Punishing Politicians on the basis of Covid-19

James Bisbee,[1‡] Haohan Chen,[1] Richard Bonneau,[1,3,4]
Joshua A. Tucker,[1,2] Jonathan Nagler[1,2]

[1]Center for Social Media and Politics, New York University

[2]Politics Department, New York University

[3]Computer Science Department, New York University

[4]Center for Data Science, New York University

[‡]To whom correspondence should be addressed: james.bisbee@nyu.edu

August 7, 2024

Abstract

Keywords:

1 Introduction

Covid-19 cleaves to political borders in American society across many dimensions. Areas that voted more Republican have lower vaccination rates and higher mortality [CITE]. Ideology is highly correlated with concern over the virus's health risks. And support for the policy solutions to the pandemic differs significantly by partisanship.

These empirical patterns have resulted in a growing consensus that the global health crisis is a politicized issue, a conclusion that is surprising given existing political science research on where and when we would expect to see political identities determine policy preferences. In theory, Covid-19 has many qualities that should make it a non-partisan issue. For example, Covid-19 is high salience, meaning that the public is less likely to simply adopt the positions of their co-partisan elites. In addition, Covid-19 carries significant individual costs, meaning that policy-preferences are less likely to be cheap talk. And it is a "valence" issue, meaning that all groups agree on a common definition of the public good.

Nevertheless, there is comprehensive evidence to suggest that Covid-19 is a politicized issue in the United States. The partisan gap is found in all areas of Covid-19, ranging from self-reported concern to observed behavior among the public; and from policy platforms to public pronouncements among elected officials [CITES]. This preponderance of evidence has strengthened one side of an ongoing debate in political science between what is known as the "Folk Theory" of democracy and partisan motivated reasoning (PMR). The Folk Theory rests on a spatial model of voting in which constituents have stable policy preferences and evaluate political candidates against these preferences, punishing those who deviate and rewarding those who are aligned. Conversely, PMR assumes that constituent policy preferences are endogenous to elite cues, meaning that the accountability mechanism assumed in the Folk Theory cannot operate.

It is natural to view the strength of the political polarization with respect to Covid-19 and conclude that this settles the debate in favor of PMR. However, simply because it is a politicized issue does not mean it is evidence of PMR. In this paper, we test whether constituents punish their elected officials for deviating from their views on Covid-19. We show that, while there is some evidence of constituents evaluating their governors' performance on Covid-19 more negatively when the governors deviated from their views of the seriousness of the pandemic in the summer of 2020, by January of 2021 we no longer observe these patterns. Similarly, we predict governor approval as a function of their constituents' views on specific policies and the extent to which those policies were implemented over the three months preceding the survey. Again, we find strong relationships between approval and the provision of constituents' preferred policies in the summer

of 2020, but no such evidence by January of 2021 with the exception of mask policies. We interpret this as dynamic evidence of the shift from initial behavior more consistent with the Folk Theory of democracy toward a new equilibrium better described by PMR.

2 Existing Research

Our contribution speaks to two related veins of research. The first is more general and focuses on the debate between the Folk Theory of democracy and partisan motivated reasoning (PMR). The Folk Theory views voters as rational actors who reward or punish incumbents based on a comparison between the voter's preferred policy and that adopted by the politician. According to this view, objective appraisal is crucial to the healthy functioning of democratic governance (**??**). In contrast, scholars of PMR have long argued that voters process new information in a biased fashion, using party affiliation as a "perceptual screen" through which they arrive at a particular conclusion (**?**). Citizens process political information with motivated reasoning, wherein they weigh information consistent with their partisanship more heavily than contradictory information (**?**). At the extreme, this perspective predicts that voters blindly adjust their beliefs and policy preferences to stay aligned with their party or partisan elites (**?**).

Despite a number of seminal contributions borne from this discussion, the debate continues. This ongoing debate is partly driven by measurement challenges. Most measures of learning rely on survey responses which may be confounded by motivations other than accuracy. These self-reported responses may conflate partisan motivated *responding* with partisan motivated *reasoning*, the latter of which captures the true updated conclusion drawn with partisan directional goals (**??**). Another explanation for the unsettled debate is that most studies examine a single explanatory factor at a time, precluding analysis of the circumstances under which both perspectives coexist (**?**).

The current consensus is that most people engage in partisan motivated reasoning most of the time. Individual-level characteristics, such as issue salience and the need for cognition, interact with system-level features, such as the clarity of party positions, to define the contexts in which we expect to find greater or weaker evidence for one side of the debate or the other (**?**). But in much of the applied research, scholars examine established political issues at a single moment in time, precluding their ability to capture variation along these theorized moderating dimensions. By focusing on Covid-19 over the course of 2020, a novel issue in January of 2020 whose salience and clarity of party positions evolved rapidly over the course of the year, we overcome these limitations.

The second vein of literature to which our research contributes is the substantive work on Covid-19. Recent contributions apply diverse methodological approaches to studying this phenomenon, using various outcome measures ranging from survey responses to GPS-tracked mobility as well as studying different policy settings ranging from the White House declaration of a state of emergency to state-specific stay-at-home orders. A powerful consensus from this work is that, at least in the context of the United States, one cannot understand Covid-19 without understanding U.S. politics. Specifically, studies find that the consumption of partisan news predict variation in social distancing behaviors (**??**). Some document a similar partisan gap in the response that is predicted by state-level policies and, importantly, the communication of these policies in the form of governor tweets (**??**). The primacy of politics in understanding the American public's response to Covid-19 is echoed in several other recent contributions (**???**), leading to a growing consensus that partisanship dominates all other factors.

