Polarization and Public Health: Partisan Differences in Social Distancing during COVID-19

Hunt Allcott, *New York University, Microsoft Research, and NBER*[*]

Levi Boxell, *Stanford University*

Jacob Conway, *Stanford University*

Matthew Gentzkow, *Stanford University and NBER*

Michael Thaler, *Harvard University*

David Yang, *Harvard University*

April 2020

Abstract

We study partisan differences in Americans' response to the COVID-19 pandemic. Political leaders and media outlets on the right and left have sent divergent messages about the severity of the crisis, and this could potentially undermine public compliance with key measures such as social distancing. We develop a simple model of a pandemic response with heterogeneous agents that clarifies the causes and consequences of divergent responses. We use location data from a large sample of smartphones to show that areas with more Republicans engage in less social distancing, controlling for other factors including state policies, population density, and local COVID cases and deaths. We then present new survey evidence of significant gaps between Republicans and Democrats in beliefs about personal risk and the future path of the pandemic.

[*]E-mail: hunt.allcott@nyu.edu, lboxell@stanford.edu, jcconway@stanford.edu, gentzkow@stanford.edu, michaelthaler@g.harvard.edu, davidyang@fas.harvard.edu.

We thank SafeGraph for providing access to the data and the SafeGraph COVID-19 response community for helpful input. We acknowledge funding from the Stanford Institute for Economic Policy Research (SIEPR), the John S. and James L. Knight Foundation, the Sloan Foundation, the Institute for Humane Studies, and the National Science Foundation (grant number: DGE-1656518). For our survey, we registered a pre-analysis plan on the AEA Registry, with ID AEARCTR-0005632. This study was approved by IRBs at NYU (IRB-FY2017-938) and Harvard (IRB17-1725).

1 Introduction

Mobilizing an effective public response to an emerging pandemic requires clear communication and trust (Holmes 2008; Taylor et al. 2009; van der Weerd et al. 2011; Vaughn and Tinker 2011). Key measures such as social distancing and self-quarantine can rarely be enforced entirely by coercion, particularly in democratic societies. The public must understand what is required of them and be persuaded of the importance of complying.

In this sense, partisan differences could be a key risk factor in the US response to the COVID-19 pandemic. Prominent officials have sent conflicting messages about the crisis, with President Trump and other Republican officials sometimes saying it was less severe, and Democrats giving more emphasis to its dangers (Beauchamp 2020; Stanley-Becker and Janes 2020; Coppins 2020; McCarthy 2020). Partisan media have tended to echo this division (Aleem 2020; Kantrowitz 2020). The result could be gaps between citizens on the right and left in their compliance with public health measures such as social distancing, possibly leading to large human and economic costs.

In this paper, we combine GPS location data from a large sample of smartphones with a new survey to study partisan differences in the early response to COVID-19. The GPS data are collected by the company SafeGraph, and record daily and weekly visits to points of interest (POIs), including restaurants, hotels, hospitals, and many other public and private businesses. Our primary analysis focuses on the period from January 26, 2020 to March 28, 2020.

We begin with two motivating facts. First, recent nationwide surveys have shown that Democrats are more concerned about the spread of COVID-19 and report taking more steps to avoid infection than Republicans. Second, counties with higher Democratic vote shares in the 2016 election saw more social distancing, as measured by location data. However, these counties also had more coronavirus cases and were in states that initiated stay-at-home policies earlier. Thus, the partisan differences in social distancing could simply be the expected result of local differences in infection risk or regulation.

We present a simple model that clarifies the potential causes and consequences of divergent social-distancing behavior. It combines a standard epidemiological model of a pandemic with an economic model of optimizing behavior by heterogeneous agents. The model clarifies that divergent responses between groups need not be inefficient. One group might engage in less social distancing because their costs of distancing are greater (e.g., they would lose more income as a result) or because their benefits of distancing are smaller (e.g., they are at lower risk of infection).

However, differences in behavior resulting from divergent *beliefs* of otherwise similar agents do suggest systematic inefficiency, as optimizing based on different beliefs means that the marginal costs of social distancing are not equated across people. In that case, society gets less social distancing at higher cost than if agents had the same beliefs.

Our main GPS results show that the strong partisan differences in social distancing behavior that emerged with the rise of COVID-19 are not merely an artifact of differences in state policies or observed risks. Controlling for state-time fixed effects to account for heterogenous policy responses by state governments only attenuates the partisan gap slightly. Including controls to proxy for health and economic variables interacted flexibly with time attenuates the gap more substantially, but it remains statistically and economically significant. After including our full set of controls, we estimate that moving from the 10th to the 90th percentile of Republican county vote share is associated with a 16.7 percent increase in the number of POI visits during the week of March 22.

Our findings are robust to the inclusion or exclusion of control variables, excluding states with early COVID-19 outbreaks, or dropping highly populated counties. Replacing the continuous measure of partisanship with discrete indicators for portions of the Republican vote share distribution or restricting the sample to counties from certain portions of the distribution does not change our qualitative conclusions. Furthermore, there is no evidence of a similar partisan gap during the same period in 2019.

Finally, we use new survey data to provide additional evidence on the differences in beliefs that may underlie the partisan gap in behavior. This survey was conducted from April 4-6, 2020. We collect participants' demographics (including party affiliation), beliefs regarding the efficacy of social distancing, self-reported distancing due to COVID-19, and predictions about future COVID-19 cases. We confirm that Democrats believe the pandemic to be more severe and report adjusting their behavior more than Republicans do on average. In our survey, we also randomly vary whether predictions are incentivized, and do not find evidence that incentives reduce the partisan gap, suggesting that these predictions are less likely to be due to partisan cheerleading (as in Bullock et al. 2015 and Prior et al. 2015), and more likely to reflect true differences in beliefs.

Several contemporaneous studies also measure partisan differences in responses to COVID-19.[1] Gadarian et al. (2020) present survey evidence showing partisan gaps in self-reported re-

[1]Coverage in the media and some studies examine partisan heterogeneity in response to COVID-19 with no or few controls for differential risk exposure or costs of social distancing (e.g., Economist 2020; Andersen 2020). Baker et al. (2020) use transaction-level data and examine heterogeneity in consumption responses to COVID-19.

sponses to the pandemic. Barrios and Hochberg (2020) show differences between Republican and Democratic areas in the frequency of COVID-related queries on Google and in movement patterns as measured in GPS data from a different source than the one we use here. Painter and Qiu (2020) examine partisan heterogeneity in response to state-level, stay-at-home orders.

Our work contributes to a broader literature on what drives responses to pandemics (e.g., Blendon et al. 2008; Vaughan and Tinker 2009; Fineberg 2014). Risk perception, behavior changes, and trust in government information sources change as pandemics progress (Ibuka et al. 2010; Bults et al. 2011). Demographic characteristics, such as gender, income, geography, or social interactions, are important determinants of the adoption of recommended public health behaviors (Bish and Michie 2010; Ibuka et al. 2010; Bults et al. 2011; Chuang et al. 2015; Shultz et al. 2016; Gamma et al. 2017).

A related literature focuses on the consequences of political polarization for health behaviors (e.g., Iyengar et al. 2019 and Montoya-Williams and Fuentes-Afflick 2019). Party affiliation is correlated with physician recommendations on politicized health procedures, enrollment in government exchanges created under the Affordable Care Act, and beliefs in the safety of vaccines (Hersh and Goldenberg 2016; Lerman et al. 2017; Sances and Clinton 2019; Trachtman 2019; Krupenkin 2018; Suryadevara et al. 2019). We show how partisan differences can lead to the inefficient allocation of public health goods, such as social distancing, during pandemics.

Our work also relates to a broader literature on partisan differences in trust and beliefs. For instance, a large body of empirical literature documents partisan differences in beliefs about factual events such as unemployment (Bartels 2002; Gaines et al. 2007; Bullock et al. 2015). There exists a growing literature on building theoretical models of opinion polarization to explain observed partisanship (Dixit and Weibull 2007; Benoit and Dubra 2015; Ortoleva and Snowber 2015; Fryer et al. 2019). Furthermore, a substantial empirical literature studies the link between media markets and political polarization (Glaeser and Ward 2006; McCarty et al. 2006; Campante and Hojman 2013; Prior 2013).

Finally, our work adds to the increasing number of papers using GPS or related data to study social interactions. For example, Dubé et al. (2017) test the effectiveness of mobile targeting with coupons to competing movie theaters based on consumers' real-time location. Hanna et al. (2017) use data from Google Maps to estimate the effects of lifting high-occupancy vehicle restrictions in Jakarta, Indonesia.[2] Chen and Rohla (2018) and Athey et al. (2019) use SafeGraph data to

[2] See also Blattman et al. (2018) and Davis et al. (2019).

measure the effects of political polarization on the length of Thanksgiving dinners and to estimate a novel measure of racial segregation, respectively.

Sections 2, 3, 4, 5, and 6, respectively, present our motivating facts, theoretical framework, data, GPS analysis, and survey results.

2 Motivating Facts

In this section, we present basic facts on differences in beliefs and social distancing by political party. Figure 1 presents results from previous national polls. Panel A shows that Democrats were consistently more concerned than Republicans about the spread of coronavirus in the United States from January 26 through the most recent polls in early April. The next three panels show differences in self-reported social distancing behaviors. Panel B presents results from a March 13th poll, showing that Democrats were more likely to say they were eating at home more often, had stocked up on food and supplies, changed travel plans, and cancelled plans to avoid crowds. Panels C and D show that throughout the month of March, Democrats were more likely than Republicans to say that they were avoiding public places and small gatherings.

Of course, one would like to know whether these self-reported behavior changes can be corroborated in observational data. Figure 2 visualizes geographic variation in this social distancing response, and compares observed variation to analogous distributions of partisanship, COVID-19 confirmed cases, and public policy responses. Panel A maps the social distancing response observed in each county, as measured by the percent decrease in SafeGraph visits between the week beginning January 26th and the week beginning March 22nd, using data described below. Panel B shades counties by their party affiliation, captured by the Republican vote share in the 2016 presidential election. Panel C maps the number of COVID-19 cases confirmed in a given county by March 28th. Panel D shades states by the effective start date for the earliest statewide "stay-at-home" order issued. Panels A and B exhibit a strong geographic correlation between the counties with weaker social distancing responses and those with higher Republican vote shares. Panel D shows that areas with stronger distancing responses also generally instituted earlier statewide, stay-at-home orders. We also observe stronger social distancing responses in counties with more COVID-19 confirmed cases (Panel C).

These results highlight the value of a framework to understand why people from the two political parties behave differently, and why that might matter.

3 Stylized Model

In this section, we present a stylized model to clarify why it might matter if different types of people choose different amounts of social distancing. We embed an epidemiological model of disease transmission into an economic model with agents who maximize utility considering the expected private cost of disease.

3.1 Epidemiological Model

We use a discrete time version of the standard SIR epidemiological model (Kermack and McKendrick 1927). In each period t, each person is in one of four states $\sigma \in \{S, I, R, D\}$, representing Susceptible, Infected, Recovered, and Deceased. The share of the population in each state at time t is s_t, i_t, r_t, and d_t. Let β represent disease infectiousness, and let c_t denote an individual's amount of risky behavior at time t—for example, the amount of travel, dining out, failing to wash hands, and other activities that increase risk of becoming Infected.

All people begin in the Susceptible state. A Susceptible person becomes Infected at time $t+1$ with probability $c_t \beta i_t$ and stays Susceptible with probability $(1 - c_t \beta i_t)$. Infected people stay Infected for one period, after which they become Deceased with probability ψ or Recovered with probability $(1 - \psi)$. Both D and R are absorbing states.

Let θ index different types of people—for example, liberals and conservatives. Let $\omega_{\theta\sigma t}$ be a state variable representing the share of type θ that is in state σ at time t. The population is of measure 1, so $\sum_\theta \sum_\sigma \omega_{\theta\sigma t} = 1$.

3.2 Individual Decisions

People of type θ earn flow utility $u_\theta(c_t; \sigma_t)$, which depends on their risky behavior c_t and their state σ_t. People discount the future at rate δ and maximize expected lifetime utility $\sum_{\tau=t}^{\infty} \delta^\tau u_\theta(c_\tau; \sigma_\tau)$. Define $V_\theta(\sigma)$ as the expected lifetime utility of a person currently in state σ; note that this also implicitly depends on current and future population states $\omega_{\theta\sigma t}$. Being infected reduces utility, so we assume $V_\theta(S) > V_\theta(I)$.

