Interdependence and the cost of uncoordinated responses to COVID-19

David Holtz[a,b,1], Michael Zhao[a,b,1], Seth G. Benzell[b,c], Cathy Y. Cao[a,b], Mohammad Amin Rahimian[a,d], Jeremy Yang[a,b], Jennifer Allen[a], Avinash Collis[b,e], Alex Moehring[a], Tara Sowrirajan[f,g], Dipayan Ghosh[a], Yunhao Zhang[a], Paramveer S. Dhillon[b,h], Christos Nicolaides[a,b,i], Dean Eckles[a,b,2], and Sinan Aral[a,b,2]

[a]Sloan School of Management, Massachusetts Institute of Technology, Cambridge, MA 02142; [b]Initiative on the Digital Economy, Massachusetts Institute of Technology, Cambridge, MA 02142; [c]Argyros School of Business and Economics, Chapman University, Orange, CA 92866; [d]Department of Industrial Engineering, University of Pittsburgh, Pittsburgh, PA 15261; [e]McCombs School of Business, The University of Texas at Austin, Austin, TX 78712; [f]Computer Science, Harvard University, Cambridge, MA 02138; [g]Media Lab, Massachusetts Institute of Technology, Cambridge, MA 02139; [h]School of Information, University of Michigan, Ann Arbor, MI 48109; and [i]School of Economics and Management, University of Cyprus, 2109 Aglantzia, Nicosia, Cyprus

Edited by Douglas S. Massey, Princeton University, Princeton, NJ, and approved July 10, 2020 (received for review May 12, 2020)

Social distancing is the core policy response to coronavirus disease 2019 (COVID-19). But, as federal, state and local governments begin opening businesses and relaxing shelter-in-place orders worldwide, we lack quantitative evidence on how policies in one region affect mobility and social distancing in other regions and the consequences of uncoordinated regional policies adopted in the presence of such spillovers. To investigate this concern, we combined daily, county-level data on shelter-in-place policies with movement data from over 27 million mobile devices, social network connections among over 220 million Facebook users, daily temperature and precipitation data from 62,000 weather stations, and county-level census data on population demographics to estimate the geographic and social network spillovers created by regional policies across the United States. Our analysis shows that the contact patterns of people in a given region are significantly influenced by the policies and behaviors of people in other, sometimes distant, regions. When just one-third of a state's social and geographic peer states adopt shelter-in-place policies, it creates a reduction in mobility equal to the state's own policy decisions. These spillovers are mediated by peer travel and distancing behaviors in those states. A simple analytical model calibrated with our empirical estimates demonstrated that the "loss from anarchy" in uncoordinated state policies is increasing in the number of noncooperating states and the size of social and geographic spillovers. These results suggest a substantial cost of uncoordinated government responses to COVID-19 when people, ideas, and media move across borders.

COVID-19 | peer effects | social spillovers | geographic spillovers

Pandemics are interdependent phenomena. Viruses and people's adherence to the government policies designed to contain them spill over from region to region. Early on, coronavirus disease 2019 (COVID-19) spread through international and domestic travel (1, 2). It is less well known, however, how behavioral responses to the pandemic and to government mitigation policies spill over from region to region due to geographic movement or social influence. As different regions begin to adopt heterogeneous reopening policies—with some opening businesses and relaxing shelter-in-place orders and others remaining closed and maintaining those orders—it is critical to understand how regional policies affect one another and the cost of adopting uncoordinated policies across regions.

Governments have enacted a variety of nonpharmaceutical interventions to reduce the spread of severe acute respiratory syndrome coronavirus 2, including social distancing policies designed to reduce high-density interactions among people in a particular region. Analyses of historical disease spread (3) and COVID-19 (4) indicate adherence to social distancing is crucial to slowing the spread of the pandemic, especially in the absence of a vaccine. But, while social distancing policies have, by and large, been left to individual cities, counties, states, and nations, uncoordinated policy interventions neglect that many geographic borders are porous and that increased social interdependence through communication media could create behavioral social influence across even distant regions.

In cases where coordination has occurred (for instance, in the northeastern United States), it has often been at the level of the "megaregion" (5). While intuitive, these local coordination efforts neglect the possibility that peoples' behaviors are influenced not just by those in their local communities but also by those with whom they are geographically distant but socially connected through mobile phones, video conferencing, and social media. These social spillovers may be even more relevant to the spread of COVID-19 as shelter-in-place orders have increased our reliance on digital connections, creating record-breaking

Significance

As local governments relax shelter-in-place orders worldwide, policy makers lack evidence on how policies in one region affect mobility and social distancing in other regions and the consequences of uncoordinated regional policies adopted in the presence of such spillovers. Our analysis suggests the contact patterns of people in one region are significantly influenced by the policies and behaviors of people in other, sometimes distant, regions. When just one-third of a state's social and geographic peer states adopt shelter-in-place policies, it creates a reduction in mobility equal to the state's own policy decisions, highlighting the need for national coordination. The paper gives governors a roadmap for coordination in the absence of national leadership and applies globally to other regions lacking coordination.

Author contributions: P.S.D., C.N., D.E., and S.A. led, directed, and oversaw the project; D.H., M.Z., C.Y.C., J.Y., J.A., A.C., and A.M. analyzed and validated measures of human mobility data; D.H. constructed a measure of geographic connectedness; C.Y.C. and J.A. constructed measures of industry-specific mobility; M.Z. led construction and analysis of historical weather data; D.H. led the difference-in-difference analysis, with preliminary analysis by M.Z.; M.Z. led the instrumental variables analysis, with J.Y., J.A., A.C., A.M., T.S., and Y.Z. contributing to defining instruments and conducting analysis; D.H. and J.Y. led statistical inference; S.G.B. and M.A.R. led development of the analytical model; D.H., S.G.B., and M.A.R. calibrated the analytical model with the empirical results; D.H. and S.A. led the writing of the manuscript; and all authors contributed to designing the research and writing the manuscript and supplementary information.

Competing interest statement: D.H., J.A., D.G., and D.E. were previously employees, interns, or contractors of Facebook, and J.A. has and, until recently, D.E. had a significant financial interest in Facebook. While this paper was under editorial review, S.A. and D.E. received a grant from Facebook for other research.

This article is a PNAS Direct Submission.

[1]D.H. and M.Z. contributed equally to this work.

[2]To whom correspondence may be addressed. Email: sinan@mit.edu or eckles@mit.edu.

This article contains supporting information online at https://www.pnas.org/lookup/suppl/doi:10.1073/pnas.2009522117/-/DCSupplemental.