However, just because politics is essential to understanding the American response to a global pandemic does not mean that it is evidence in favor of partisan motivated reasoning. As suggested by **?**, differences in the public's risk appetites that are correlated with ideology might produce strong correlations between Covid-19 beliefs and behaviors that are highly correlated with partisanship, without actually being causal. In this paper, we explicitly test the predictive power of the spatial model of voting in the context of Covid-19, investigating whether this core component of the Folk Theory of democracy describes the public's approval of their governors. We find that the spatial model's predictions are upheld in the summer of 2020, showing that constituents disapprove of governors who diverge from their views on the seriousness of the virus and the appropriateness of several different policy solutions to the pandemic. However, these patterns disappear by January of 2021, replaced by pure politics where constituents' approval is based solely on whether their governor adopts an ideologically congruent position.

Taken together, our results provide crucial texture to the growing consensus that partisan motivated reasoning dominates the Covid-19 crisis specifically, and contemporary American politics writ large. As suggested by **?**, our findings support the expectation that PMR should become more apparent as issues decline in salience and as party positions become more clearly delineated.

3 Empirical Setting: Covid-19 in the U.S.

We focus on analysis on Covid-19 in the United States, which we argue should be a uniquely hard test of partisan motivated reasoning (PMR). According to **?**, there are a series of scope conditions that limit

where and when we should expect to observe partisan motivated reasoning. Specifically, PMR should be less visible on valence issues where everyone agrees on the policy goal, less visible on high-salience issues that have direct consequences for the public welfare, and less visible where party positions are uncertain or not clearly differentiated. In the context of Covid-19, no one wants to be sick, a global pandemic is extremely high-salience where individual welfare is literally a matter of life and death, and the novelty of a once-in-a-century outbreak means that partisan stances were not already established, at least in the first half of 2020. However, by 2021, many of these dimensions had shifted. While the general agreement of not wanting to be sick persisted, the development of vaccines and treatments combined with the declining health risks of newer strains of the virus made the welfare consequences less existential, and Democrats and Republicans had grown quite clearly differentiated with respect to all aspects of the pandemic. Against this backdrop, we investigate whether the spatial model or partisanship better predicts public approval of governors in July of 2020 and January of 2021.

Dependent Variable

Our measure of governor evaluations are based on the Nationscape survey [CITE], which we aggregate to either the July 2020 waves or the January 2021 waves. We rely on a question that asks respondents to indicate their approval of their governor's handling of the Covid-19 pandemic, ranging from 1 indicating strong approval to 4 indicating strong disapproval. (We rescale this such that higher values indicate stronger approval.) This outcome is recorded in both waves of the Nationscape survey.

Independent Variables

Our primary predictors of interest are based on Twitter data, collected daily over the course of 2020 for governors and a random sample of the public which are geolocated to their state and divided into 6 ideological categories (Very Liberal, Liberal, Moderate, Conservative, Very Conservative, and not estimable or 'NA') using the method described in **?**. We manually labeled a set of 20,000 tweets that were about Covid-19 for whether the tweet expressed concern about the health risks of Covid-19 or whether it expressed skepticism and downplayed the severity of the pandemic. We then used these as training data for a transformer classifier which we used to generate predicted labels for each tweet. Our primary measure of interest is the difference between the predicted probability that the tweet takes the pandemic seriously minus the predicted probability that the tweet downplays the severity of the pandemic.

Figure 1 plots the average daily concern expressed in these tweets by group over the course of 2020. The patterns are broadly consistent with other findings on the influence of ideology on concern about Covid-19. Liberals take the pandemic more seriously overall, compared to conservatives. These political patterns are echoed in the descriptive plot of governor concern over the same period, with Democrats (blue) consistently expressing more concern online than Republicans (red). Nevertheless, there is powerful evidence of correlated changes in concern, as illustrated by all groups growing more or less concerned (roughly) in concert with each other. Furthermore, on average governors take the pandemic more seriously than their constituents, regardless of ideology.

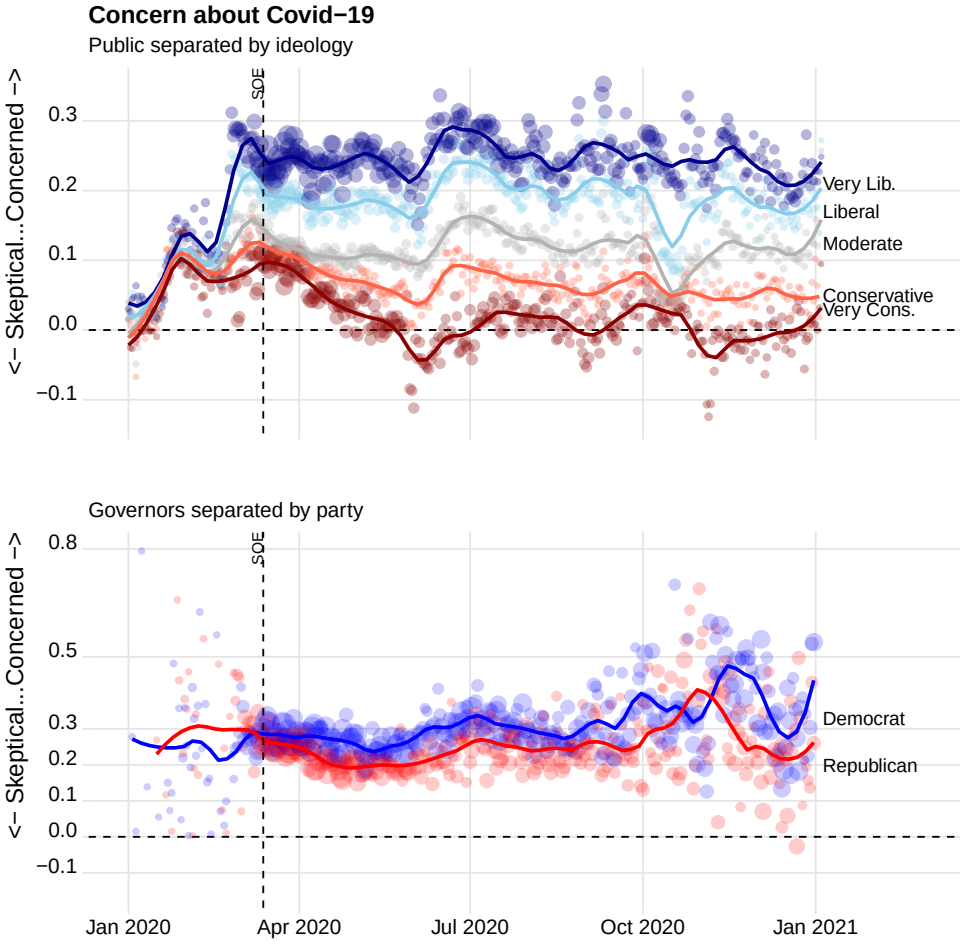

Figure 1: Overtime concern by the public (top) and governors (bottom), shaded by ideology (top) or partisanship (bottom). Points sized by the number of tweets written by each group about Covid-19 by day.