We focus on Susceptible people, as they comprise most of the population during the period we study. We can write their maximization problem as a Bellman equation, in which people maximize the sum of utility from risky behavior today and expected future utility:

$$V_\theta(S_t) = \max_{c_t} \left\{ \underbrace{u_\theta(c_t; S_t)}_{\text{current utility from risky behavior}} + \underbrace{\delta \left[c_t \beta i_t V_\theta(I) + (1 - c_t \beta i_t) V_\theta(S) \right]}_{\text{expected future utility}} \right\}. \tag{1}$$

The first-order condition for privately optimal risky behavior is

$$\underbrace{u'_\theta}_{\text{marginal utility of risk}} = \underbrace{\beta i_t}_{\text{marginal infection probability}} \underbrace{\delta \left(V_\theta(S) - V_\theta(I) \right)}_{\text{private cost of infection}}. \tag{2}$$

The first-order condition shows that people choose their risky behavior to equate marginal benefit (more utility today) with private marginal cost (higher risk of infection, which reduces future utility). The equation illustrates that there are three reasons why risky behavior might vary across types. First is the marginal utility of risk (or equivalently, the marginal cost of social distancing): for example, people vary in how much they like travel and dining out, as well as in how easy it is to work from home. Second is the marginal infection probability: for example, local infection rate i_t differs across geographic areas. Third is the private cost of infection: for example, infection is more harmful for people who are older or have underlying health conditions.

3.3 Social Optimum

It is difficult to know for sure whether people take too many or too few steps to reduce disease transmission during our study period. Thus, we do not consider the optimal consumption of c. Instead, we hold constant the total amount of risky behavior and ask whether the allocation across types is optimal. Tangibly, this means that we are not asking, "how much social distancing should people be doing?" Instead, we are asking, "holding constant the amount of social distancing people are doing, would some people ideally be doing less, and others ideally be doing more?"

Social welfare is the sum of utility across all people in all states:

$$W_t = \sum_\theta \sum_\sigma \omega_{\theta\sigma t} V_\theta(\sigma_t). \tag{3}$$

Let C_t denote the total risky behavior at time t across all people. The (constrained) socially optimal outcome results from maximizing W_t subject to the constraint that $C_t = \bar{C}_t$. Let λ be the shadow price on that constraint; this reflects the loss from having too much or too little social distancing overall.

Consuming c imposes two types of externalities. First, it imposes a positive pecuniary externality, as travel, dining out, and other risky activities help keep firms in business and workers employed. Second, it imposes a negative externality by increasing the person's infection probability, which increases the expected stock of infected people in the next period (i_{t+1}), which increases other Susceptible people's infection risk. Let ϕ_t denote the net externality per unit of consumption, which may be positive or negative; this becomes more negative as the contagion externality grows. We assume that these externalities are constant across people, and that people do not account for them when setting their c_t^*.

In the constrained social optimum, Susceptible people's consumption of c_t would satisfy the following first-order condition:

$$0 = \underbrace{u'_\theta - \beta i_t \delta \left(V_\theta(S) - V_\theta(I)\right)}_{\text{private marginal utility}} + \underbrace{\phi_t}_{\text{externality}} + \underbrace{\lambda}_{\text{shadow price}} . \tag{4}$$

3.4 Heterogeneous Risk Misperceptions

We now allow people to misperceive risks. These misperceptions cause people to choose too much or too little risky behavior relative to their private optimum, and heterogeneous misperceptions cause transfers across types and efficiency losses.

We now add θ subscripts to explicitly denote different parameters by type. Let $\mu_{t\theta} := \beta i_t \delta \left(V_\theta(S) - V_\theta(I)\right)$ denote type θ's expected utility cost due to infection from an additional unit of risky consumption. Let $\tilde{\mu}_{t\theta}$ denote type θ's *perception* of that cost. Susceptible type θ consumers then set $c_{t\theta}$ according to the following modified first-order condition:

$$u'_\theta = \tilde{\mu}_{t\theta}, \tag{5}$$

giving consumption denoted $c^*_{t\theta}$.

For illustrative purposes, imagine that there are two types $\theta \in \{a, b\}$ in equal proportion, and that period t marginal utility is linear and the same for both types, so $u'_\theta(c) = u'(c)$ for both types and u'' is a constant. Finally, without loss of generality, assume that type a perceives greater risk, so $\tilde{\mu}_{a\theta} > \tilde{\mu}_{b\theta}$. Our survey data show that Democrats perceive greater risk, so one can think of Democrats as type a. We do not take a stand on which type perceives risk more correctly or which type's behavior is closer to the unconstrained social optimum.

Define $\bar{\tilde{\mu}}_t := \frac{1}{2}\left(\tilde{\mu}_{ta} + \tilde{\mu}_{tb}\right)$ as the average risk perception. With homogeneous risk perceptions, both types would set c_t such that $u' = \bar{\tilde{\mu}}_t$, giving homogeneous consumption denoted $\bar{c}_t$. With heterogeneous misperceptions, type a consumes more and type b consumes less; the consumption difference is $c^*_{tb} - c^*_{ta} = \frac{\tilde{\mu}_{ta} - \tilde{\mu}_{tb}}{-u''}$. These consumption differences cause both transfers across types and efficiency losses.

Risk perceptions affect risky consumption, and risky consumption causes externalities, so the heterogeneous misperceptions cause transfers across groups. The net transfer from type a to type b from heterogeneous instead of homogeneous misperceptions is

$$\underbrace{\frac{\tilde{\mu}_{ta} - \tilde{\mu}_{tb}}{-u''}}_{\text{consumption difference}} \cdot \underbrace{\phi_t}_{\text{externality}} . \tag{6}$$

If $\phi_t > 0$, i.e. the positive pecuniary externality from risky consumption outweighs the negative contagion externality, then heterogeneous misperceptions cause a net transfer from type b to type a. Intuitively, we would say that Republicans are doing more to keep the economy going. On the other hand, if $\phi_t < 0$, i.e. the negative contagion externality outweighs the positive pecuniary externality, then heterogeneous misperceptions cause a net transfer from type a to type b. Intuitively, we would say that Democrats are doing more to reduce the spread of disease.

The efficiency cost in period t from heterogeneous instead of homogeneous misperceptions are the two deadweight loss triangles around $\bar{c}_t$, with total area:

$$\Delta W_t = \frac{s_t}{2} \cdot \frac{\left(\overbrace{\tilde{\mu}_{ta} - \bar{\tilde{\mu}}_t}^{\text{misperception}}\right)^2}{\underbrace{-u''}_{\text{slope of private marginal utility}}}. \tag{7}$$

Intuitively, type a people (Democrats) are doing too much social distancing, and type b (Republicans) too little, relative to the (constrained) social optimum with homogeneous risk perceptions. The marginal cost of social distancing is increasing: it's easy to start by avoiding going to a bar once a week, but eventually one's only contact with people is going to the grocery store for food, and it is quite costly to stop buying food. Thus, society could achieve the same amount of social distancing at lower cost if type a did less and type b did more.

This model informs the empirical tests in the rest of the paper. In Section 6, we ask if Democrats and Republicans have different risk perceptions, which would generate the transfers and efficiency costs described above. In doing so, we control for factors such as population density that could generate difference in actual risks across types. In Sections 5 and 6, we ask if Democrats and Republicans are reducing risk by different amounts. In doing so, we use proxies to control for differences in actual risks and marginal costs of risk reduction that could cause differential risk reduction to be socially optimal.

4 Data

4.1 SafeGraph Mobile GPS Location Data

Our analysis uses GPS data from SafeGraph, aggregating GPS pings from numerous mobile applications to measure foot traffic patterns to a collection of points-of-interest (POIs). POIs include retail shops, restaurants, movie theaters, hospitals, and many other public locations individuals may choose to go when leaving their house. For each POI, SafeGraph reports its geographic location, industry, and the total number of visitors in their mobile device panel that have visited each day.[3] SafeGraph also provides the number of visitors traveling from a given home census block group (CBG) to a given POI over the course of a week or month.[4]

Our primary analysis uses data from a period of nine weeks, from January 26 to March 28, 2020. We aggregate visits across all POIs in a given county for a given week. We also separately aggregate visits by 2-digit NAICS code for each county and week. In a placebo analysis, we analyze data over earlier time periods (starting in January 2019).

We also use data from the SafeGraph Social Distancing data released as a part of their COVID-19 response. This data is available since February 1, 2020 and updated regularly. We use data through March 30, 2020, which gives us nine weeks of data.[5] This data contains alternative measures of social distancing beyond POI visits, such as the number of devices leaving their assigned geohash-7 home or the median time spent away from home across devices.

See the Appendix for additional information on the SafeGraph data construction.

We supplement the SafeGraph data with various other sources of county and census block group data. For demographic information on age, race, education, income, and poverty status at

[3] SafeGraph removes POIs with fewer than five visitors in a given month for data through February 2020. For the March 2020 data, SafeGraph has released data on a weekly basis, rather than a monthly basis, and include all POIs with at least 1 visitor for these weekly releases.

[4] To preserve anonymity, SafeGraph removes home CBG↔POI pairs with fewer than five visitors in a given month for data through February 2020. For the March 2020 data, SafeGraph has released data on a weekly basis, rather than a monthly basis, and again excludes CBG↔POI pairs with fewer than 5 visitors over this (now weekly) time horizon.

[5] The last week of our social distancing data only contains March 28th through March 30th.

the county-level, we aggregate census block group data from SafeGraph Open Census to the county level.[6] For each county, we define county partisanship to be the proportion of total votes received by President Donald Trump in the 2016 election (MIT Election Data and Science Lab 2018). We use county-level data on COVID-19 cases and deaths from The New York Times (2020).

4.2 Survey

To supplement these data, we run an online survey with a sample of American adults to study partisan gaps in beliefs about and responses to COVID-19 at the individual level. The survey was conducted from April 4-6 on CloudResearch's Prime Panels, a market research firm with access to 50 million participants. As of April 6, we have recruited 1,665 participants to complete the study. Participants are broadly representative of U.S. adults in terms of party affiliation, age, gender, and race.[7] Subjects who completed the survey are paid a show-up fee from CloudResearch and have the chance to earn additional bonus incentives of up to $100.

Participants were asked for their party affiliation on a seven-point scale, ranging from "Strongly Democrat" to "Strongly Republican." We interpret party continuously, where 0 represents "Strongly Democrat" and 6 represents "Strongly Republican." We also classify participants into Republican (including independents who lean Republican) and Democrats (including independents who lean Democrat) for some analyses; there are 40 percent and 47 percent of these groups respectively.

The survey asked for demographic information (zipcode, age, race, gender, income, education, number of children, and health). It then asked about news consumption habits and trust before and during COVID-19. Then, there are several questions about social distancing: self-reported social distancing in response to COVID-19, beliefs about the risk of not distancing, and the appropriate trade-off between going out more to help the economy versus going out less to avoid spreading COVID-19.

We next elicited beliefs about the number of new COVID-19 cases that would be confirmed in the US in April, 2020, as well as the approval rating of Donald Trump's response to the pandemic on April 30, and randomly vary whether these are incentivized or not. 847 (51 percent) of subjects make incentivized predictions in which they earn more money if they are closer to the correct answer. They are told that we will randomly select 10 participants who will receive a payment

[6]The SafeGraph Open Census data is derived from the 2016 5-year ACS at the census block group.

[7]When noted, we weight observations so that age, gender, and race averages match the 2010 Census data, and party affiliation matches the Gallup survey from March 13-22, 2020 (Gallup 2020).

of (\$100 $-\Delta$) where Δ is the percentage point difference between their answer and the true value. The remaining 818 (49 percent) of subjects are not incentivized.

More details and additional questions are discussed in our pre-analysis plan (AEA Registry, ID: AEARCTR-0005632).

5 SafeGraph Empirical Specification and Results

Figure 4 reports trends in social distancing and COVID-19 prevalence separately for Republican and Democratic counties, defined to be counties above or below the median 2016 Republican vote share respectively. Panel A shows that the overall number of POI visits is relatively constant until COVID-19 cases begin emerging in the United States in March. During this same period, Democratic counties exhibited a sharper drop in weekly POI visits than their Republican counterparts. As Panel B demonstrates, Democratic counties also exhibited a much sharper rise in COVID-19 cases and deaths—accounting for nearly all verified COVID-19 cases and deaths through March 22. Appendix Figure A1 shows that these declining and differential POI trends are not present over the same time period in 2019.

Our main empirical specification takes the following form

$$\log(c_{it}) = \alpha_t \rho_i + X_{it} \cdot \gamma_t + \varepsilon_{it},$$

where c_{it} is the number of POI visits in county i during week t, α_t are the time-varying coefficients on county partisanship ρ_i, X_{it} are potentially time-varying controls, and ε_{it} is the county-specific error term.[8] In choosing our control variables X_{it}, we chose variables to flexibly control for the four channels of divergent behavior highlighted in equation (2). Standard errors are clustered at the county-level throughout unless specified otherwise.