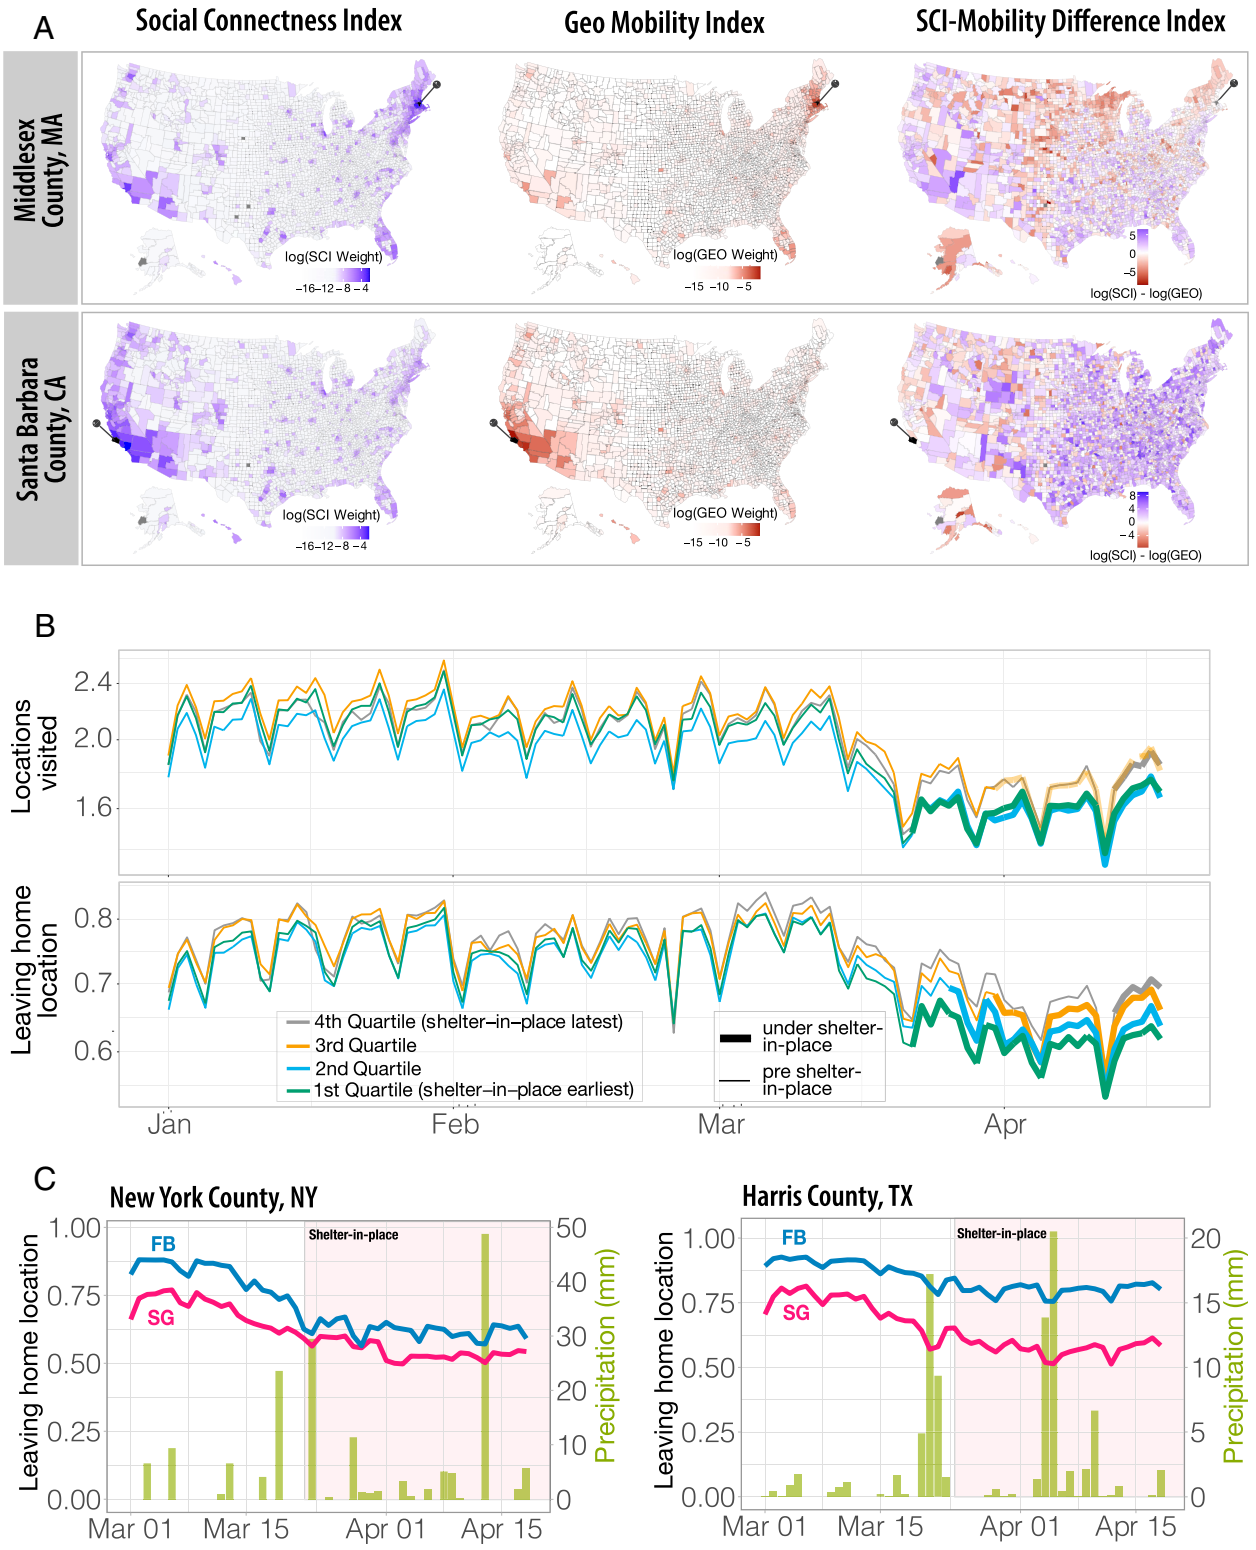

Fig. 1. (*A*) The social and geographic adjacency weights for two counties and the difference between them. For each county, geographic weights are generally stronger for nearby counties, whereas social weights are stronger for geographically distant counties. (*B*) The time series trends for the number of locations visited per device and the fraction of devices leaving home across county quartiles determined by the time at which each county introduced a shelter-in-place policy (if at all). Thicker lines correspond to periods of time where shelter-in-place was in effect. (*C*) The fraction of devices leaving home for two counties, along with the amount of precipitation in each county. Areas of the graph shaded in red correspond to periods during which shelter-in-place was in effect. In general, fewer devices leave home when it is raining, providing visual evidence of the strength of our weather instruments.

usage of social media and video conferencing to maintain our social ties across geographic distance (6).

Recent studies have used population-scale digital trace data (7) to measure the impact of social distancing policies on mobility, interaction intensity, and, in some cases, COVID-19 infections and their associated morbidity and mortality (8–11). These studies found adherence to social distancing policies is moderated by demographic attributes such as political affiliation (8, 10), age, gender, educational attainment (12), income, and access to high-speed internet (13). Unfortunately, our understanding of the impact of social distancing policies on mobility, infection rates, morbidity, and mortality is limited because existing research has not credibly accounted for social and geographic spillovers, which, if large, could substantially alter our perceptions of the effectiveness of local policies.