As these measures of concern about Covid-19 expressed by both the public and governors are generated by the same model, we can directly compare them to each other.[1] Specifically, we construct a "concern gap"

[1]There may be substantive reasons to be suspicious of a direct comparison between the public expressions of politicians on

measure that is the difference between the public's concern and that of their governor, such that positive values indicate the public takes the pandemic more seriously while negative values indicate that the governor takes the pandemic more seriously.

Descriptive Patterns

We start by looking at an overall scatterplot (Figure 2, left panel) in which we compare the relationship between the concern gap (x-axes) and evaluations of the governor (y-axes). Each point is an ideological subset of the constituents in a given state, surveyed in either the summer 2020 wave or the January 2021 wave, and sized by the total number of respondents available for this group in the Nationscape data.

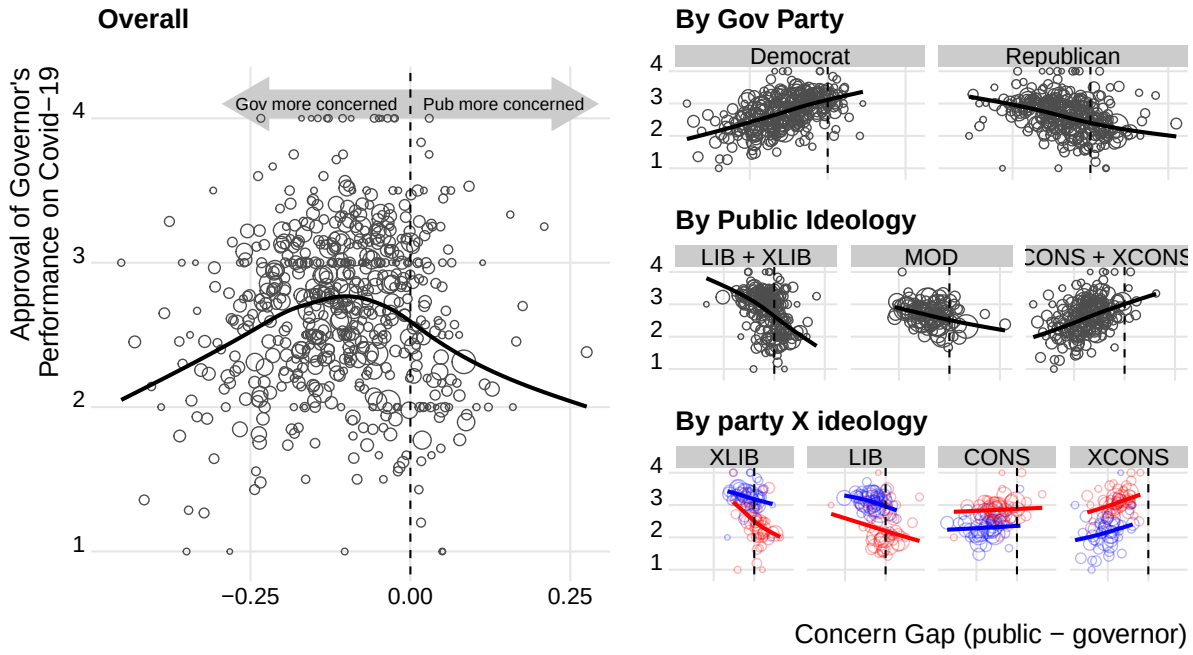

Figure 2: Scatter plot of relationship between concern gap (x-axes) and evaluations of governor's performance on Covid-19 (y-axes). Points are sized by the total number of respondents in the survey. Marginal histograms indicate distributions of approval and the concern gap.

The left panel of Figure 2 suggests that the spatial model is in fact a very useful tool for describing how constituents evaluated their governor on the basis of Covid-19. The local smoother line of best fit (in black) suggests that approval declines as the gap increases, just as the spatial model would predict. But in contrast with the spatial model, note that the maximum approval according to the black line is not at 0, but

Twitter and the expressions of the general public. In our Methods section below, we describe an interacted specification that relieves us this potentially heroic assumption.

rather at roughly -0.1, suggesting that the public has a preference for governors who take Covid-19 slightly more seriously than them.

However, by aggregating over all states and ideologies, this simplified plot ignores the role played by politics entirely. The subsequent panels on the right in Figure 2 demonstrate the importance of incorporating ideology and partisanship into our empirical analyses, by plotting lines of best fit separately by governor party (top right panel), ideological group (middle right panel), and their interaction (bottom right panel). In isolation, one might conclude that Democrat governors benefit from taking the pandemic less seriously relative to the public while Republican governors should take it more seriously, based on the top right panel. But these aggregate patterns obscure systematic differences among ideological groups. The positive slopes for conservatives in the middle right panel indicate that these ideological subsets always reward governors who take the pandemic less seriously, regardless of partisanship. Conversely, the negative slopes for liberals indicate the opposite – regardless of how much more concerned are governors relative to even their most liberal constituents, they are consistently rewarded by this group with higher approval ratings. Finally, the bottom right panel indicates that the responsiveness of liberals to the concern gap is attenuated among those living in states with a Democratic governor (blue lines in the bottom right panel of Figure 2). The same patterns are found among conservatives reacting less strongly to Republican governors (red lines in the bottom right panel). In sum, while the aggregate data is consistent with the spatial model, disaggregating by politics in fact suggests the opposite – ideological groups appear to condition approval not on the difference between their concern and that of their governor, but rather on a simple calculation of whether the governor is promoting the party line (i.e., Democrats take the pandemic seriously while Republicans don't).