Figure 5 reports our estimates of α_t under various sets of covariates chosen to incrementally control for the mechanisms highlighted by our model.

In Panel A, we only include county and time fixed effects. This measures the extent to which these two groups' behavior diverges with the rise of COVID-19 via any of the aforementioned channels. Throughout February, there are no significant partisan differences in POI visits relative to the January 26 week baseline. However, as COVID-19 begins to emerge in the United States, partisan differences arise and grow throughout the weeks of March.

[8]In implementing, we normalize α_t relative to the first week.

These results do not control for differences in state policies, which themselves may be a function of the partisan leanings of government officials. In Panel B, adding state-time fixed effects to control for state-level policies in response to COVID-19 along with other state-level temporal shocks causes the partisan differences to attenuate slightly.

In Panel C, we flexibly control for various health[9] and economic[10] characteristics of the county. We view the health controls as proxies for the marginal infection probability and the private cost of infection, and we view the economic controls as proxies for the marginal cost of social distancing, though each group of controls could proxy for other factors as well. We allow the coefficient on these variables γ_t to vary flexibly across time.

Although these controls attenuates the partisan differences to some degree, they remain economically and statistically significant. By the week of March 22, our estimate of α_t is 0.419. This implies that going from a county with the 10th to the 90th percentile in Republican vote share is associated with an 16.7 percent increase in the number of POI visits during the week of March 22.[11] Appendix Figure A2 shows that these strong partisan differences do not appear over the same time period in 2019. We view these results as evidence of behavioral differences driven by partisan misperceptions of risks at the group-level, consistent with the survey evidence.

In Appendix Figure A3, we report sensitivity to various alternative specifications. Panels A and B use alternative sets of controls. Panel C replaces the measure of partisanship with a discrete indicator for certain quantiles of the Republican vote share distribution. Panel D drops counties with populations above half a million or states with early COVID-19 outbreaks (California, Washington, and New York). Panel E restricts the sample to counties from certain portions of the Republican vote share distribution. And, Panel F weights observations by the county's population, uses standard errors clustered at the state-level, and examines sensitivity to the start date. None of the alternative specifications change the central conclusion regarding partisan differences in social distancing in March.

Figure 6 examines heterogeneity across 2-digit NAICS codes by re-aggregating POI visits to the county level after restricting to certain NAICS codes. Consistent with the narrative around COVID-19, we see the strongest partisan differences emerge with POIs in the accommodations

[9]Health controls include log of one plus the number of confirmed COVID-19 cases in the county, the log of one plus the number of COVID-19 deaths in the county, the log of one plus the county population density (individuals per square kilometer) and the share of the population 65 years or older

[10]Economic controls include the share of the population with at least a bachelor's degree, the share in poverty, and the shares of white, black, and asians.

[11]The difference between the 90th and 10th percentile of Republican vote share is 0.806 - 0.407 = 0.399.

and food, entertainment, and retail industries. There are no significant partisan differences in visits to health care POIs.

Figure 7 repeats Panel C of 5, but using POI visits aggregated at the day level. The partisan differences emerge in March for both weekdays and weekends, suggesting these differences are not driven solely by differences in work-from-home policies.

Figure 8 considers various alternative measures of social distancing derived from SafeGraph's Social Distancing data release as described in Section 4. Statistically significant partisan differences emerge in March for the log number of devices leaving home, the share of devices leaving home, and the total number of active devices.[12] For the log of the median time away from home, we see positive, but insignificant, point estimates.

6 Survey Results

Turning to the results of our survey, we first consider beliefs about future COVID-19 cases in the US. We tell participants the number of cases by March 31 and ask them to predict the number of cases in April. For half of subjects, these predictions are incentivized. We find that Democrats anticipate more future COVID-19 cases. On average, Republicans predict 188,678 cases (s.e. 8,069 cases) and Democrats predict 13 percent more cases (213,761 cases; s.e. 8,211 cases).[13] The difference between these predictions is 25,083 cases (s.e. 11,512 cases; p = 0.030).

Next, we consider how Republicans and Democrats differ in their perceptions of the risk of not socially distancing, and find that that Democrats believe that the probability of catching COVID-19 in one month without any social distancing is higher than Republicans do. On average, Republicans assess this probability to be 19.8 percent (s.e. 1.1 percent), and Democrats assess this probability to be 24.8 percent (s.e. 1.1 percent). The difference between these estimates is 5.0 percentage points (s.e. 1.5 percentage points; p = 0.001). We see also see small differences in self-reported actions. On average, Republicans report reducing their contact with others over one month by 70.1

[12] A key issue with the SafeGraph social distancing data is sample attrition. SafeGraph restricts the panel to devices with observed location pings in a given time period. For some applications, the frequency of location pings depends on device mobility. If devices are immobile at home or turned off, they may not generate location pings and would then be dropped from the sample. The total number of active devices changes over our sample period in a manner consistent with sample attrition. Given these issues, we prefer measures of social distancing derived solely from *external* activity (e.g., POI visits) that do not contain the same measurement error problems. We attempt to correct for the differential attrition in our measure of the share of devices leaving home (see Figure 8 footnotes for correction; see Panel G of Appendix Figure A3 for estimates using the uncorrected measure).

[13] As pre-specified, these averages are calculated after winsorizing at the 5-percent level to account for outliers.

percent (s.e. 1.0 percent), and Democrats report reductions of 71.5 percent (s.e. 0.8 percent). The difference between these reductions is 1.4 percent, but not statistically significant (s.e. 1.3 percent; p = 0.277).

We next consider Republicans' and Democrats' perceptions of the tradeoff between going outside and supporting the economy, versus staying inside and avoiding the spread of COVID-19. Participants were asked to answer which was more important on a seven-point scale, which we normalize to the unit interval. We find that a majority of participants of both parties find it more important to stay inside to stop the virus from spreading, but that they differ in their strength of opinion. Republicans on average give an answer that leans 80.6 percent of the way towards staying inside (s.e. 1.0 percent), while Democrats give an answer that leans 83.8 percent towards staying inside (0.9 percent). The difference between these answers is 3.2 percentage points (s.e. 1.3 percentage points; p = 0.017).

These results seem to indicate that Democratic participants generally expect COVID-19 to spread more, and be a greater threat, as compared to Republican participants. However, as with past surveys, there is a confound that Democrats tend to live in areas that are generally more susceptible to COVID-19. There may also be other demographic, health, or location differences between Republicans and Democrats that contribute to these gaps in beliefs and attitudes.

In Figure 9, we normalize participants' responses on the previous four questions, control for a large set of control variables (including population density), and weight data to mimic a nationally-representative sample on age, race, gender, and party. We find that the results persist, and that the control variables do not qualitatively change the predictions from the raw data.

Lastly, we ask whether the partisan gap in beliefs about the number of US cases shrinks when subjects are given incentives for accuracy. Bullock et al. (2015) and Prior et al. (2015) show that partisan differences on factual questions often shrink under incentives, and interpret this as evidence that the differences in survey responses are partly due to "partisan cheerleading" rather than differences in true beliefs. As shown in Figure 10, we find that on an explicitly political question (Trump's approval rating for his handling of COVID-19), incentives reduce the partisan gap, consistent with the findings in Bullock et al. (2015) and Prior et al. (2015). However, the partisan gap in predicted future cases if anything widens with incentives. This supports the view that Democrats and Republicans genuinely differ in their beliefs about the severity of the outbreak.

7 Conclusion

Divergent messages from political leaders and media outlets about the severity of COVID-19 have the potential to seriously undermine the country's response to the pandemic. If citizens on the left and the right disagree about the potential risks, they may also diverge in their compliance with critical social distancing measures. Many lives could be lost relative to a benchmark where the same overall level of social distancing was efficiently allocated between the two groups.

Our empirical results show that partisan gaps in beliefs and behavior are real. GPS evidence reveals large partisan gaps in actual social distancing behaviors. Survey evidence shows substantial gaps between Republicans and Democrats in beliefs about severity and the importance of social distancing. The raw partisan differences partly reflect the fact that Democrats are more likely to live in the dense, urban areas hardest hit by the crisis, and to be subject to policy restrictions—in other words, to face stronger individual incentives for social distancing. Even after controlling carefully for such factors, however, the partisan gaps remain statistically and economically significant. While our evidence does not permit us to conclusively pin down the ultimate causes of partisan divergence, the patterns are consistent with the messaging from politicians and media having played an important role.

References

Aleem, Zeeshan. 2020. A new poll shows a startling partisan divide on the dangers of the coronavirus. *Vox*. March 15. https://www.vox.com/2020/3/15/21180506/coronavirus-poll-democrats-republicans-trump.

Andersen, Martin. 2020. Early evidence on social distancing in response to COVID-19 in the United States. *Working Paper.*

Athey, Susan, Billy Ferguson, Matthew Gentzkow, and Tobias Schmidt, 2019. Experienced Segregation. *Working Paper.*

Baker, Scott R., R.A. Farrokhnia, Steffen Meyer, Michaela Pagel, Constantine Yannelis. 2020. How does household spending respond to an epidemic? Consumption during the COVID-19 pandemic. *Working Paper*.

Barrios, John M. and Yael V. Hochberg. 2020. Risk perception through the lens of politics in the time of the COVID-19 pandemic. *Working Paper.*

Bartels, Larry M. 2002. Beyond the running tally: Partisan bias in political perceptions. *Political Behavior.* 24(2): 117-150.

Beauchamp, Zack. 2020. The stunning contrast between Biden and Trump on coronavirus. *Vox.com*. https://www.vox.com/policy-and-politics/2020/3/12/21177135/coronavirus-covid-19-pandemic-trump-biden-speeches.

Benoit, Jean-Pierre, and Juan Dubra. 2014. A theory of rational attitude polarization. *Working Paper.*

Bish, Alison, and Susan Michie. 2010. Demographic and attitudinal determinants of protective behaviours during a pandemic: A review. *British Journal of Health Psychology*. 15(4): 797-824.

Blattman, Christopher, Donald Green, Daniel Ortega, and Santiago Tobón. 2017. Place-based interventions at scale: The direct and spillover effects of policing and city services on crime. *Working Paper.*

Blendon, Robert J., Lisa M. Koonin, John M. Benson, Martin S. Cetron, William E. Pollard, Elizabeth W. Mitchell, Kathleen J. Weldon, and Melissa J. Herrmann. 2008. Public response to community mitigation measures for pandemic influenza. *Emerging Infectious Diseases*. 14(5): 778.

Bullock, John G., Alan S. Gerber, Seth J. Hill, and Gregory A. Huber. 2015. Partisan bias in factual beliefs about politics. *Quarterly Journal of Political Science*. 10(4): 519-578.

Bults, Marloes, Desirée JMA Beaujean, Onno de Zwart, Gerjo Kok, Pepijn van Empelen, Jim E van Steenbergen, Jan Hendrik Richardus, and Hélène ACM Voeten. 2011. Perceived risk, anxiety, and behavioural responses of the general public during the early phase of the Influenza A (H1N1) pandemic in the Netherlands: results of three consecutive online surveys. *BMC Public Health*. 11(1): 2.

Campante, Filipe R., and Daniel A. Hojman. 2013. Media and polarization: Evidence from the introduction of broadcast TV in the United States. *Journal of Public Economics.* 100: 79-92.

Chen, M. Keith, and Ryne Rohla. 2018. Politics Gets Personal: Effects of Political Partisanship and Advertising on Family Ties. *Science.* 360: 1020-1024.

Chuang, Ying-Chih, Ya-Li Huang, Kuo-Chien Tseng, Chia-Hsin Yen, and Lin-hui Yang. 2015. Social capital and health-protective behavior intentions in an influenza pandemic. *PloS One*. 10(4).

Coppins, McKay. 2020. Trump's Dangerously Effective Coronavirus Propaganda. *TheAtlantic*.com. https://www.theatlantic.com/politics/archive/2020/03/trump-coronavirus-threat/607825/.

Davis, Donald R., Jonathan I. Dingel, Joan Monras, and Eduardo Morales. 2019. How segregated is urban consumption? *Journal of Political Economy.* 127(4): 1684-1738.

Dixit, Avinash K., and Jörgen W. Weibull. 2007. Political polarization. *Proceedings of the National Academy of Sciences.* 104(28): 7351-7356.

Dubé, Jean-Pierre, Zheng Fang, Nathan Fong, and Xueming Luo. 2017. Competitive price targeting with smartphone coupons. *Marketing Science*. 36(6): 944-975.