Researchers have causally identified the existence of social contagion in offline behaviors such as exercise (14), product adoption (15), and voting (16). Others have shown the potential for local policies to cause geographic spillovers to neighboring communities (17). Given these empirical regularities, it is likely that an individual's mobility and adherence to social distancing are impacted not only by policies in their own regions but also by the policies of neighboring regions and distant regions in which their social network connections reside. Put differently, a local government's social distancing policy may significantly impact the health outcomes of other communities, including those that are geographically proximate or those that are geographically distant but socially proximate. The existence of such spillovers could imply substantial health and economic consequences to adopting uncoordinated policies across socially and geographically connected regions.

Here, we measure mobility across borders, adherence to social distancing, and high-density interactions between people in physical space using population-scale digital trace mobility data from Safegraph (18) and Facebook (19). The Safegraph data record the location and movement of over 22 million mobile devices, including the fraction of mobile devices staying home each day in every US county, the average number of locations visited by mobile devices each day in every US county, and the number of visits to distinct points of interest each day in every county. The Facebook data, which cover over 27 million mobile devices, also record the fraction of mobile devices staying home each day in every US county and the average number of locations mobile devices visit each day in each US county.

We augmented these mobility datasets with an index of the degree to which different US counties are socially connected on Facebook (20), temperature and precipitation data from the National Oceanic and Atmospheric Administration's global historical climatology network database (21), census counts of each US county's total population, and a detailed database of the timing of COVID-related government interventions in every US county (22). This combination of data allowed us to causally estimate the direct effect of government social distancing policies on local mobility, the indirect effects of other governments' social distancing policies on local mobility, and the mediation of these effects by social influence and geographic proximity across the entire United States. Fig. 1 highlights various attributes of these datasets that will be key to our analyses. We specifically focus on mobility outcomes, rather than health outcomes such as mortality rates and hospitalizations, as social distancing policies directly target mobility behavior and because there are known data quality issues for health-related outcomes. We focus our analysis on the 2,502 US counties appearing in both the Facebook and Safegraph data from March 1, 2020 to April 18, 2020, during which the vast majority of social distancing policies were implemented in the United States.

We first estimated a difference-in-differences (DiD) model that considered the direct county-level effect of social distancing policies, but did not account for geographic or social spillovers (*SI Appendix*, section S2). Consistent with previous studies (8, 23), we found that implementing a shelter-in-place policy led to a 3.2% ($P < 0.001$) decrease in the fraction of devices leaving their homes and a 6.0% decrease ($P < 0.001$) in the number of locations visited.* While this specification suggested that social distancing policies are effective at curbing mobility when enacted by focal states, it fails to account for geographic spillovers and social spillovers, and may therefore overstate the effectiveness of any one county's or state's policy.

We therefore estimated DiD models that account for geographic spillovers, according to a geographic adjacency matrix, and social spillovers, according to a social adjacency matrix. We constructed the geographic adjacency matrix using the Safegraph data from January and February 2020 to calculate the fraction of Census block group visits to county i from people living in county j (*SI Appendix*, section S1). We constructed the social adjacency matrix by combining Facebook's Social Connectedness Index with Census data to calculate the fraction of county i's Facebook ties that are to friends in county j (*SI Appendix*, section S1).

We began by estimating a DiD model that quantifies geographic spillovers, but not social spillovers (*SI Appendix*, section S2). In some sense, this specification accounts for spillovers in the same way mayors and governors who are currently coordinating at the "megaregion" level may account for them—by considering geographically proximate peer regions. Results from this model, shown in Fig. 2*A*, suggest that, when accounting for geographic spillovers, a focal county implementing a shelter-in-place order reduced the average number of locations visited in that county by 4.0% ($P < 0.001$). But, when half of the county's geographic alters also implemented shelter-in-place orders, it further reduced the average number of locations visited in the focal county by 2.3% ($P < 0.001$). A focal county implementing a shelter-in-place order reduced the fraction of devices leaving home in that county by 2.0% ($P < 0.001$). But, when half of the county's geographic alters also implemented shelter-in-place orders, it further reduced the fraction of devices leaving home in the focal county by 1.4% ($P < 0.001$).

When we estimated the impact of these geographic spillovers on mobility in a dyadic DiD model (*SI Appendix*, section S2), we found that, when only one county in a dyad implemented a shelter-in-place policy, travel from that county to the nonimplementing county increased by 0.55% ($P = 0.05$) on average, while travel from the county not implementing the policy to the county implementing the policy decreased by 1.2% ($P < 0.01$). When both counties implemented shelter-in-place orders, travel between the counties decreased by 0.51% ($P < 0.001$). The results in Fig. 2 *A* and *B* validate the importance of coordinating geographically connected regions to, for example, reduce travel across borders from counties in which businesses are closed to neighboring counties in which businesses are open. But, they fail to account for social spillover effects.

When we estimated a DiD model that distinguished between within-state and across-state alters, we found that, when considering both social and geographic spillovers, the estimated spillover effect from 100% of alter states implementing a shelter-in-place policy was a 13% reduction ($P < 0.001$) in locations visited and a 9.1% reduction ($P < 0.001$) in the fraction of devices leaving home (Fig. 2*C*) (*SI Appendix*, section S2). Under this model, policy spillover estimates, when accounting for social spillovers, are over 2 times larger than when only considering

*Although we estimate similar effects for both Safegraph and Facebook outcomes, we refer explicitly to the Safegraph estimates in the text. However, both Safegraph and Facebook results are shown in Figs. 1–3*A*, and throughout the *SI Appendix*.