Despite these patterns, one might be concerned that, by focusing not on a specific policy preference but rather on a more subjective sense of concern about Covid-19's health risks, we are submitting the spatial model to an unfair test. Therefore, we also look at specific policies related to the pandemic which were asked in the Nationscape waves. We link these self-reported measures of support for different policies (i.e., mask mandates, school closures, etc.) with a measure of whether the state had such a policy implemented. To merge these data with our July 2020 and January 2021 waves of the Nationscape survey, we calculated the average concern for each group over the three months preceding the survey date (i.e., April 1, 2020 through July 1, 2020 for the July 2020 wave). Similarly, for the specific policy dimensions, we tally the total number of days over the preceding three months that a state had a given policy in place.

We again start with a purely descriptive look at the data, dichotomizing both public approval and policy implementation to be 1 if the public either somewhat or strongly approves of the policy, and to be

1 if the public's state had one or more days of the policy implemented over the preceding 3 month period. Figure 3 focuses on four policies which we have survey-based measures of support for, and plots approval of the governor's performance on Covid-19 as a function of whether the public supports the policy (x-axes) and whether the policy has been implemented over the preceding three months (y-axes). The tiles are sized based on governor performance approval and shaded relative to the overall mean of 2.6 on a scale ranging from 1 to 4. The diagonals capture congruence between public policy preferences and policy implementation, while the off-diagonals capture incongruence. [NEED TO UPDATE]

In general, there is descriptive support for the spatial model, with pro-policy constituents living in states where the policy is implemented expressing higher levels of approval for their governor's performance on Covid-19. Interestingly, the lowest levels of approval are consistently found among constituents who oppose a certain policy but live in states where the policy is implemented (top-right quadrants of each panel). While suggestive, these patterns are consistent with a story in which the presents of an unwanted policy is more activating for constituent disapproval than the absence of a desired policy.

4 Methods:

The preceding descriptive results suggest general support for the spatial model, albeit a pattern that obtains when the public is aggregated. When we disaggregate to accommodate partisanship, we find less support for a spatial model within the behavior of subsets of the public. To formally evaluate the predictive power of the spatial model, we turn to regression analysis.

Specifically, we predict constituent evaluations of their incumbent governor as a function of incumbent politicians' stance on Covid-19, and the constituents' concern about the pandemic. Our July 2020 and January 2021 data are organized by state (s) and ideological subset of the public (i), with columns indicating the public's evaluations of their politician (either $\text{approve}_{s,i}$ or $\text{vote}_{s,i}$); the governor's concern about Covid-19's health risks (gov_concern_s) and the public's concern about Covid-19's health risks ($\text{pub\_concern}_{s,i}$), both averaged over the preceding three months; the public's self-reported preferences for a battery of Covid-related policies ($\text{pub\_policy}_{s,i}$); the total days in which at least one of the associated policies was enforced in their state (policy_s); and a set of controls measured at the state level (average weekly change in Covid-19 deaths over the preceding three months deaths_s, total tweets written by all groups $\text{allTweets}_{s,i}$, and total Covid-19 tweets written by all groups $\text{covidTweets}_{s,i}$). Our main predictor of interest is the "concern gap" which, as described above, measures the difference in the degree to which governors and their constituents took

Figure 3: Approval of governor performance on Covid-19 indicated with the size of the tiles and the color, relative to an overall average of 2.6 on a scale from 1 (strongly disapprove) to 4 (strongly approve). Each tile represents support among constituents who either support or oppose a policy (x-axes) living in a state with 0 days with the policy implemented versus those living in a state with 1 or more days with the policy implemented over the preceding three months. Diagonals capture congruence between public preferences and policy implementation (i.e., either the public supports a policy and lives in a state where it is implemented, or the public opposes a policy and lives in a state where it is not implemented). Off-diagonals capture incongruence.

Covid-19 seriously in the three months prior to the survey. We construct $\text{gap}_{s,i}$ as the absolute difference between the public's concern and their governor's concern, and examine the sensitivity of our findings to alternative measures, including the raw difference, the squared difference, and two measures that truncate the gap measure to capture only variation in which the public takes the pandemic more seriously than their governor, and vice versa.

To evaluate the extent to which constituent evaluations reflect alignment between their views on the pandemic and the views of their incumbent governor, we run an OLS regression model of the following specification:

$$\text{approve}_{s,i,t} = \delta_t + \beta_1 \text{gap}_{s,i,t} + \boldsymbol{\gamma_1} \mathbb{I}(\text{ideo}_{s,i,t}) + \boldsymbol{\gamma_2} \mathbb{I}(\text{party}_{s,t}) + \boldsymbol{\gamma_3} \mathbb{I}(\text{ideo}_{s,i,t}) * \mathbb{I}(\text{party}_{s,t}) + \boldsymbol{\lambda} \mathbf{X}_{s,i,t} + \varepsilon_{s,i,t} \tag{1}$$

where $\mathbb{I}(\text{ideo}_{s,i})$ is an indicator for the ideological subset of the public i in state s (ranging from very liberal to very conservative), $\mathbb{I}(\text{party}_s)$ is an indicator for the party affiliation of the politician in state s (either Democrat, Independent/Third Party or Republican), and $\mathbf{X}_{s,i}$ is the battery of controls described above. Our main results include the period fixed effect for the stacked data δ_t, although we also run these regressions on the July 2020 and January 2021 data separately.