Economist. 2020. Democrats seem to take social distancing more seriously than Republicans. *Economist.com* April 4. https://www.economist.com/united-states/2020/04/04/democrats-seem-to-take-social-distancing-more-seriously-than-republicans.

Fineberg, Harvey V. 2014. Pandemic preparedness and response—lessons from the H1N1 influenza of 2009. *New England Journal of Medicine*. 370(14), 1335-1342.

Fryer Jr, Roland G., Philipp Harms, and Matthew O. Jackson. 2019. Updating beliefs when evidence is open to interpretation: Implications for bias and polarization. *Journal of the European Economic Association.* 17(5): 1470-1501.

Gadarian, Shana Kushner, Sara Wallace Goodman, and Thomas B. Pepinsky. 2020. Partisanship, health behavior, and policy attitudes in the early stages of the COVID-19 pandemic. *Working Paper.*

Gaines, Brian J., James H. Kuklinski, Paul J. Quirk, Buddy Peyton, and Jay Verkuilen. 2007. Same facts, different interpretations: Partisan motivation and opinion on Iraq. *The Journal of Politics*. 69(4): 957-974.

Gallup. 2020. Party Affiliation. *Gallup.* https://news.gallup.com/poll/15370/party-affiliation.aspx.

Gamma, Anna E., Jurgita Slekiene, Gregor von Medeazza, Fredrik Asplund, Placido Cardoso, and Hans-Joachim Mosler. 2017. Contextual and psychosocial factors predicting Ebola prevention behaviours using the RANAS approach to behaviour change in Guinea-Bissau. *BMC Public Health*. 17(1): 446.

Glaeser, Edward L., and Bryce A. Ward. 2006. Myths and realities of American political geography. *Journal of Economic Perspectives.* 20(2): 119-144.

Hanna, Rema, Gabriel Kreindler, and Benjamin A. Olken. 2017. Citywide effects of high-occupancy vehicle restrictions: Evidence from "three-in-one" in Jakarta. *Science.* 357(6346): 89-93.

Hersh, Eitan D. and Matthew N Goldenberg. 2016. Democratic and Republican physicians provide different care on politicized health issues. *Proceedings of the National Academy of Sciences.* 113(42): 11811–11816.

Holmes, Bev J. 2008. Communicating about emerging infectious disease: The importance of research. *Health, Risk & Society.* 10(4): 349-360.

Ibuka, Yoko, Gretchen B. Chapman, Lauren A. Meyers, Meng Li, and Alison P. Galvani. 2010. The dynamics of risk perceptions and precautionary behavior in response to 2009 (H1N1) pandemic influenza. *BMC Infectious Diseases*. 10(1): 296.

Iyengar, Shanto, Yphtach Lelkes, Matthew Levendusky, Neil Malhotra, and Sean J. Westwood. 2019. The origins and consequences of affective polarization in the United States. *Annual Review of Political Science.* 22: 129-146.

Kermack, William Ogilvy, and Anderson G. McKendrick. 1927. A contribution to the mathematical theory of epidemics. *Proceedings of the Royal Society of London. Series A, Containing Papers of a Mathematical and Physical Character.* 115(772): 700-721.

Kantrowitz, Alex. 2020. Conservative media still isn't sure what to think about the coronavirus. *Buzzfeed News.* March 18. https://www.buzzfeednews.com/article/alexkantrowitz/conservative-media-still-isnt-sure-coronavirus

Krupenkin, Masha. 2018. Does partisanship affect compliance with government recommendations? *Working Paper.*

Lee, Alicia. 2020. These states have implemented stay-at-home orders. Here's what that means for you. *CNN*. https://www.cnn.com/2020/03/23/us/coronavirus-which-states-stay-at-home-order-trnd/index.html. Accessed on April 4, 2020.

Lerman, Amy E., Meredith L. Sadin, and Samuel Trachtman. 2017. Policy uptake as political behavior: Evidence from the Affordable Care Act. *American Political Science Review.* 111(4): 755–770.

Marist. 2020. March 13th and 14th survey of American adults. http://maristpoll.marist.edu/wp-content/uploads/2020/03/NPR_PBS-NewsHour_Marist-Poll_USA-NOS-and-Tables_2003151338.pdf.

McCarthy, Tom. 2020. Disunited states of America: Responses to coronavirus shaped by hyper-partisan politics. *The Guardian*. Mar 29. https://www.theguardian.com/us-news/2020/mar/29/america-states-coronavirus-red-blue-different-approaches.

McCarty, Nolan, Keith T. Poole, and Howard Rosenthal. 2006. *Polarized America: The Dance of*

Political Ideology and Unequal Riches. MIT Press: Cambridge, MA.

MIT Election Data and Science Lab. 2018. County Presidential Election Returns 2000-2016. *Harvard Dataverse*. https://doi.org/10.7910/DVN/VOQCHQ.

Montoya-Williams, Diana and Elena Fuentes-Afflick. 2019. Political determinants of population health. *JAMA Network Open.* 2(7): e197063.

Ortoleva, Pietro, and Erik Snowberg. 2015. Overconfidence in political behavior. *American Economic Review.* 105(2): 504-35.

Piacenza, Joanna. 2020. Tracking public opinion on the coronavirus. *Morning Consult*. https://morningconsult.co public-opinion-on-the-coronavirus/.

Prior, Markus. 2013. Media and political polarization. *Annual Review of Political Science*. 16: 101-127.

Prior, Markus, Gaurav Sood, and Kabir Khanna. 2015. You cannot be serious. The impact of accuracy incentives on partisan bias in reports of economic perceptions. *Quarterly Journal of Political Science*. 10(4): 489–518.

Saad, Lydia. 2020. Americans step up their social distancing even further. *Gallup*. https://news.gallup.com/opinio step-social-distancing-even-further.aspx.

Sances, Michael W. and Joshua D. Clinton. 2019. Who participated in the ACA? Gains in insurance coverage by political partisanship. *Journal of Health Politics, Policy and Law.* 44(3): 349–379.

Shultz, James M., Janice L. Cooper, Florence Baingana, Maria A. Oquendo, Zelde Espinel, Benjamin M. Althouse, Louis Herns Marcelin et al. 2016. The role of fear-related behaviors in the 2013–2016 West Africa Ebola virus disease outbreak. *Current Psychiatry Reports*. 18(11): 104.

Stanley-Becker, Isaac and Chelsea Janes. 2020. As virus takes hold, resistance to stay-at-home orders remains widespread – exposing political and social rifts. *Washington Post*. April 2. https://www.washingtonpost.com/politics/as-virus-takes-hold-resistance-to-stay-at-home-orders-remains-widespread–exposing-political-and-social-rifts/2020/04/02/d87314e0-7436-11ea-85cb-8670579b863d_story.html

Suryadevara, Manika, Cynthia A. Bonville, Donald A. Cibula, Joseph B. Domachowske, and Amar C. Suryadevara. 2019. Associations between population based voting trends during the 2016 US presidential election and adolescent vaccination rates. *Vaccine*. 37(9): 1160–1167.

Taylor, Melanie, Beverley Raphael, Margo Barr, Kingsley Agho, Garry Stevens, and Louisa Jorm. 2009. Public health measures during an anticipated influenza pandemic: factors influencing willingness to comply. *Risk Management and Healthcare Policy*. 2: 9.

The New York Times. 2020. Coronavirus (Covid-19) Data in the United States. https://github.com/nytimes/covid-19-data. Accessed on April 2, 2020.

Trachtman, Samuel. 2019. Polarization, participation, and premiums: How political behavior helps explain where the ACA works, and where it doesn't. *Journal of Health Politics, Policy and Law.* 44(6): 855–884.

van der Weerd, Willemien, Daniëlle RM Timmermans, Desirée JMA Beaujean, Jurriaan Oudhoff, and Jim E. van Steenbergen. 2011. Monitoring the level of government trust, risk perception and intention of the general public to adopt protective measures during the influenza A (H1N1) pandemic in the Netherlands. *BMC Public Health*. 11(1): 575.

Vaughan, Elaine, and Timothy Tinker. 2009. Effective health risk communication about pandemic influenza for vulnerable populations. *American Journal of Public Health*. 99(S2): S324-S332.

Figure 1: Partisan Differences in Perceived Risk and Social Distancing

Note: This figure shows responses to nationally representative polls by political affiliation. Panel A shows the share of people concerned about coronavirus spreading to the United States (Piacenza 2020). Panel B shows self-reported behavior change as of March 13-14 (Marist 2020). Panel C shows the share of people avoiding public places, such as stores and restaurants (Saad 2020). Panel D shows that share of people avoiding small gatherings, such as with friends and family (Saad 2020).

Figure 2: Geographic Variation in Social Distancing, Partisanship, COVID-19, and Public Policy

Panel A: % Change in SafeGraph Visits

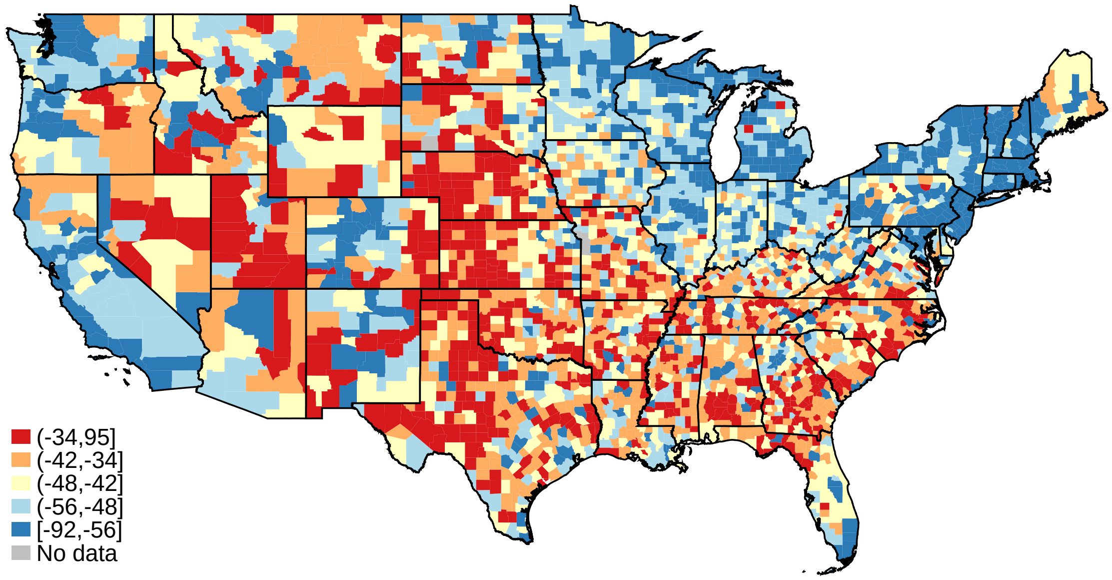

Panel B: 2016 Republican Vote Share

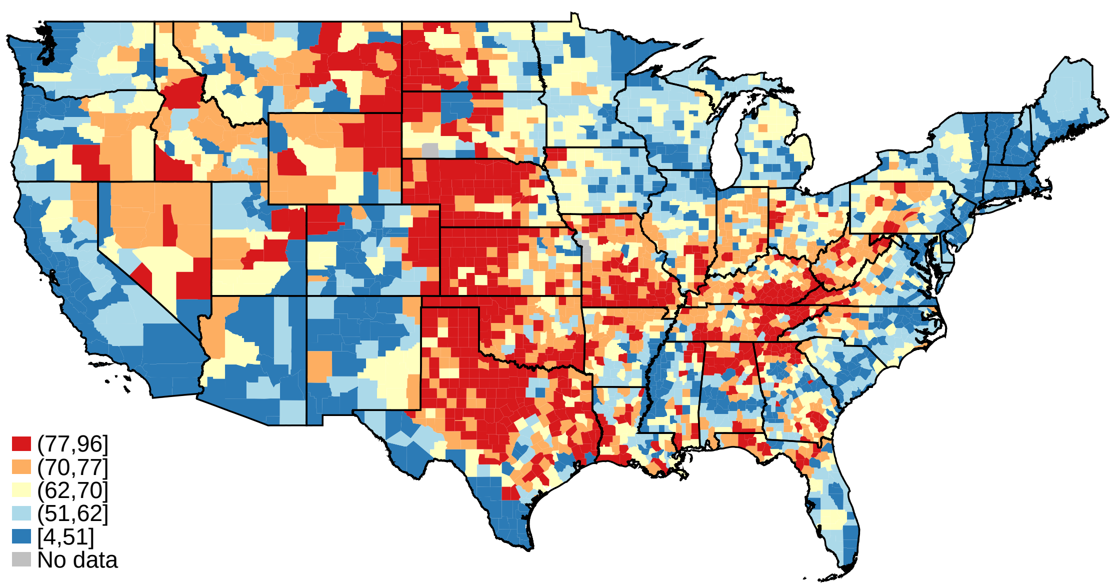

Note: This figure shows the U.S. geographic distribution of social distancing, political affiliation, COVID-19, and public policy responses. Panel A shows for each county the percent change in aggregate visits between the week beginning January 26, 2020 and the week beginning March 22, 2020. Blue shading denotes a more negative percent change in visits during the latter week relative to the former. Red shading indicates an increase or a smaller decrease in visits. These visits are sourced from SafeGraph's mobile device location data. Panel B maps counties by the percentage of votes Donald Trump received in the 2016 presidential election. Red shading in this panel indicates more Republican counties (higher Trump vote share), and blue shading indicates more Democratic counties (lower Trump vote share).