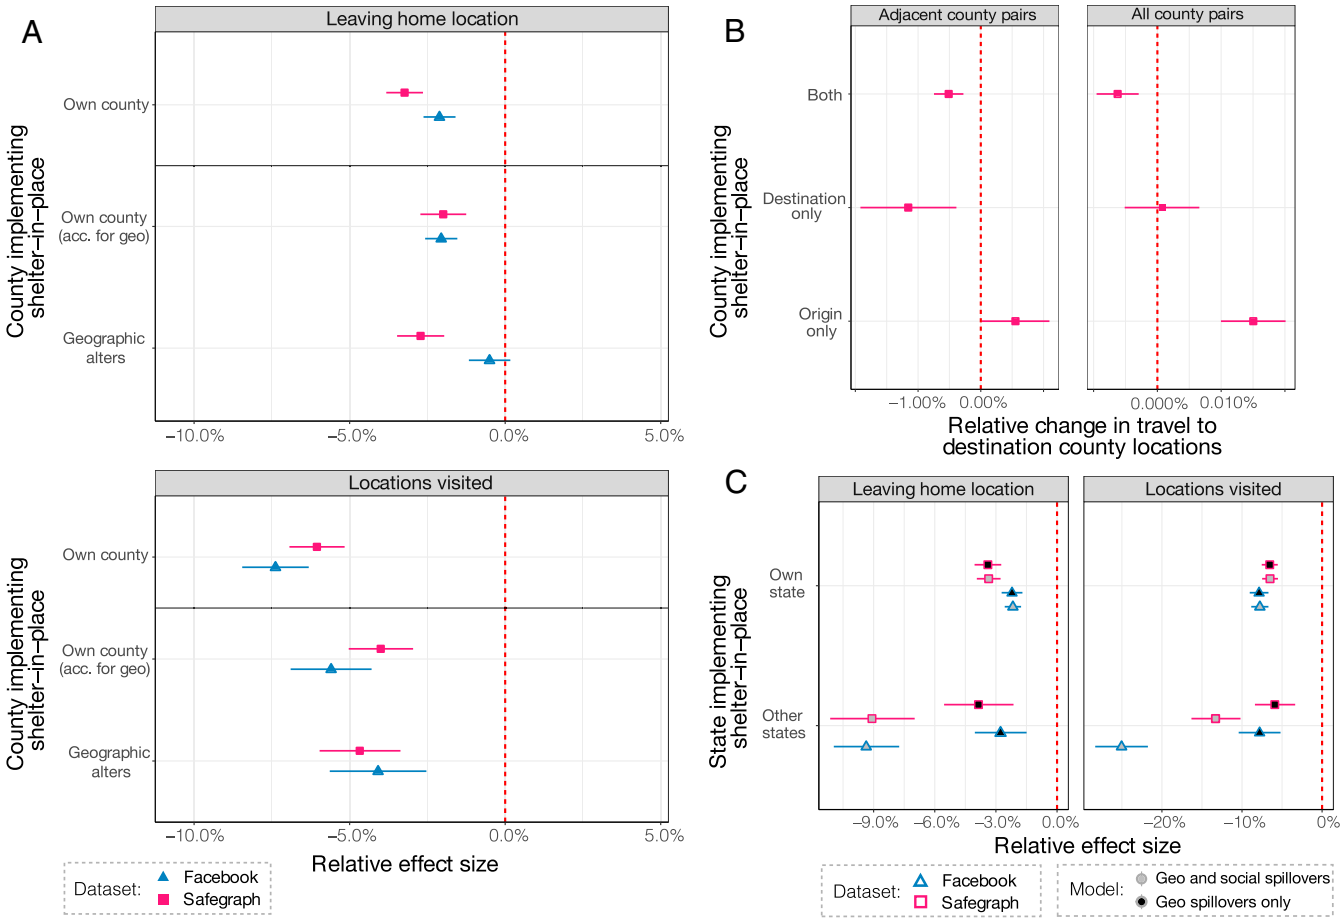

Fig. 2. (*A*) Comparison of the results of our DiD model that ignores spillovers and estimates the effect of a policy on its "own county" and our DiD model that includes geographic spillovers and separates the effects of a policy on its own county ("own county [acc. for geo]") from the effects of the policies of geographically connected counties ("Geographic alters"). For both the fraction of devices leaving home and the number of locations visited, the geographic spillovers are approximately equal in magnitude to the direct effects of ego county shelter-in-place policies. (*B*) The results of a county-level dyadic DiD model using either all county pairs or only adjacent county pairs. When only an "origin" county implements a shelter-in-place policy, outbound travel to the destination county increases. Only when either the destination county or both counties implement shelter-in-place does travel to the destination county decrease. (*C*) Comparison of our estimates of the direct effect of shelter-in-place, as well as the spillover effects of other states' shelter-in-place policies, with and without accounting for social spillovers. When we account for social spillovers, the magnitude of our spillover estimates increases by over a factor of 2.

geographic spillovers. In other words, it is not only the policy decisions of geographically proximate states that affect outcomes in a focal state but also the communities to which that state is socially connected through communication technology. Results from this model also suggest that 36% of a state's geographic and social peer states implementing shelter-in-place policies is as effective at reducing mobility as the focal state implementing its own shelter-in-place policy.

Our analyses thus far establish the importance of two types of connections that contribute to spillovers: geographic proximity and social influence. However, although our DiD estimates show that social spillovers are an important determinant of a focal county's mobility levels, these estimates do not establish the underlying mechanisms that drive these effects. It is unclear whether changes in focal county mobility levels are driven by knowledge of peers' counties' policies, changes in the behavior of socially connected peers, or another mechanism. To identify the extent to which this effect is driven by changes in peer behaviors, we employed our instrumental variables (IV) estimation framework and, while controlling for peers' shelter-in-place policies, instrumented for the behavior of peers in socially connected counties using weather, shifts in industry visit shares, and their interaction with peers' shelter-in-place policies (*SI Appendix*, section S3). We estimated that a 3.0% reduction in the number of peer locations visited leads to a 5.6% reduction in the number of locations visited in a focal county ($P < 0.001$) and that a 1.5% reduction in the number of peers leaving their home location leads to a 2.4% reduction in the number of focal county devices leaving their home location ($P < 0.001$) (Fig. 3*A*). These effect sizes suggest that social spillovers are substantially mediated by peer behavior. In other words, people in a focal state are significantly influenced by the behavior of their peers in other states when calibrating their own social distancing behaviors and choices.

We also combined our social and geographic adjacency matrices with the point estimates obtained from our DiD with spillovers model to estimate the strength of interdependence between each pair of US states to understand, for example, how much mobility would go down in state j if state i implemented a shelter-in-place policy (*SI Appendix*, section S5). Fig. 3*C* shows the ego networks for eight US states chosen from across the country (we report results for all 50 states and Washington, D.C. in *SI Appendix*). Generally speaking, each state's mobility outcomes are impacted by the policy decisions of not just geographically proximate states but also socially connected, distant states. For instance, Florida's mobility is most affected by