While our empirical setting precludes causal identification, it is nevertheless instructive to think through the causal implications of the debate between the Folk Theory of democracy (embodied in the spatial model of voting) and partisan motivated reasoning (PMR). Fundamentally, we are interested in the degree to which the gap between constituent and politician policy preferences influence the constituent's approval of the politician, indicated with pathway (a) in Figure 4. According to the spatial model, as the gap widens, voter utility declines according to $-||\theta_i - x_j||^\alpha$, which we assume translates into a lower evaluation of the politician.

Conversely, PMR posits that both constituent preferences and politician evaluations are informed by co-partisan elite cues, which we indicate with pathway (b) in Figure 4. Substantively, this means that both a voter's beliefs and her approval of a politician are determined by her ideology / party affiliation, meaning that a Democrat voter will approve of a Democrat governor and adopt similar policy preferences, meaning that the gap will be small and the evaluations will be positive. Statistically, PMR acts as an omitted variable in this setting, rendering both politician and constituent preferences endogenous to partisanship / ideology. As such, the first-order test to adjudicate between the spatial model and PMR is to compare our model's estimates with and without this omitted variable. If its inclusion removes all predictive power of the preference gap measure, we will conclude that PMR wholly explains constituent behavior in the context of governor evaluations based on Covid-19.

To evaluate whether the spatial model or PMR better explains variation in governor approval, we compare the β_1 coefficient estimated using the full specification described above in Equation 1 (denoted β_1^{full}), to the β_1 coefficient estimated without the inclusion of the γ_3 interaction term ($\beta_1^{restricted}$). This

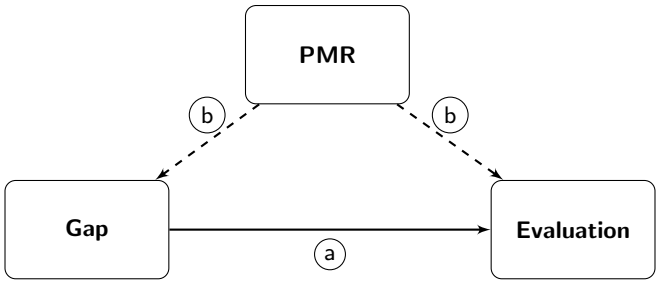

Figure 4: The Folk Theory (a) versus PMR (b) as a causal directed acyclic graph. According to the Folk Theory, the gap in policy preferences between constituents and governors influences constituent evaluations of the politician, as indicated by pathway (a). Conversely, according to PMR, both policy preferences and evaluations are endogenous to partisanship, as indicated by pathway (b).

comparison illuminates the degree to which PMR operates as an omitted variable, predicting both the public's evaluation of their governor and the degree to which they are aligned in their views on the pandemic. If $\beta_1^{restricted} << \beta_1^{full}$, we will conclude in favor of PMR. We also make this comparison on the July 2020 and January 2021 data subsets separately to test whether the relative novelty of the virus in the summer of 2020 reduced the influence of PMR (**?**).

Table 1: Summary Statistics

Variable	N	Mean	Std. Dev.	Min	Pctl. 25	Pctl. 75	Max
Wave	594						
... jan_2021_fall	297	50%					
... summer_2020	297	50%					
Ideology	588						
... MOD	98	16.7%					
... XLIB	98	16.7%					
... LIB	98	16.7%					
... CONS	98	16.7%					
... XCONS	98	16.7%					
... NA	98	16.7%					
Gov. Party	594						
... Democrat	292	49.2%					
... Republican	302	50.8%					
Approve Gov. Perf.	575	2.683	0.565	1	2.31	3.072	4
Concerned about Covid-19	578	3.214	0.459	1	3	3.544	4
Support Closing Business	577	3.015	0.544	1	2.667	3.4	4
Support Social Distancing	578	3.052	0.551	1	2.722	3.435	4
Support Closing Schools	578	2.844	0.584	1	2.439	3.25	4
Support Mask Requirements	576	3.324	0.485	1	3.052	3.657	4
Concern Gap (public - Gov.)	585	-0.106	0.109	-0.452	-0.175	-0.031	0.276

5 Results:

5.1 Approval of Governor's Performance on Covid-19

We start with the specification described in Equation 1, and examine how the public's approval of their governor's handling of the pandemic is correlated with the gap between the degree to which the public and their governor took the pandemic seriously in the third quarter of 2020. Column 1 drops all controls and calculates just the bivariate relationship between approval of governor's performance on Covid-19 and the absolute gap in concern by the governor and the public. Column 2 adds the control for average weekly deaths over the preceding 3 months. Column 3 includes dummies for public ideology and governor party, while column 4 interacts these. Columns 5 and 6 re-run column 4's specification on the July 2020 and January 2021 data individually. Standard errors are clustered at the state level, and all predictors are either dummies (ideology and partisanship) or scaled, allowing for direct comparison of magnitudes of coefficients, which capture the change in a standard deviation increase in the predictor on the raw change in governor approval on a scale ranging from 1 to 4 (mean = 2.32, sd = 0.53).

As illustrated the β_1 coefficient (row 1) is negative and statistically significant across the first 3 columns, indicating that the larger the difference in the level of concern of governors and their constituents, the lower the public's approval of their governor. However, when we properly control for partisan motivated reasoning via the interaction between public ideology and governor partisanship in column 4, the negative association is reduced by two-thirds and is no longer statistically significant. As per the discussion in Section **??** above, this comparison would suggest that the spatial model is ill-equipped to explain variation in approval of governor performance on Covid-19 after controlling for politics. However, when we subset the data to examine July 2020 and January 2021 separately, we highlight that the negative coefficient on the absolute gap reappears in the earlier period, while it is an even smaller null in the later period (and is even of the opposite sign). These patterns are consistent with the scope conditions described by **?** and indicate that the debate between the Folk Theory and PMR should be recast as about where and when we might expect one framework to dominate.