Figure 3: Geographic Variation in Social Distancing, Partisanship, COVID-19, and Public Policy cont.

Panel C: COVID Cases Confirmed by 3/28/2019

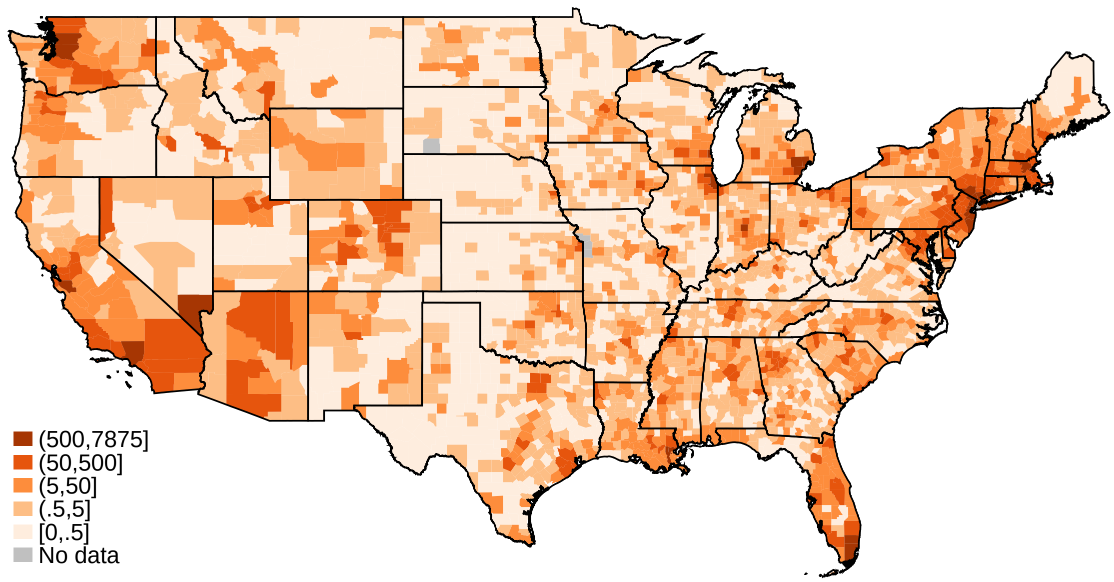

Panel D: State "Stay at Home" Order Start Date

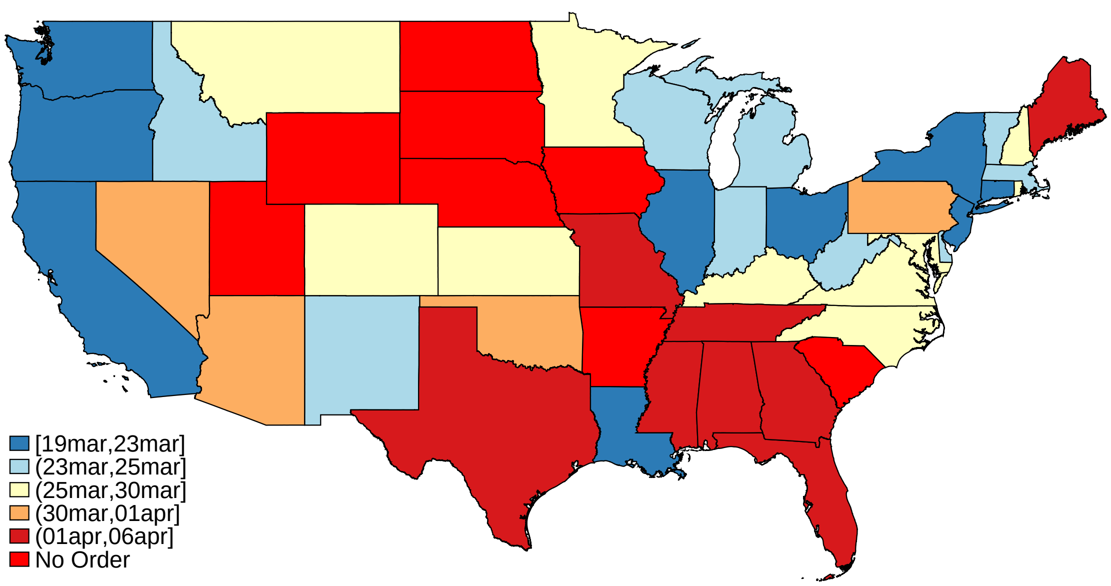

Note: This figure shows the U.S. geographic distribution of social distancing, political affiliation, COVID-19, and public policy responses. Panel C shows for each county the number of COVID-19 cases confirmed by March 28, 2020 (sourced from the New York Times). Panel D shades U.S. states by the effective start date for the earliest statewide "stay-at-home" order issued (Lee 2020). Blue shading indicates an earlier order, while red shading indicates that an order was issued later or was never issued.

Figure 4: Social Distancing and COVID-19 Prevalence

Panel A: POI Visits

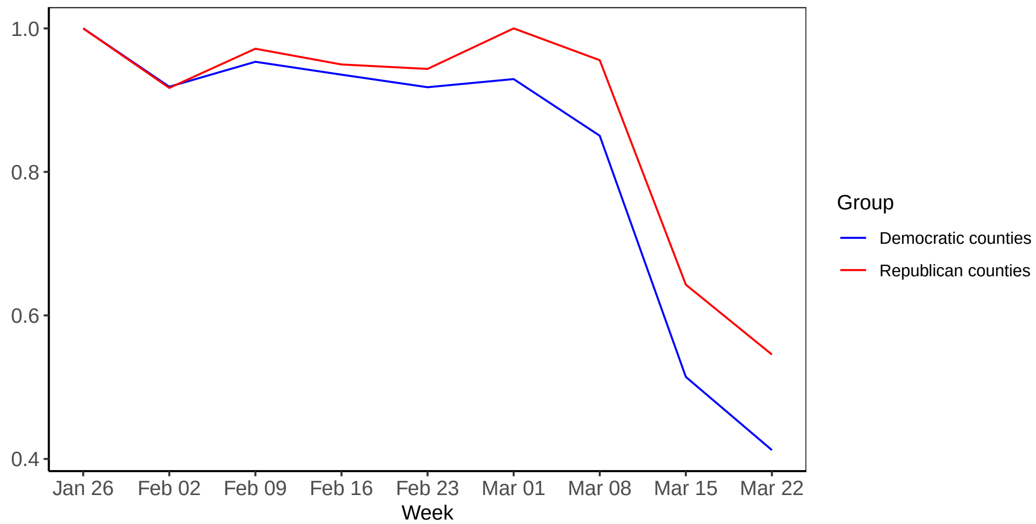

Panel B: COVID-19 Cases and Deaths

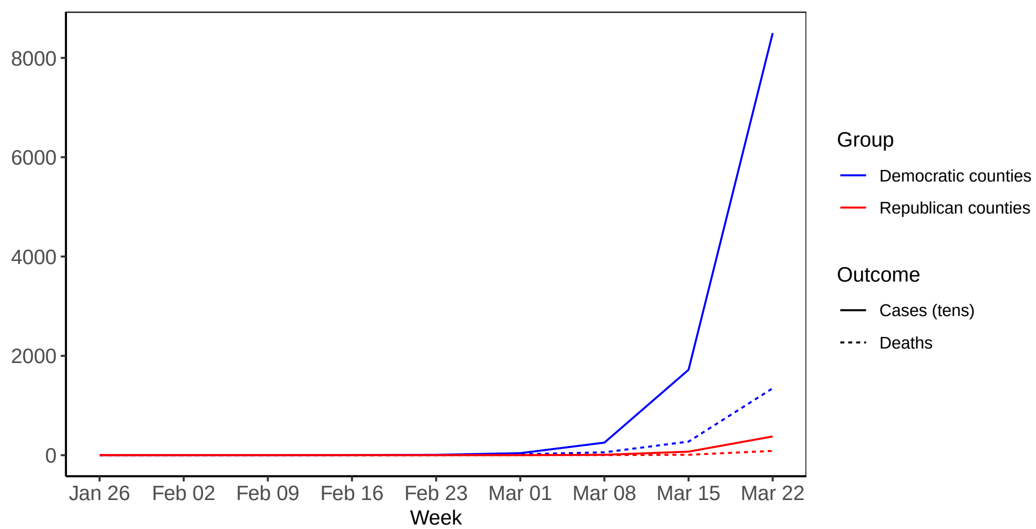

Note: Panel A shows the number of visits (normalized to one) to SafeGraph POIs for each week since January 26, 2020 for Republican counties and Democratic counties separately. Panel B is analagous but plots COVID-19 cases (in tens) and COVID-19 deaths. Republican counties are defined to be those whose 2016 Republican vote share is greater than the median vote share across the counties in our sample.

Figure 5: Partisan Differences in Social Distancing

Panel A: Only County & Time FE

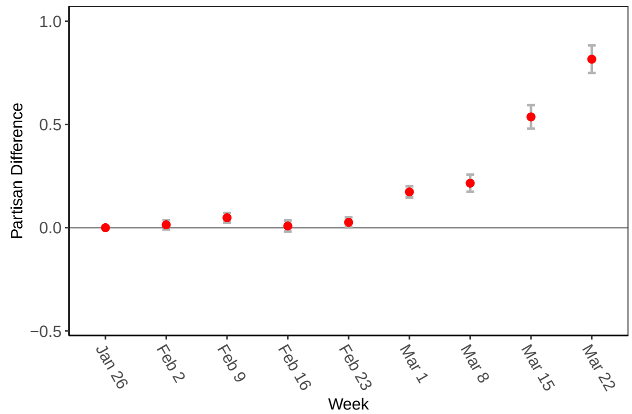

Panel B: Adds State-Time FE

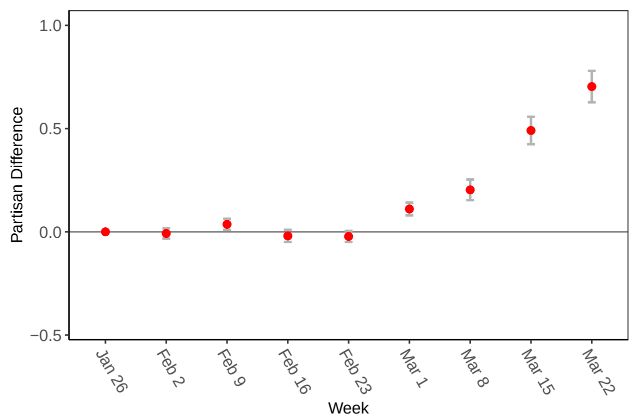

Panel C: Adds Health + Econ Controls

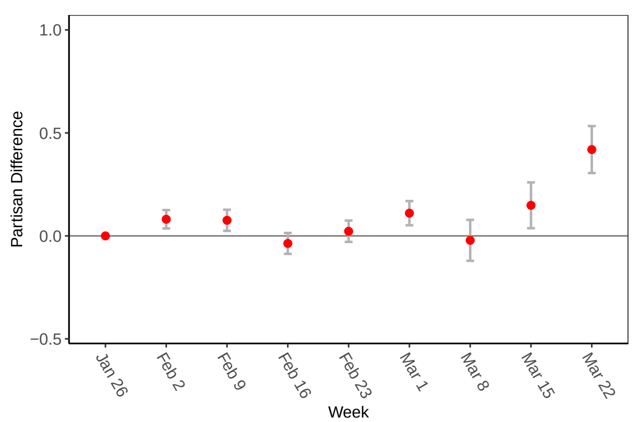

Note: Figure shows the estimated coefficients for county partisanship ρ_i on the log number of POI visits in the county using the specification outlined in the main text. For Panel A, only county and time fixed effects are included as controls. Panel B is the same as Panel A except state-time fixed effects replace the time fixed effects. Panel C is the same as Panel B except the health and economic covariates are included. The grey error bars indicate 95 percent confidence intervals constructed using standard errors clustered at the county-level.