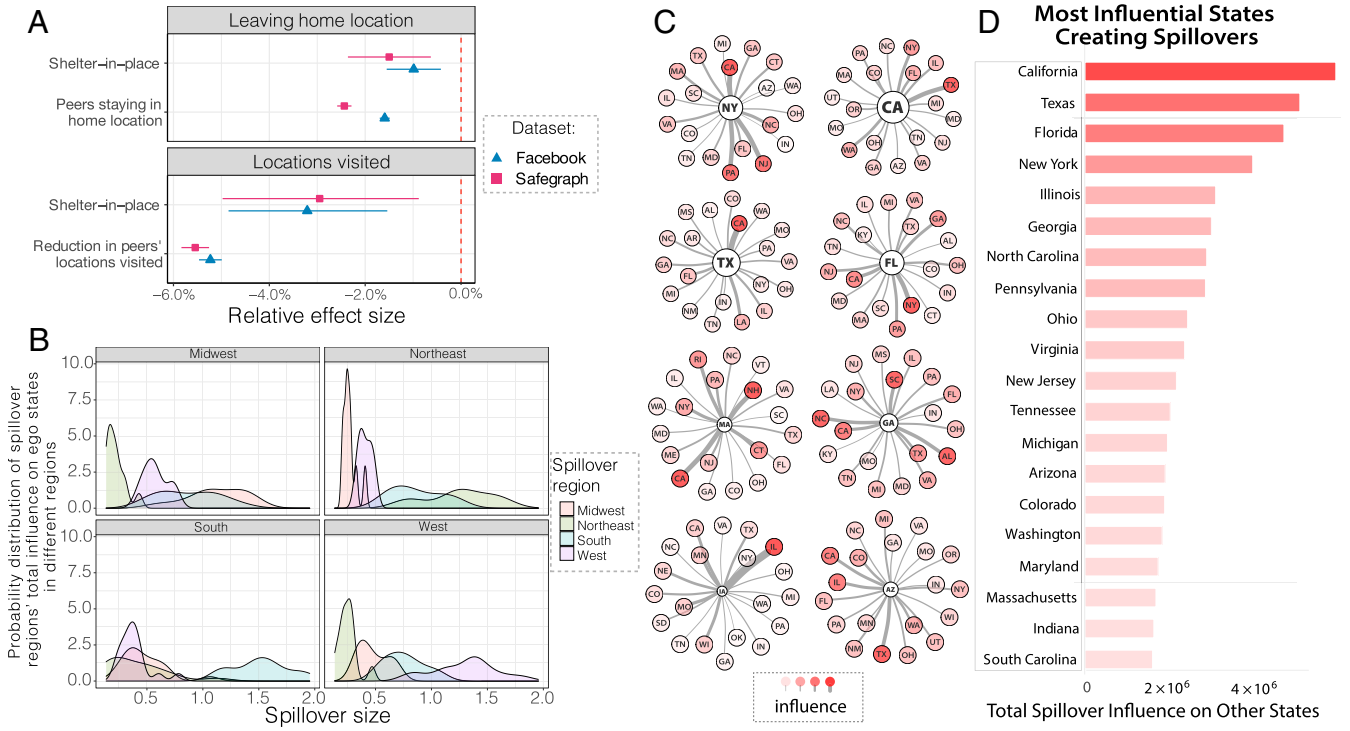

Fig. 3. (*A*) The causal effect of endogenous mobility levels in Facebook alter counties on mobility levels in focal counties estimated using an IV framework. The magnitudes of endogenous peer effects are scaled to the direct effect of shelter-in-place orders in their own state. (*B*) The region-level probability distribution functions for the size of the total spillover effect from alter states in each US region, relative to the direct effect of each focal state's own shelter-in-place policy. (*C*) The ego networks for eight different US states. For each state, we display the 20 alter states whose own shelter-in-place policy causes the largest reductions in mobility levels in the ego state, according to our DiD model. Both the edge weight and the alter node color correspond to the amount of influence the alter exerts on the ego. (*D*) The 20 states whose shelter-in-place policies cause the greatest reduction in devices leaving home across the United States, according to our DiD model.

New York implementing shelter-in-place, presumably through digitally mediated social influence or travel, despite the states being distant. New Hampshire has a strong influence on adjacent Massachusetts, despite being a small state. These interdependence estimates can also be combined with state population levels to estimate the US states whose shelter-in-place policies would lead to the greatest reductions in mobility across the rest of the United States (Fig. 3*C*). A state's total spillover influence is highly correlated, but not equivalent to, that state's population size. This highlights the need for states across the country to coordinate, even if they are not near one another, and our results suggest which states should be coordinating with which other states based on the strength of spillovers between them.

Finally, we used our empirical estimates to calibrate a simple game-theoretic model of the inefficiency created by states failing to coordinate over social and geographic spillovers (*SI Appendix*, section S6). In the model, each states' social distancing outcomes depend on their own (linearly costly) mobility policies as well the policy of other states. We further assume that each state has a specific (exogenous) "target" mobility they aim to achieve. When states are uncoordinated, they play a one-shot game without transfers where each state chooses its own level of mobility policy restrictiveness by balancing the direct policy cost and a quadratic loss function for missing their own mobility target. We compare the aggregate utility achieved under the Nash equilibrium of this game to the aggregate utility achieved under optimal coordination by a social planner for varying levels of spillover intensity. The difference between the Nash equilibrium outcome and the socially optimal outcome characterizes losses from uncoordinated policies, while the choices states make in equilibrium characterize the free-riding and compensation for other states' negligence that take place in the absence of coordination.

When spillovers or the cost of implementing policies are low, welfare under coordination through a social planner is not much higher than in a Nash equilibrium. But, when spillovers and costs are high, the lack of coordination can be quite costly. Utility can be up to 69% lower when states fail to coordinate in the presence of spillovers as large as those we detect in our empirical analyses (Fig. 4*A*). Furthermore, when spillovers are high, states' policies diverge, as one state needs to compensate for the neglect of another state's loose restrictions by imposing even stricter, more costly policies than necessary to achieve their desired mobility target. When states coordinate, however, social and geographic spillovers actually help them achieve their targets more efficiently because they essentially provide "free treatments" as the cooperative behaviors of peer states positively influence social distancing behaviors in focal states.

This work is not without limitations. First, while our estimates of peer influence utilize weather and shift share instruments, our estimates of the interdependence between individual US states rely on our DiD analysis, which may miss some state- or dyad-level heterogeneity. Furthermore, although our analysis of lags and leads suggests the robustness of our analysis (presented in *SI Appendix*), our DiD estimates may not capture all anticipatory behavior (e.g., people stocking up on groceries before government policies take effect). In *SI Appendix*, we more carefully examine the robustness of our estimates to these and other challenges, but this work nonetheless relies on assumptions that are not fully testable. Additionally, our DiD analyses do not control for factors such as COVID-19–related hospitalizations and/or