Model:	Vanilla (1)	+PID (2)	+Ideo (3)	PID Inter (4)	Ideo Inter (5)
Variables					
Sq Gap	-0.026**	-0.028**	-0.050***	-0.021*	-0.020**
	(0.011)	(0.011)	(0.012)	(0.011)	(0.010)
Public: Indep		-0.252***		-0.486***	
		(0.045)		(0.028)	
Public: Rep		0.046		-0.494***	
		(0.095)		(0.048)	
Public: Lib			0.169**		0.581***
			(0.072)		(0.038)
Public: Mod			0.063		0.316***
			(0.046)		(0.031)
Public: X-Cons			0.078*		-0.070*
			(0.040)		(0.041)
Public: X-Lib			0.308***		0.719***
			(0.073)		(0.044)
Public: Indep × Gov: Rep				0.481***	
				(0.049)	
Public: Rep × Gov: Rep				1.09***	
				(0.086)	
Public: Lib × Gov: Rep					-0.848***
					(0.069)
Public: Mod × Gov: Rep					-0.504***
					(0.049)
Public: X-Cons × Gov: Rep					0.292***
					(0.063)
Public: X-Lib × Gov: Rep					-0.855***
					(0.069)
Fixed-effects					
stab	Yes	Yes	Yes	Yes	Yes
month	Yes	Yes	Yes	Yes	Yes
Fit statistics					
Observations	16,927	16,778	14,494	16,778	14,494
R^2	0.055	0.079	0.072	0.141	0.132
Within R^2	0.0009	0.025	0.016	0.091	0.081

Clustered (stab) standard-errors in parentheses
*Signif. Codes: ***: 0.01, **: 0.05, *: 0.1*

Table 2: Regression results predicting approval of Governor performance on Covid-19 as a function of the absolute difference in concern expressed by the public and their governor.

6 Policy Representation

The preceding results define Democratic accountability in terms of the gap between the degree to which the ideological public took Covid-19 seriously relative to their governor. However, we also have data on the public's support for specific policies related to the pandemic, including social distancing, mask mandates, business closures, etc. As illustrated in Figure 5, the policy response to Covid-19 varied dramatically by state, with DC recording an average of 10 Covid-related policies per day and Nebraska recording an average of 1.3 over the same period. Across types of policies, the most prevalent are social distancing policies, followed by policies pertaining to masks (right panel of Figure 5).

We adopt a similar regression specification to those described above in Equation 1, except here we are predicting approval of the Governor's performance on Covid-19 as a function of the interaction between the public's support for a given policy and the total number of days their state had one or more instances of the associated policy enforced. Formally,

$$\begin{aligned}\text{approve}_{s,i,t} = \delta_t + \beta_1 \text{pol\_support}_{s,i,t} + \beta_2 \text{pol\_enact}_{s,t} + \beta_3 \text{pol\_support} * \text{pol\_enact} \\ + \gamma_1 I(\text{ideo}_{s,i,t}) + \gamma_2 I(\text{party}_{s,t}) + \gamma_3 I(\text{ideo}_{s,i,t}) * I(\text{party}_{s,t}) + \lambda \mathbf{X}_{s,i,t} + \varepsilon_{s,i,t}\end{aligned}$$

Positive β_3 coefficients indicate that the marginal effect of public support for a given policy on governor approval is higher in states with more vigorously implemented policies supported by the public.

We plot the marginal effects of this analysis in Figure 6, revealing systematic evidence of a positive association between support for a policy and governor approval among respondents living in states where the policy was more robustly enacted. To give a substantive example using the face mask policies in January of 2021 (top right panel), constituents who support mask mandates but who live in states with no mask policies over the preceding three months are significantly less supportive of their governor. Conversely, pro-mask constituents living in states with mask policies enacted continuously over the preceding three months prior to the survey are significantly more supportive of their governor. We interpret these positive interaction terms as evidence consistent with the spatial model of voting. Importantly, these positive interaction effects are strongest in the July 2020 sample and – with the exception of mask mandates – all but disappear by January of 2021, providing further support for our claim that the spatial model was supported earlier in the pandemic but later.

7 Discussion

In this paper, we measure the degree to which voters evaluate their governors on the basis of Covid-19. We broadly investigate two competing frameworks of voter behavior: the first based on a spatial model of voting in which voters punish politicians who deviate from their position, and the second based on a partisan motivated reasoning (PMR) model in which voters update their positions based on co-partisan elite cues. Our results support the scope conditions described by **?**, who emphasizes that we should expect to see more evidence of PMR where the issue is lower salience and where party positions are more clearly defined. In line with this assumption, we find empirical patterns consistent with the spatial model in the summer of 2020 when the issue of Covid-19 was higher salience and when party positions were less clearly defined. Conversely, by January of 2021, empirical support for the spatial model all but disappeared, and approval of the governor's performance on the pandemic was wholly predicted by the ideology of the voters and the partisanship of their governor.

Substantively, our findings speak to a growing body of research that examines Covid-19, a high salience, novel issue that nevertheless has grown increasingly politically polarized in the United States. Our conclusion that both the spatial model and PMR operate in the context of the pandemic offers an important caveat to a flurry of findings emphasizing the pivot role of politics in defining the public's response to a once-in-a-century public health crisis. While we do not discount the importance of politics, we push back on the emerging consensus that PMR completely dominates public discourse. In the earlier part of the pandemic, constituents behaved more like rational actors in holding governors accountable.

The implications for our work for the normative state of American politics cuts both ways. On the one hand, the evidence we document of the spatial model in the summer of 2020 comports with precisely when we should expect to see the influence of PMR at its weakest. On the other hand, politics was already highly influential even in this period, and evolved to completely determine voters' evaluations by January of 2021. That even an issue as important to individual welfare still grew to be best explained with partisan motivated reasoning is perhaps a cause for concern, and further evidence of the increasing political polarization at all levels of American politics.