Figure 6: Partisan Differences in Social Distancing by 2-Digit NAICS Code Industry

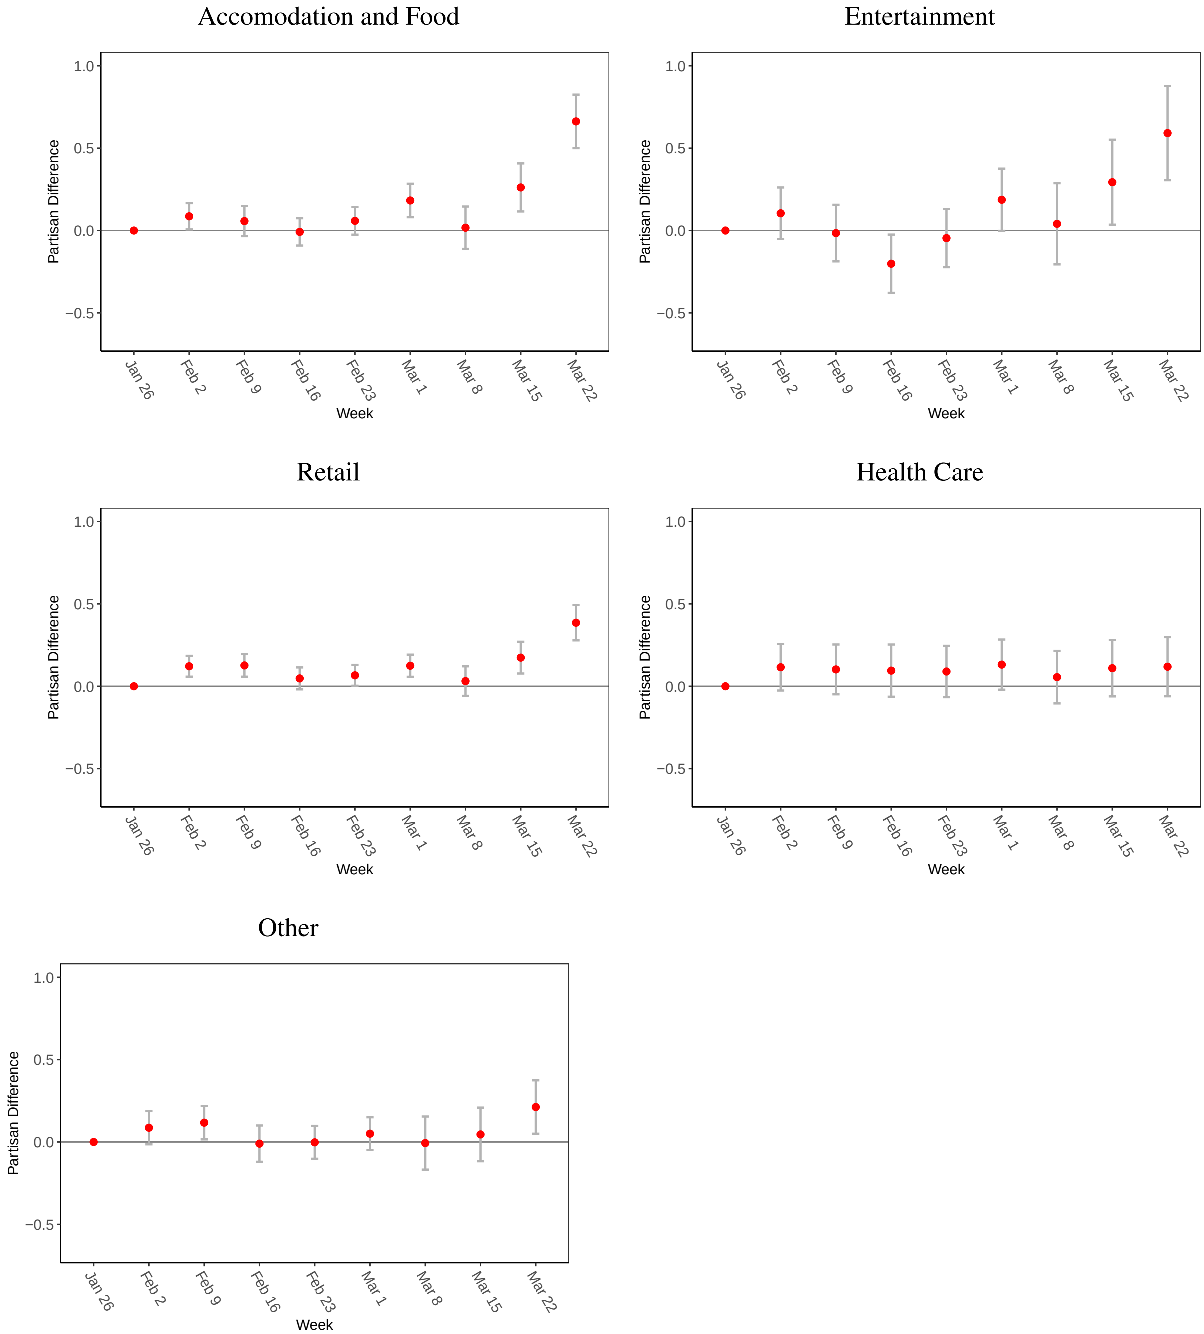

Note: Figure shows the estimated coefficients for county partisanship ρ_i on the log number of POI visits in the county after restricting POI visits to various 2-digit NAICS codes. The NAICS code groups are: Accomodation and Food (NAICS 72), Entertainment (NAICS 71), Retail Trade (NAICS 44 and 45), Health Care (NAICS 62), and Other Industries (All NAICS codes not previously used). The same controls are used as in Panel C of Figure 5. The grey error bars indicate 95 percent confidence intervals constructed using standard errors clustered at the county-level.

Figure 7: Partisan Differences in Social Distancing, Daily

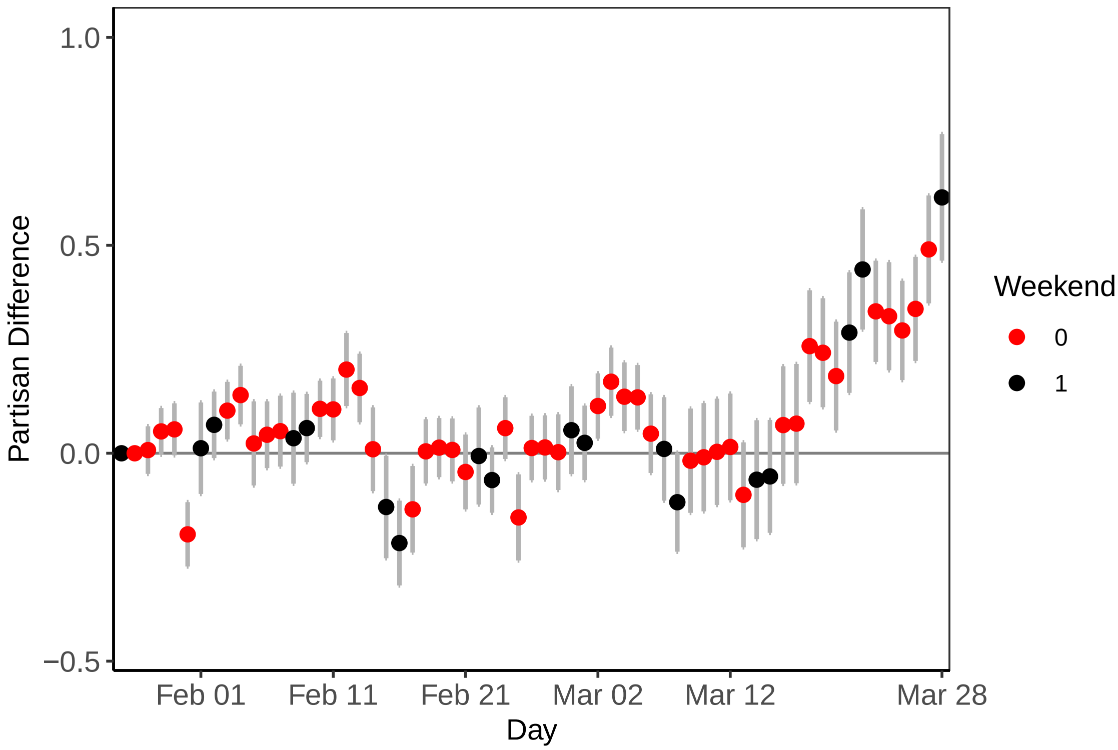

Note: Figure shows the estimated coefficients for county partisanship ρ_i on the log number of POI visits in the county. The same controls as in Panel C of Figure 5 are used except that state-time fixed effects occur at the day level and the weekday and weekend series are normalized separately. The grey error bars indicate 95 percent confidence intervals constructed using standard errors clustered at the county-level.

Figure 8: Partisan Differences in Social Distancing, Alternative Measures

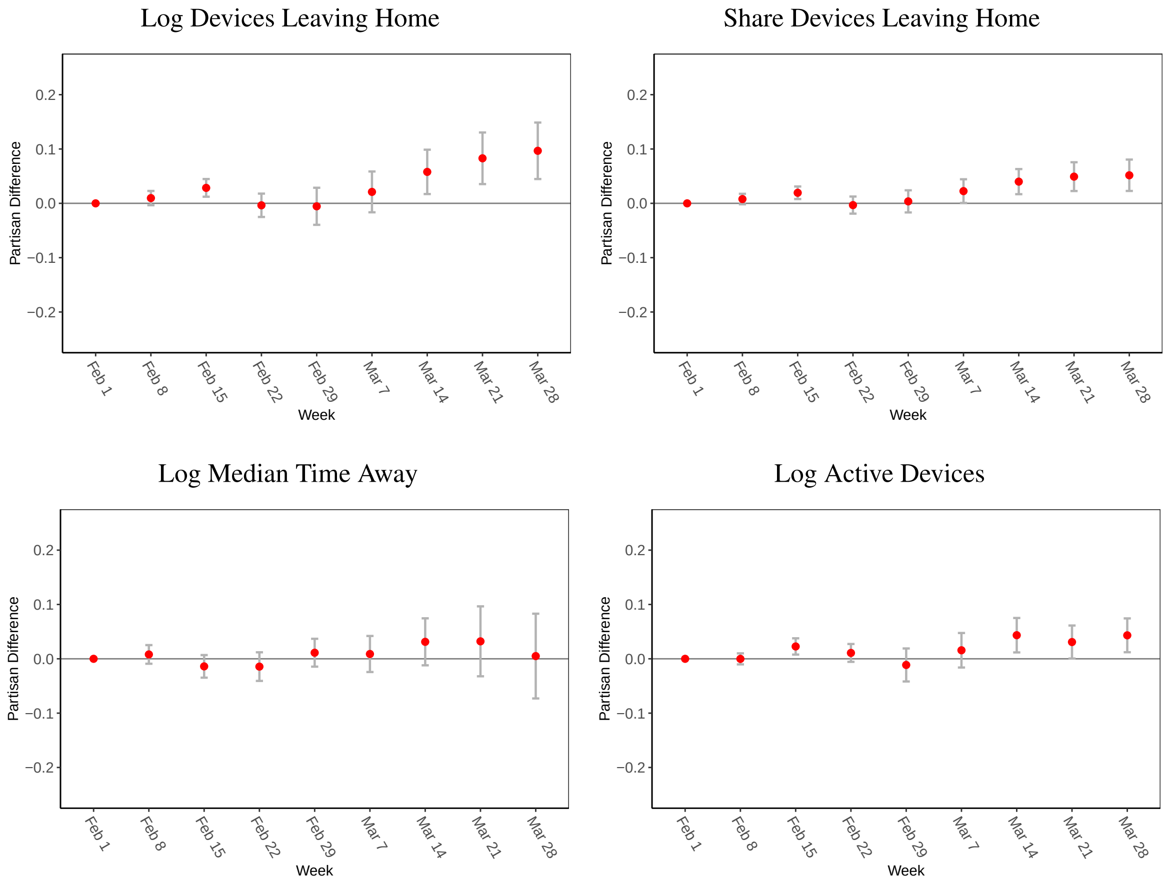

Note: Figure shows the estimated coefficients for county partisanship ρ_i on various alternative outcomes constructed from the Daily Social Distancing dataset from SafeGraph. 'Log Devices Leaving Home' is the log of one plus the number of active devices in the panel minus the active devices never observed leaving their geohash-7 home. 'Share Devices Leaving Home' is defined to be $1 - \frac{\max\{0, \text{home devices} + (\text{initital device count} - \text{current device count})\}}{\text{initial device count}}$, where 'home devices' are active devices never observed leaving their geohash-7 home, 'initial device count' is the number of active devices for the week of February 1, and 'current device count' is the number of active devices for the current week. 'Log Median Time Away' is $\log(1 + 1440 - \text{time home})$ where 'time home' is the median observed time at home across devices. 'Log Active Devices' is the log of one plus the number of active devices in the panel. The same controls are used as in Panel C of Figure 5. The grey error bars indicate 95 percent confidence intervals constructed using standard errors clustered at the county-level.