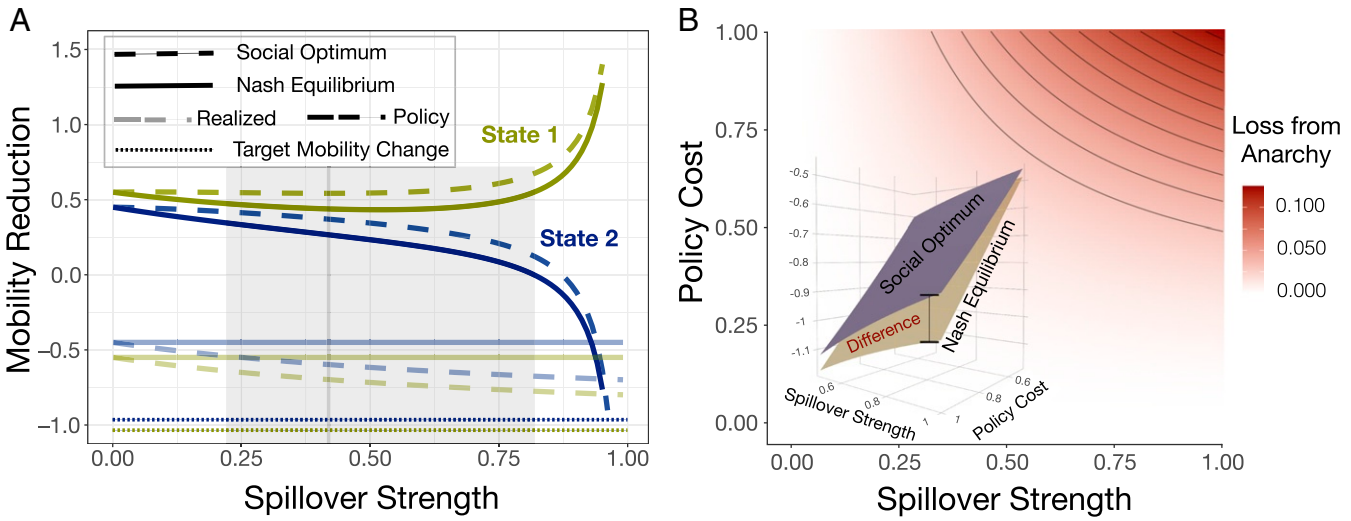

Fig. 4. Mobility reduction targets, optimal policy choices, equilibrium mobility reductions, and utility under anarchy (Nash Equilibrium) and coordination (Social Optimum). (*A*) Mobility outcomes achieved (faint) under varying levels of spillover strength (*x* axis) for a pair of states with similar but not identical reduction targets (dash-dot-dotted lines) and their resulting policy choices (dark lines). The gray shading and gray vertical line correspond to the minimum, median, and maximum spillover strengths we observe in our DiD estimates of between state spillovers (the model producing the gray shaded estimates in Fig. 2*C*). (*B*) When spillover strength is low, Nash equilibrium policies for both states are similar, and there is no loss from anarchy compared to coordination. As spillovers get stronger, policies for two states diverge. Under coordination, this divergence decreases equilibrium mobility toward the target, but, under anarchy, the states' actions wastefully offset, leaving outcome mobility unchanged as spillover strength increases. (*Inset*) Utility under both equilibria. Maximum utility is increasing in spillovers because they, in a sense, create "free" reductions in mobility. The loss from anarchy is increasing in both the size of spillovers and the cost of mobility reductions. Spillovers are assumed to be symmetrical. In *A*, the cost of implementing policies is set at 1.

deaths, which may affect the timing with which social distancing policies were imposed. However, it is not clear how to appropriately control for these factors, since the "correct" control is likely policy makers' perceptions of COVID-19–related health outcomes at the time of decision-making, as opposed to the actual number of hospitalizations or deaths. Finally, when treatment is staggered, the DiD estimand is a weighted average of every possible pairwise treatment effect. The weighting function used to construct this estimand has recently been characterized for the case of a binary treatment variable (24), but is still not well understood when the treatment variable is continuous and its effects are dynamic.

As government officials around the world begin to calculate the costs and benefits associated with lifting social distancing policies, it is crucial we accurately estimate these policies' effects. Our findings indicate that any given government's decision to lift a social distancing policy will likely affect the behavioral and health outcomes of not only their own citizens but also the citizens of geographically and socially proximate communities. These results suggest there are significant negative welfare repercussions from uncoordinated government social distancing policies, which suffer from a coordination problem resembling the price of anarchy (25). This implies that it is important for federal governing bodies (e.g., the United States federal government and the European Union) to coordinate policy action, even in cases where final policy decisions are in the hands of local governments. In the absence of coordination by federal governing bodies, we recommend that individual countries, states, and counties coordinate with the countries, states, and counties to which they are the most strongly geographically and socially connected. In the United States, our estimates provide governors with direct guidance on which other states are influencing their states the most (see *SI Appendix*, section S5 for coordination maps for all 50 states and Washington, D.C.). These coordination maps could also be created for localities around the world, using our methods. As states begin to reopen, we recommend that governors use these maps to establish direct coordination between influencing states, to keep each other abreast of changing policies, to model the effects of other states' actions on outcomes in their own states, and to coordinate regional and superregional policies to maximize the effectiveness of local policies. Our model suggests that spillovers can benefit states in the presence of coordination. We therefore hope our work inspires a greater level of such coordination between local government officials when determining policies related to social distancing and future research into the indirect effects of these policies.

Materials and Methods

We first estimated the causal effects of county-level shelter-in-place orders on their own county's population mobility, measured by the fraction of mobile devices leaving home and the mean number of locations visited per device, as well as their effects on mobility in counties to which they are geographically connected through physical proximity or socially connected through social media on Facebook, using the following DiD model specification:

$$Y_{it} = \delta_1 D_{it} + \delta_2 D^{geo}_{-it} + \delta_3 D^{social}_{-it} + f(W_{it}) + \alpha_i + \tau_t + \epsilon_{it}, \quad [1]$$

where Y_{it} denotes the social distancing outcome, D_{it} indicates whether shelter-in-place has been enacted in county i in time period t, D^{geo}_{-it} is the geographic adjacency weighted average of peer county policies, D^{social}_{-it} is the social adjacency weighted average of peer county policies, and $f(W_{it})$ is a term that flexibly controls for the potential nonlinear impact of weather using a "double machine learning" approach (26). α_i and τ_t represent a set of county and time fixed effects, and ϵ_{it} denotes the error term. Our statistical inference allows for correlations between counties that are socially or geographically connected or located in the same US states using adjacency- and cluster-robust standard errors (27). Although not explicitly indicated in this notation, we estimate DiD models that treat all alter counties equivalently, and also DiD models that distinguish between same-state counties and different-state counties. Here we report results for shelter-in-place policies, which typically supersede business closures. We report results for both shelter-in-place policies and business closures in *SI Appendix*.

While the DiD analysis allowed us to measure the effect of connected counties' policies on focal counties' population mobility, the effect of connected counties' policies could be driven by awareness of the policies of nearby US counties and states, changes in friends' behavior, or the amount

of intercounty and interstate travel between regions. We therefore used IV analysis to estimate the mechanisms driving geographic and social spillovers by separately measuring the effects of connected county policies and the effects of peer behavior within connected counties on mobility in focal counties. Our IV analysis uses exogenous variation in weather (14, 28, 29) and the extent to which different counties are exposed to national changes in industry visit behavior based on prepandemic data (30, 31) as exogenous shocks to peer behaviors in connected counties to identify their causal influence on mobility behavior in focal counties.

