403*** (0.098)
Two-Party Republican Vote Share (10s)	0.011 (0.053)
Republican x 1970	0.126 (0.152)
Republican x 1980	0.530+ (0.293)
Republican x 1990	0.317 (0.226)
Republican x 2000	0.166 (0.147)
Republican x 2010	0.331* (0.148)
Republican x 2020	0.930*** (0.266)
Rep. Vote x 1970	0.104 (0.083)
Rep. Vote x 1980	−0.209 (0.167)
Rep. Vote x 1990	0.194 (0.159)
Rep. Vote x 2000	0.155+ (0.077)
Rep. Vote x 2010	0.194* (0.081)
Rep. Vote x 2020	0.285* (0.112)
Election Year	0.078 (0.064)
Governor Running for Re-Election	−0.070 (0.075)
Unified State Government	−0.002 (0.067)
Fixed Effects	
State	✓
Year	✓
Num.Obs.	2427
R2 Adj.	0.366
R2 Within Adj.	0.205

+ p < 0.1, * p < 0.05, ** p < 0.01, *** p < 0.001

Note: Coefficients come from an ordinary least squares model where the dependent variable is the probability a speech is given by a Republican governor. Standard errors are clustered at the state level.

The models in Table B2 replicate the model in column 6 of Table 4 in the main text with results split by party. We can see that there is a correlation between what governors in both parties say and constituency partisanship in the post-1994 period.

Table B2: Interaction with Decade		
	Rep. Only	Dem. Only
	(1)	(2)
Two-Party Republican Vote Share (10s)	−0.078 (0.052)	−0.035 (0.099)
Two-Party Rep. Vote (10s) x Post-1994	0.248* (0.103)	0.334* (0.126)
Election Year	−0.024 (0.106)	0.144+ (0.085)
Governor Running for Re-Election	−0.105 (0.116)	−0.048 (0.093)
Unified State Government	0.040 (0.116)	0.108 (0.080)
Fixed Effects		
State	✓	✓
Year	✓	✓
Num.Obs.	1221	1206
R2 Adj.	0.309	0.294
R2 Within Adj.	0.011	0.040

Note: Coefficients come from an ordinary least squares model where the dependent variable is the probability a speech is given by a Republican governor. Standard errors are clustered at the state level.

References

Butler, Daniel M. and Joseph L. Sutherland. 2022. "Have State Policy Agendas Become More Nationalized?" *The Journal of Politics* p. 000–000.