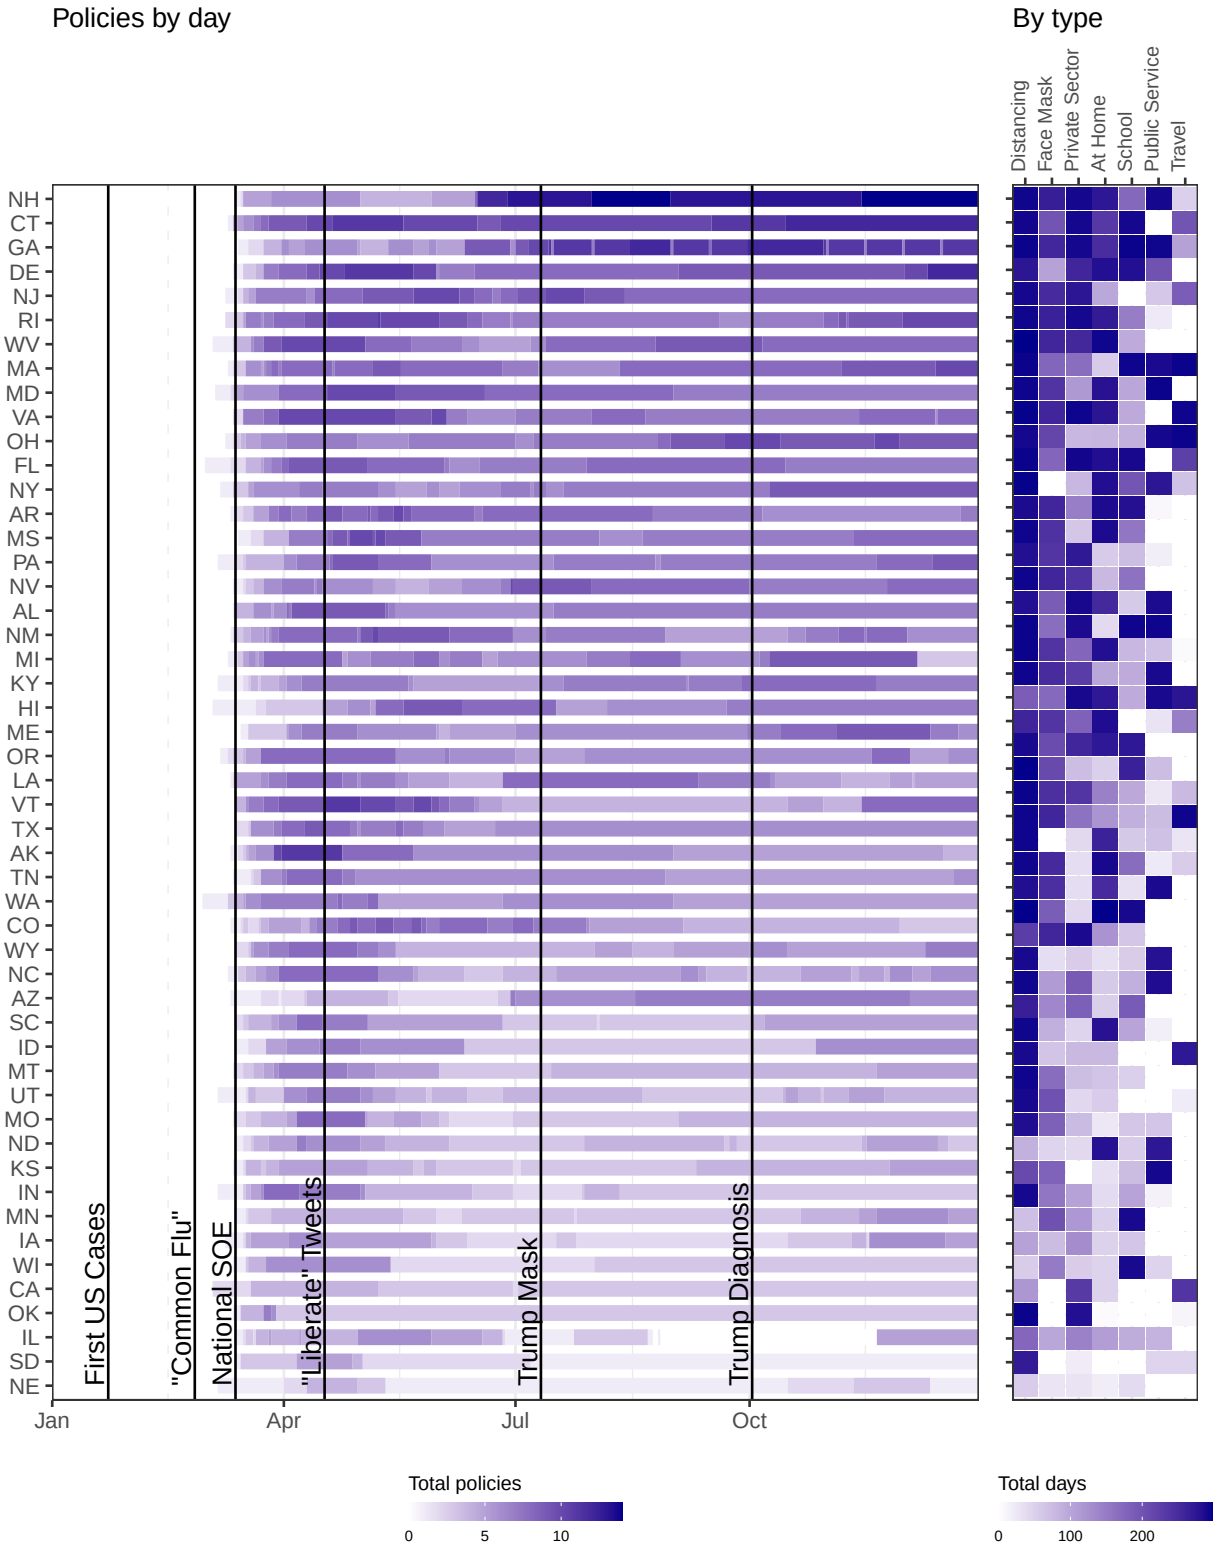

Figure 5: Total daily policies related to Covid-19 by state (left panel) and total days with one or more policies by type (right panel).