We estimate the following main and first-stage model specifications:

$$Y_{it} = \beta Y_{-it} + \delta_1 D_{it} + \delta_2 D^{geo}_{-it} + \delta_3 D^{social}_{-it} + \psi S_{it} + f(W_{it}) + \alpha_i + \tau_t + \epsilon_{it} \quad [2]$$

$$Y_{-it} = \gamma_1 D_{it} + \gamma_2 D^{geo}_{-it} + \gamma_3 D^{social}_{-it} + \pi S_{it} + g(W_{it}) + h(D^{social}_{-it}, S_{-it}, W_{-it}) + \alpha_{-i} + \tau_t + \nu_{-it}, \quad [3]$$

where Y_{it}, D_{it}, D^{geo}_{-it}, D^{social}_{-it}, $f(W_{it})$, α_i, τ_t, and ϵ_{it} are as they were in Eq. **1**. Y_{-it} denotes the social adjacency weighted average mobility behaviors of individuals in other counties, and the main parameter of interest β represents the endogenous peer effect of that behavior. S_{it} is the set of industry shift-shares for county i. In the first stage, $g(\cdot)$ is also a function that captures the nonlinear effects of W_{it}. D_{-it} and the social adjacency weighted averages of alter counties shift-shares and weather, S_{-it} and W_{-it}, and their interactions, form the set of candidate instruments. The associated function $h(\cdot)$ is a post-least absolute shrinkage and selection operator (post-LASSO) (32) that selects a smaller set of instruments. Lastly, ν_{-it} denotes the first-stage error term. We report adjacency- and cluster-robust standard errors. Further details are provided in *SI Appendix*. This research was reviewed and classified as exempt by the Massachusetts Institute of Technology (MIT) Committee on the Use of Humans as Experimental Subjects (i.e., MIT's Institutional Review Board), because the research was secondary use research involving the use of de-identified, aggregate data.

Data Availability. Code and information regarding data access are available on GitHub (https://github.com/mfzhao/covid_interdependence).

ACKNOWLEDGMENTS. We are grateful to the MIT Initiative on the Digital Economy and the MIT Social Analytics Lab for their support. We thank Safegraph and Facebook for providing data. We thank E. Bakshy, B. J. Fogg, A. Fradkin, A. Ghose, M. G. Hudgens, M. Kaptein, D. Rand, A. Simonov, S. J. Taylor, C. Tucker, A. Volfovsky, and J. Ugander for helpful comments.

1. J. S. Jia *et al.*, Population flow drives spatio-temporal distribution of COVID-19 in China. *Nature* **582**, 389–394 (2020).
2. M. Chinazzi *et al.*, The effect of travel restrictions on the spread of the 2019 novel coronavirus (COVID-19) outbreak. *Science* **368**, 395–400 (2020).
3. J. Adda, Economic activity and the spread of viral diseases: Evidence from high frequency data. *Q. J. Econ.* **131**, 891–941 (2016).
4. M. U. Kraemer *et al.*, The effect of human mobility and control measures on the COVID-19 epidemic in China. *Science* **368**, 493–497 (2020).
5. R. Florida, T. Gulden, C. Mellander, The rise of the mega-region. *Camb. J. Reg. Econ. Soc.* **1**, 459–476 (2008).
6. A. Schultz, J. Parikh, *Keeping our services stable and reliable during the COVID-19 outbreak*. *Facebook* (24 March 2020) https://about.fb.com/news/2020/03/keeping-our-apps-stable-during-covid-19/. Accessed 7 May 2020.
7. C. O. Buckee *et al.*, Aggregated mobility data could help fight COVID-19. *Science* **368**, 145.2–146 (2020).
8. M. Painter, T. Qiu, Political beliefs affect compliance with COVID-19 social distancing orders. https://papers.ssrn.com/sol3/papers.cfm?abstract_id=3569098. Accessed 7 May 2020.
9. M. Bakker, A. Berkel, M. Grohl, S. Pentland, E. Moro, Effect of social distancing measures in the New York City metropolitan area. https://connection.mit.edu/sites/default/files/publication-pdfs/Effect%20of%20social%20distance%20measures%20in%20social%20distancing%20in%20the%20NY%20area.pdf. Accessed 7 May 2020.
10. H. Allcott *et al.*, "Polarization and public health: Partisan differences in social distancing during the coronavirus pandemic" (Tech. Rep. 26946, National Bureau of Economic Research, 2020).
11. N. Oliver *et al.*, Mobile phone data for informing public health actions across the COVID-19 pandemic life cycle. *Sci. Adv.* **6**, eabc0764 (2020).
12. A. L. Olsen, F. Hjorth, Willingness to distance in the COVID-19 pandemic. https://osf.io/xpwg2/. Accessed 7 May 2020.
13. L. Chiou, C. Tucker, "Social distancing, internet access and inequality" (Tech. Rep. 26982, National Bureau of Economic Research, 2020)
14. S. Aral, C. Nicolaides, Exercise contagion in a global social network. *Nat. Commun.* **8**, 1–8 (2017).
15. S. Aral, D. Walker, Identifying influential and susceptible members of social networks. *Science* **337**, 337–341 (2012).
16. R. M. Bond *et al.*, A 61-million-person experiment in social influence and political mobilization. *Nature* **489**, 295–298 (2012).
17. S. G. Bronars, J. R. Lott, Criminal deterrence, geographic spillovers, and the right to carry concealed handguns. *Am. Econ. Rev.* **88**, 475–479 (1998).
18. M. K. Chen, R. Rohla, The effect of partisanship and political advertising on close family ties. *Science* **360**, 1020–1024 (2018).
19. P. Maas *et al.*, "Facebook disaster maps: Aggregate insights for crisis response and recovery" in *Proceedings of the 16th International Conference on Information Systems for Crisis Response and Management*, Z. Franco, J. J. González, J. H. Canós, Eds. (Information Systems for Crisis Response and Management, Valencia, Spain, 2019), pp. 19–22.
20. M. Bailey, R. Cao, T. Kuchler, J. Stroebel, A. Wong, Social connectedness: Measurement, determinants, and effects. *J. Econ. Perspect.* **32**, 259–80 (2018).
21. M. J. Menne, I. Durre, R. S. Vose, B. E. Gleason, T. G. Houston, An overview of the Global Historical Climatology Network-Daily database. *J. Atmos. Ocean. Technol.* **29**, 897–910 (2012).
22. B. D. Killeen *et al.*, A county-level dataset for informing the United States' response to COVID-19. arXiv:2004.00756 (1 April 2020).
23. A. Brzezinski, G. Deiana, V. Kecht, D. Van Dijcke, "*The COVID-19 pandemic: Government vs. Community action across the United States*" (Tech. Rep. 2020-06, Institute for New Economic Thinking, 2020).
24. A. Goodman-Bacon, "*Difference-in-differences with variation in treatment timing*" (Report 25018, National Bureau of Economic Research, 2018).
25. E. Koutsoupias, C. Papadimitriou, "Worst-case equilibria" in *Annual Symposium on Theoretical Aspects of Computer Science*, C. Meinel, S. Tison, Eds. (Springer, 1999), pp. 404–413.
26. V. Chernozhukov *et al.*, Double/debiased machine learning for treatment and structural parameters. *Econom. J.* **21**, C1–C68 (2018).
27. T. G. Conley, GMM estimation with cross sectional dependence. *J. Econom.* **92**, 1–45 (1999).
28. L. Coviello *et al.*, Detecting emotional contagion in massive social networks. *PloS One* **9**, e90315 (2014).
29. S. Aral, M. Zhao, Social media sharing and online news consumption. https://papers.ssrn.com/sol3/papers.cfm?abstract_id=3328864. Accessed 7 May 2020.
30. T. J. Bartik, *Who Benefits from State and Local Economic Development Policies?* (W.E. Upjohn Institute for Employment Research, 1991).
31. P. Goldsmith-Pinkham, I. Sorkin, H. Swift, "*Bartik Instruments: What, when, why, and how*" (Report 24408, National Bureau of Economic Research, 2018).
32. A. Belloni, D. Chen, V. Chernozhukov, C. Hansen, Sparse models and methods for optimal instruments with an application to eminent domain. *Econometrica* **80**, 2369–2429 (2012).