Figure 9: Partisan Differences in Beliefs and Actions

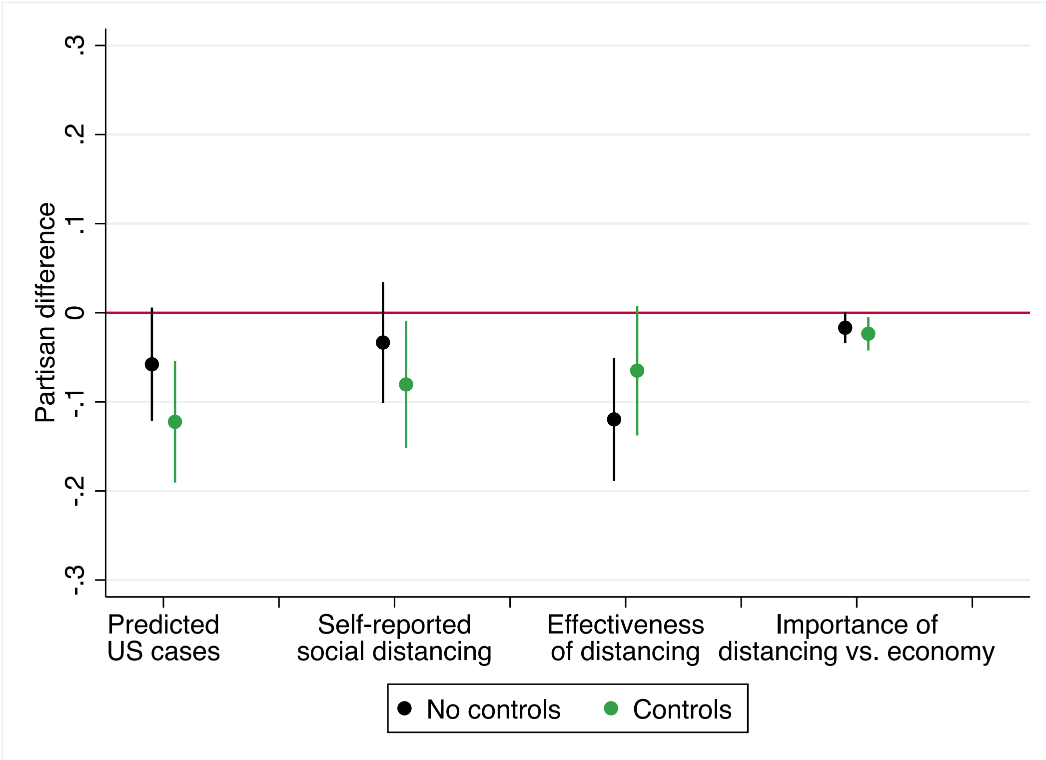

Note: This figure shows coefficient plots of regressing normalized measures of beliefs and actions on party. Positive values indicate less concern about COVID-19 or social distancing. Demographic controls are age, race, income, education, number of children, ZIP code logged population density, state. 2 percent of observations are dropped for not including a valid ZIP code. Predicted US cases are predictions about the number of new COVID-19 cases in the US in April; self-reported social distancing is the percent reduction in contact with others over one month; effectiveness of distancing is the estimated likelihood of catching COVID-19 in one month without social distancing; importance of distancing vs. economy is subjects' perception of whether it is more important to go out and stimulate the economy versus staying in and preventing the spread of COVID-19. Error bars represent 95 percent confidence intervals.

Figure 10: Effect of Incentives on Beliefs

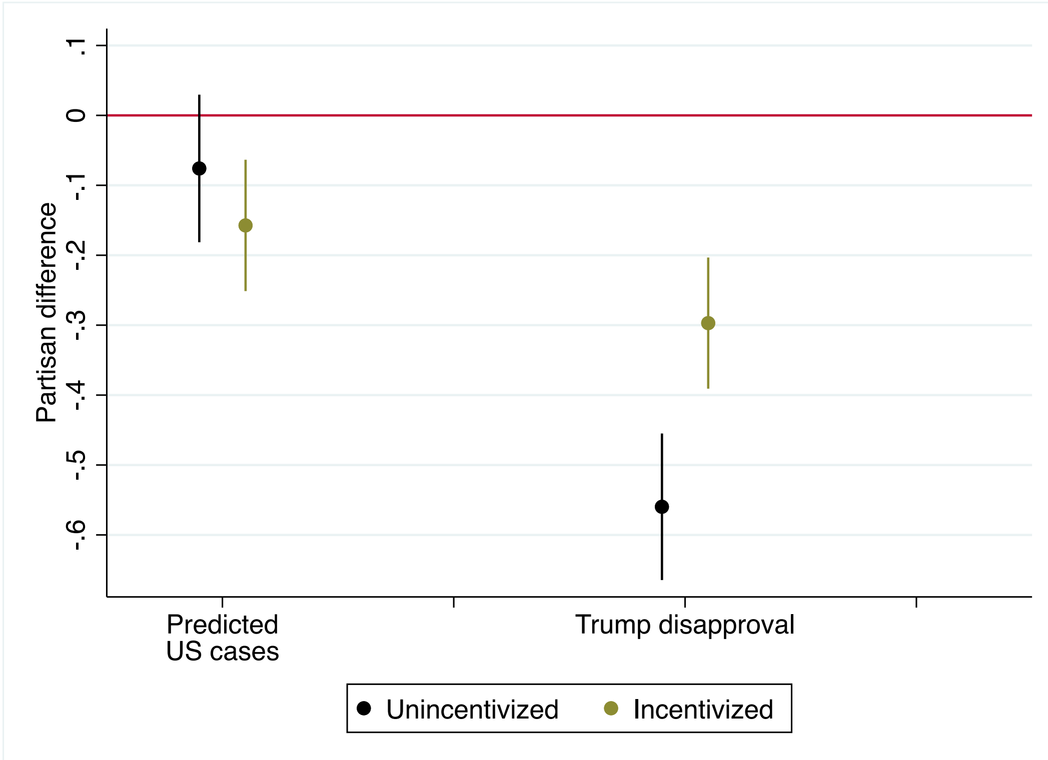

Note: This plot shows coefficient plots of regressing beliefs on party, with and without incentives for getting close to the correct answer. Trump disapproval is a low-stakes question that is susceptible to partisan cheerleading (Bullock et al. 2015; Prior et al. 2015). These results show that predicting COVID-19 cases does not appear susceptible to the same behavior. Error bars represent 95 percent confidence intervals.

A Appendix

A.1 Data Details

A.1.1 County-Level Data Build (POI)

To construct the county-level POI dataset used in the analysis, we proceeded as follows:

1. We use county data on 2016 Presidential votes shares (MIT Election Data and Science Lab 2018). We define the Republican vote share to be the share of votes received by the Republican candidate over the sum of votes across all candidates. We exclude Alaska, and merge with the 2010 TIGER county shapefile.[14] Two counties in the shapefile do not have valid vote data (FIPS: 15005, 51515).

2. We then use the latitude and longitude in the the Core POI dataset from SafeGraph to match POIs to counties. We successfully assign more than 99.9 percent of the POIs to a county.

3. We merge the output from (2) with the Patterns dataset from SafeGraph using the safegraph-place-id variable. We drop all observations with invalid vote shares at this stage. Using the March 22 week as an example, we only lose 0.34 percent of daily POI observations with these steps.

4. We use the Open Census data from SafeGraph to construct a county-level dataset of demographic information at the county level. We do this by aggregating up the data given at the census block group level to the county level. We then merge the county demographic information with the output from (3). We successfully match all but one county in this merge.

5. We then merge The New York Times COVID-19 tracking data onto our output from (4). We assume zero cases and deaths for the observations not observed in The New York Times data. We drop the five counties associated with New York City and the four counties which overlap with Kansas City (MO), because The New York Times lists these as geographic exceptions where it either does not assign cases to these counties or excludes cases occurring within the city.

A.1.2 County-Level Data Build (Social Distancing)

To construct the county-level social distancing dataset used in this analysis, we proceeded as follows:

[14] Downloaded from https://www.census.gov/geo/maps-data/data/cbf/cbf_counties.html on July 24, 2018.

1. We use the Daily Social Distancing SafeGraph data with observations at the census block group-day level for February 1 through March 30. We drop duplicate observations and exclude Alaska. We restrict our sample to census block groups with active devices throughout the entire time period. We also drop one census block group with anomalous behavior as notified by SafeGraph (FIPS: 190570010001).

2. We then aggregate to the county level. For the 'device count' and 'completely home device count' variables, we take the sum. For the 'median home dwell time' variable we take the mean weighted by 'the device count' in the census block group.

3. We then follow steps (4) and (5) described in Section A.1.1.

4. Lastly, we merge on 2016 Presidential vote shares, only keeping observations with valid vote shares.

Appendix Figure A1: POI Visits in 2019

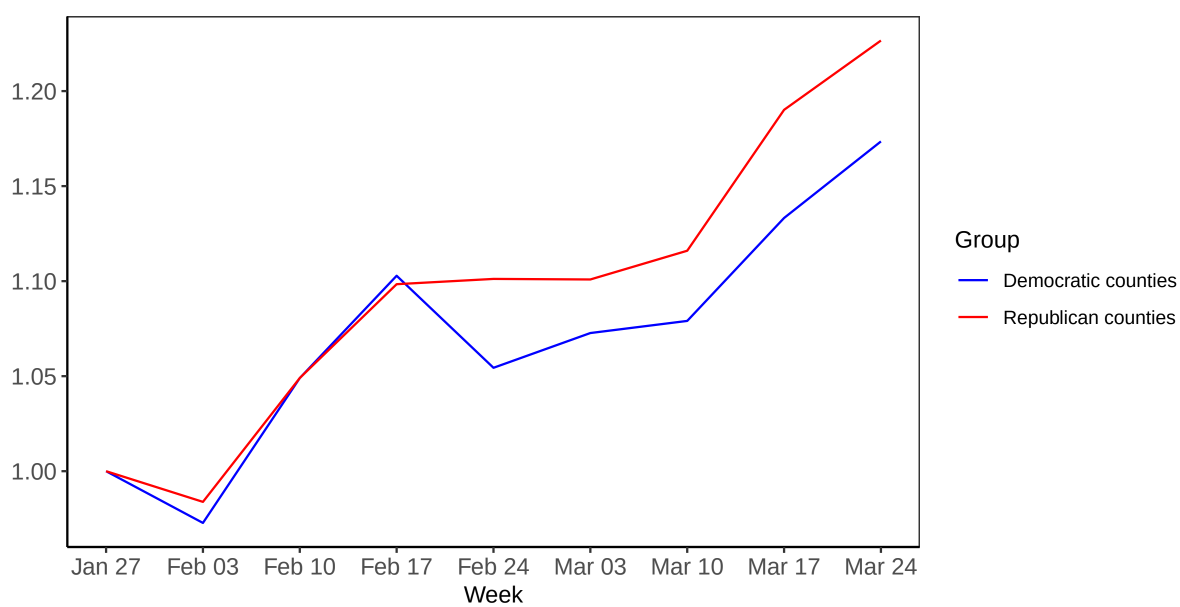

Note: Figure shows the aggregate number of POI visits (normalized to one) for nine weeks starting on January 27, 2019 for Republican counties and Democratic counties. Republican counties are defined to be those whose 2016 Republican vote share is greater than the median vote share across the counties in our sample.