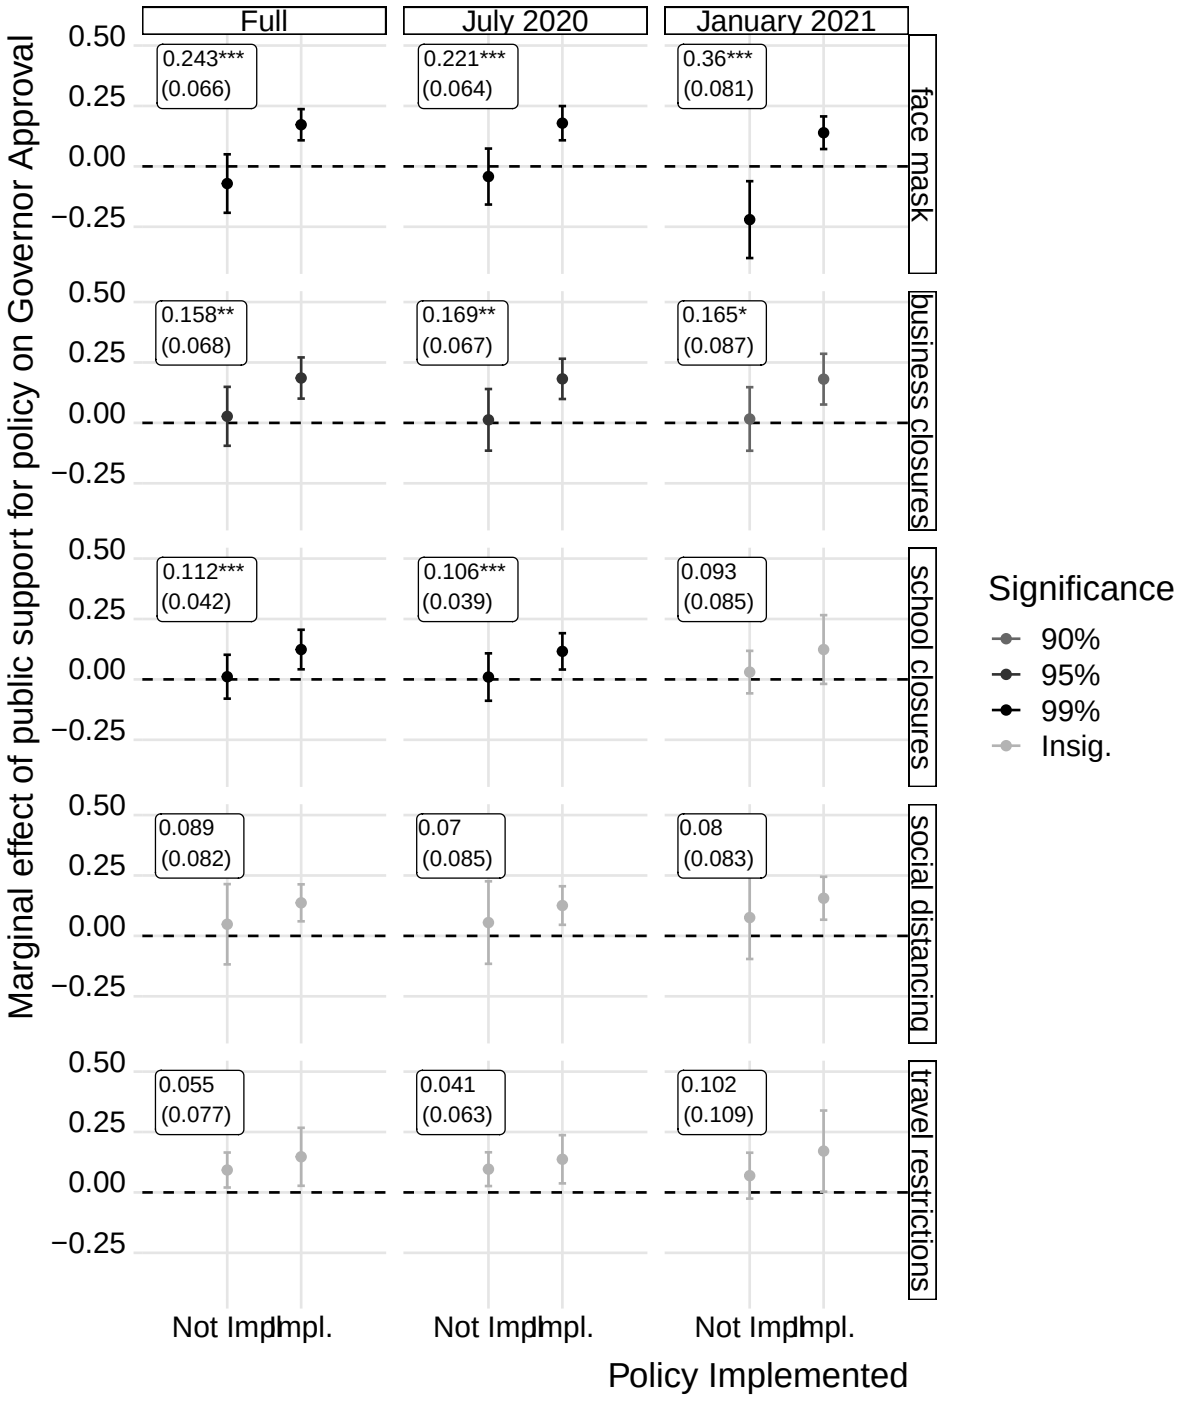

Figure 6: Marginal effects (y-axes) of support for a particular policy (rows) on approval of governor's performance on Covid-19 across different levels of policy prevalence in the respondents' state, measured as the total days in which one or more associated policies was enacted over the preceding three months. Columns summarize results for the full survey (left), the July 2020 sample (center), or the January 2021 sample (right).