Appendix Figure A2: Partisan Differences in Social Distancing, 2019

Panel A: Only County & Time FE

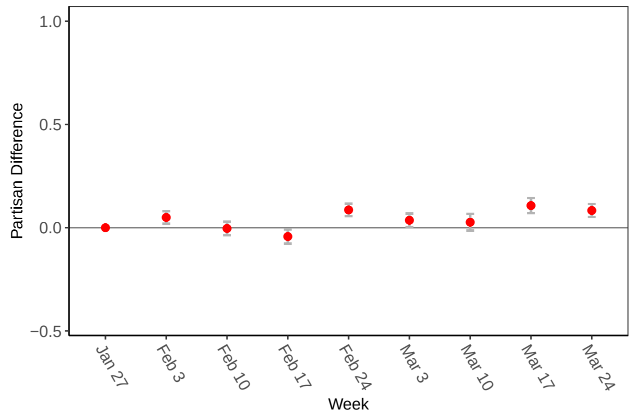

Panel B: Adds State-Time FE

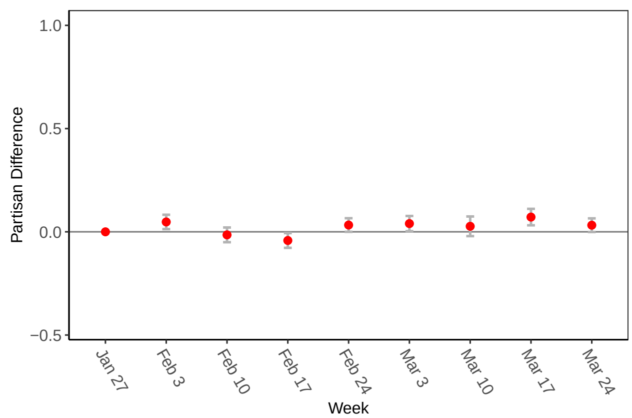

Panel C: Adds Health + Econ Controls

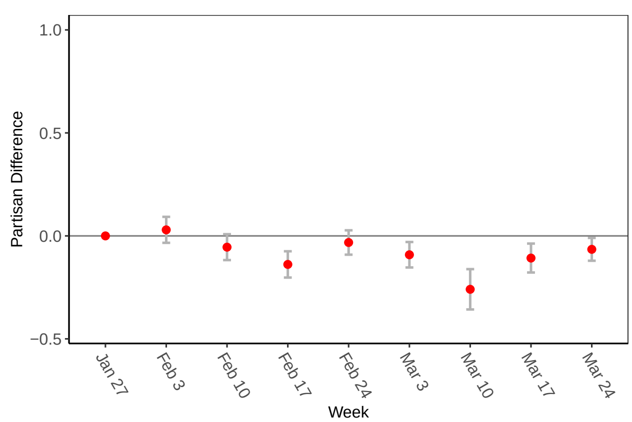

Note: Figure shows the estimated coefficients for county partisanship ρ_i on the log number of POI visits in the county as in Figure 5, except that eight weeks of data from January 27, 2019 are used instead of January 26, 2020. For Panel A, only county and time fixed effects are included as controls. Panel B is the same as Panel A except state-time fixed effects replace the time fixed effects. Panel C is the same as Panel B except the health and economic covariates are included. The grey error bars indicate 95 percent confidence intervals constructed using standard errors clustered at the county-level.

Appendix Figure A3: Partisan Differences in Social Distancing, Robustness

Panel A: Dropping Controls

Drops COVID-19 Controls | Drops Economic Controls | Drops All Health Controls

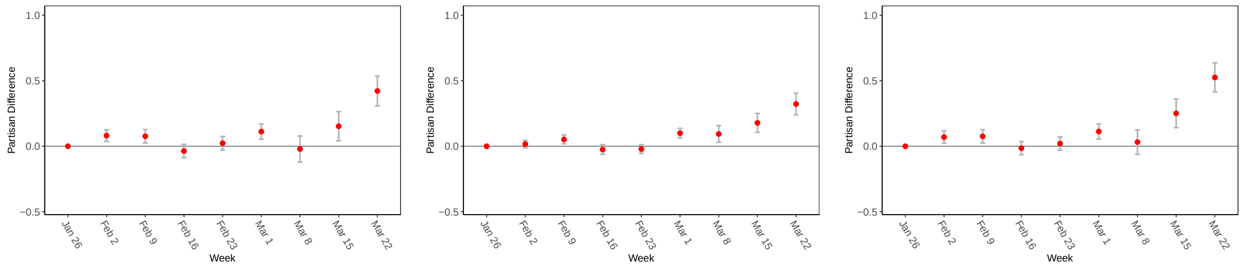

Panel B: Additional Specifications

Linear Controls | Adds Hispanic and Income | Drops State-Time FE

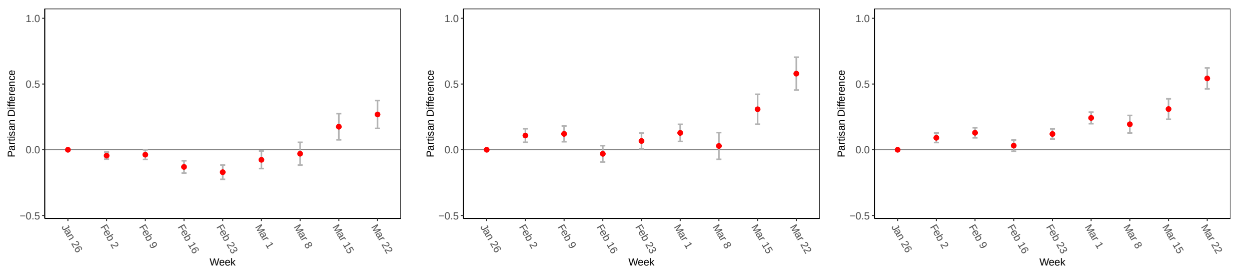

Panel C: Partisanship Indicators

Above or Below Median | Top or Bottom Quartile | Top or Bottom Decile

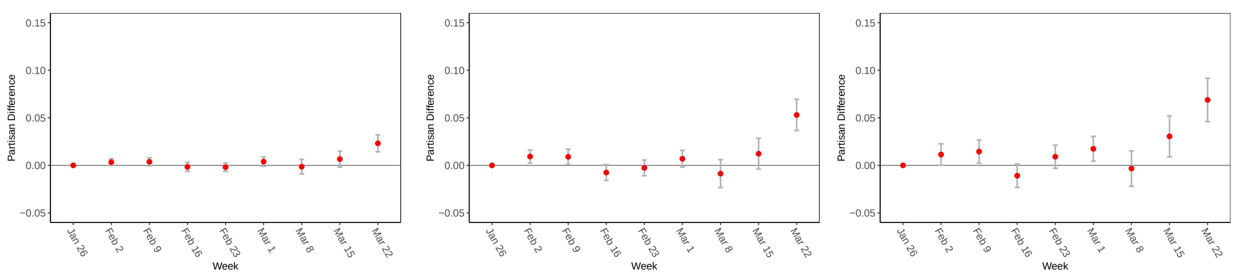

Note: Figure shows the estimated coefficients for county partisanship ρ_i on the log number of POI visits in the county. The specifications are analogous to our baseline in Panel C of Figure 5 except with the following deviations.

- Panel A: The first plot drops the COVID-19 cases and deaths controls; the second plot drops the economic controls; and the third plot drops all of the health controls, including the COVID-19 ones.
- Panel B: The first plot does not allow the coefficients on the controls to vary over time and interacts time-invariant controls with a linear time trend; the second plot adds the share Hispanic and the share with income less than 60k with time-varying coefficients; and the third plot drops the state-time fixed effects.
- Panel C: The first plot defines partisanship ρ_i to be 1 if Trump's vote share is greater than the median and -1 otherwise; the second plot defines partisanship ρ_i to be 1 if Trump's vote share is in the top quartile, -1 if in the bottom quartile, and 0 otherwise; and the third plot defines partisanship ρ_i to be 1 if Trump's vote share is in the top decile, -1 if in the bottom decile, and 0 otherwise.

Appendix Figure A3: Partisan Differences in Social Distancing, Robustness cont.

Panel D: Sample Restrictions and First Differences

Population below 500,000 | Drop CA, WA, and NY | First Differences

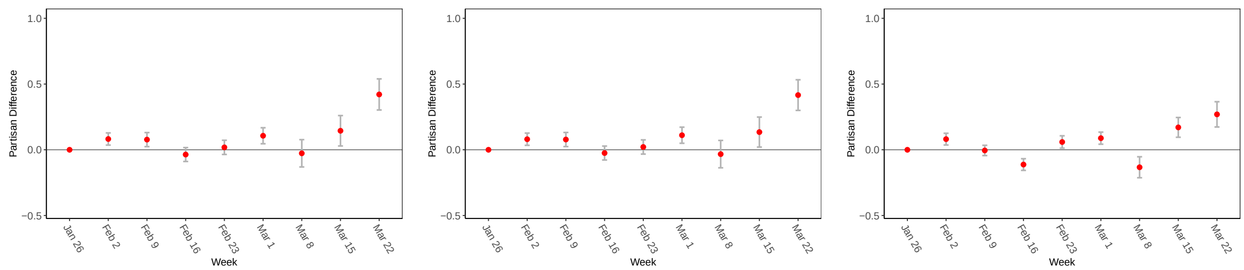

Panel E: Sample Restrictions by Vote Shares

Above Median Vote Share | Below Median Vote Share | Drop Bottom and Top Deciles

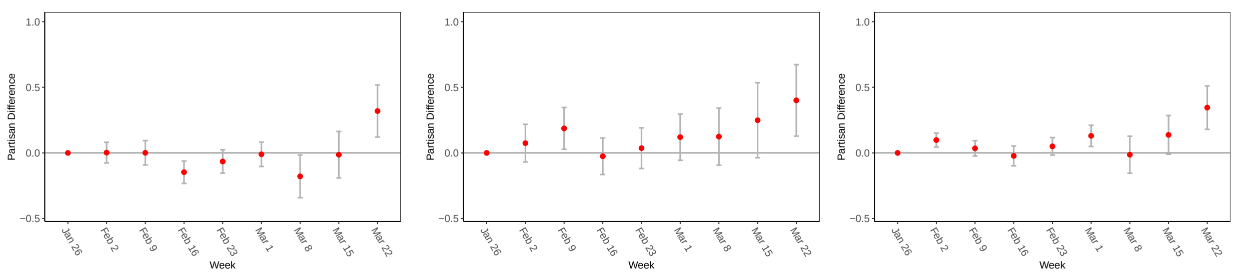

Panel F: Weighting, State Clustering, and Alternative Start Date

Weight by Population | Cluster SEs by State | Drop Week of January 26

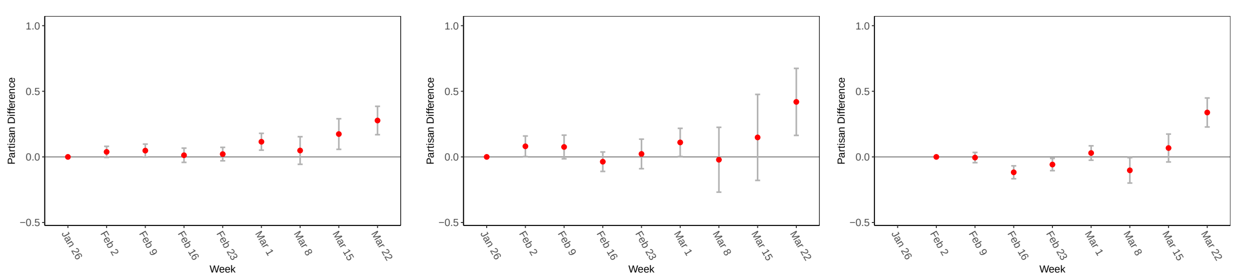

Note: Figure shows the estimated coefficients for county partisanship ρ_i on the log number of POI visits in the county. The specifications are analogous to our baseline in Panel C of Figure 5 except with the following deviations.

- Panel D: The first plot only keeps counties with a population below 500,000; the second plot drops California, Washington, and New York; and the third plot shows the estimated coefficients for county partisanship ρ_i on the *change* in the log number of POI visits in the county while dropping county fixed effects.

- Panel E: The first plot keeps counties for which Trump's vote share is greater than the median; the second plot keeps counties for which Trump's vote share is less than or equal to the median; and the third plot drops counties for which Trump's vote share was in the bottom or top decile.

- Panel F: The first plot weights observations by the county's population. The second plot clusters standard errors at the state-level. The third plot drops the week of January 26 and normalizes the estimates relative to the week of February 2.

Appendix Figure A3: Partisan Differences in Social Distancing, Robustness cont.

Panel G: Alternative Measures

Share Devices Leaving Home, No Correction

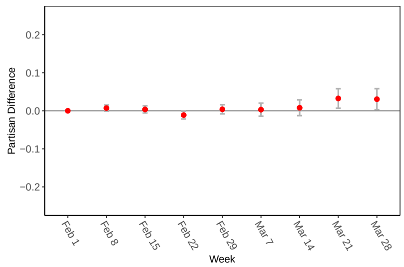

Note: Figure shows the estimated coefficients for county partisanship ρ_i on the log number of POI visits in the county. The specifications are analogous to our baseline in Panel C of Figure 5 except with the following deviations.

- Panel G: The first plot is analogous to 'Share Devices Leaving Home' in Figure 8, except that it does not account for differential sample attrition. Specifically, the outcome is defined to be number of devices observed leaving home divided by the share of devices in the panel for the same period.