Supplementary Information

Interdependence and the Cost of Uncoordinated Responses to COVID-19

www.pnas.org/cgi/doi/10.1073/pnas.2009522117

David Holtz,[1,2,†] Michael Zhao,[1,2,†] Seth G. Benzell,[3,2] Cathy Y. Cao,[1,2]
Mohammad Amin Rahimian,[4,1] Jeremy Yang,[1,2] Jennifer Allen,[1] Avinash Collis,[5,2]
Alex Moehring,[1] Tara Sowrirajan,[6,7] Dipayan Ghosh,[1] Yunhao Zhang,[1]
Paramveer S. Dhillon,[8,2] Christos Nicolaides,[9,1,2] Dean Eckles,[1,2*] Sinan Aral [1,2*]

[1]MIT Sloan School of Management
100 Main Street, Cambridge, MA 02142, USA
[2]MIT Initiative on the Digital Economy
100 Main Street, Cambridge, MA 02142, USA
[3]Chapman University, Argyros School of Business and Economics
1 University Drive, Orange, CA 92866, USA
[4]Department of Industrial Engineering, University of Pittsburgh
4200 Fifth Avenue, Pittsburgh, PA 15260, USA
[5]McCombs School of Business, The University of Texas at Austin
2110 Speedway, Austin, TX 78705, USA
[6]Harvard University
29 Oxford Street, Cambridge, MA 02138, USA
[7]MIT Media Lab
75 Amherst Street, Cambridge, MA 02139, USA
[8]School of Information, University of Michigan
105 S. State Street, Ann Arbor, MI 48104 USA
[9]School of Economics & Management, University of Cyprus
1 Panepistimiou Avenue, 2103 Aglantzia, Cyprus

[†]These authors contributed equally to this work.
[*]To whom correspondence should be addressed; E-mail: eckles@mit.edu, sinan@mit.edu.

This PDF file includes:
Materials and Methods
Figures S1 to S40
Tables S1 to S11
Additional References

Contents

S1 Data **7**

S1.1 Safegraph Data 7

S1.1.1 Social Distancing Measures 7

S1.1.2 Industry Visits Shares and Shifts 8

S1.1.3 Geographic Adjacency Matrix 9

S1.2 Facebook Data 10

S1.2.1 Facebook Mobility Metrics 12

S1.2.2 Facebook Social Connectedness Index 12

S1.3 Weather Data 13

S1.4 Government Intervention Data 15

S1.5 Census Data 19

S1.6 Variation in Alter Policies 19

S1.7 Comparison of Outcome Metrics 20

S2 Difference-in-differences analysis **20**

S2.1 Analysis without spillovers 21

S2.2 Analysis with Geographic Spillovers 23

S2.3 Analysis with Geographic and Social Spillovers 24

S2.4 State-Level Policy Contrasts 26

S2.5 Dyad-level Difference-in-differences 27

S2.6 Robustness Checks 37

S2.6.1 Leading and Lagging Policy Effects 37

S2.6.2 Randomization Inference 39

S3 Identifying Endogenous Social Effects **42**

S3.1 Weather Instruments 45

S3.2 Industry Visits Shift-Share Instrument 47

S3.3 Model Specification 47

S3.3.1 Flexibly Controlling for the Effect of Weather 52

S3.3.2 Sparse Instrumental Variable Selection Procedure 54

S3.4 Results 55

S3.4.1 Visualization of the Instrumental Variables 56

S3.5 Robustness Checks 57

S3.5.1 Varying the Number of Selected instruments 57

S3.5.2 Only Weather Instruments 57

S3.5.3 Only Shift-share Instruments 57

S3.5.4 Randomization Inference 63

S3.5.5 Excluding Counties with Correlated Weather 64

S4 Asymptotic Inference with Adjacency- and Cluster-robust Standard Errors **66**

S5 Ranking States' Influence **71**

S6 Analytical Model of the Loss From Anarchy **72**

S6.1 Nash Equilibrium Solution 77

S6.2 Socially Optimal Solution 79

S6.3 Loss from Anarchy 80

S6.3.1 Model Calibration 84

S7 Software **85**

References **87**

Materials and Methods

S1 Data

S1.1 Safegraph Data

This paper uses data provided by Safegraph, Inc.[1] which collects location data from approximately 22 million unique devices each month. These data consist of "pings," which identify the coordinates of a particular smartphone at a given moment in time. Each user of these smartphones has given permission for his or her location to be tracked by a variety of mobile applications. Safegraph in turn partners with many of these mobile application services that obtain affirmative opt-in consent from users to collect their location data. To preserve anonymity, the data is aggregated to the Census block group (CBG) level and all CBGs with fewer than five observations are omitted. We use this data to construct social distancing measures, industry visit shares and shifts, and a geographic adjacency matrix to capture cross-county mobility.

S1.1.1 Social Distancing Measures

We use data from Safegraph's "Social Distancing Metrics" and "Weekly Patterns" covering the period from January 1, 2020 through April 18, 2020 to construct two measures of social distancing compliance. First we construct "Fraction of Devices Not at Home" (NHDF) as[2] $1 - completely\_home\_device\_count/device\_count$ which is simply the fraction of devices not detected to be entirely at home on a given day in a Census block group. A state-aggregated visualization of this data can be seen in Fig. S2A. As can be seen in this figure, there is a strong pattern of weekly cyclicality, where peaks in the time series correspond to weekends. For our analysis, we aggregate this data to the county level to conform with the rest of our analysis. To allow us interpret our results in relative terms, we apply a hyperbolic inverse sin transformation,

[1] `https://www.safegraph.com`

[2] This transformation is so that it goes in the same direction as our other movement based measure.

leading to one of our four main social distancing measures, asinh(NHDF). We use the arcsinh transformation because it is defined at zero, and, except for very small values of y, $\sinh^{-1}(y) \approx \log(2y)$. Thus, arcsinh-transformed variables can be interpreted similarly to log-transformed variables (*1*). The second metric we construct is "Census Block Groups Visited" (dCBGVs), which measures the mean number of census block groups an individuals within a specific have visited on that day. Specifically, this is formed by summing *non_home_cbg_visits_within_county*, *cbg_visits_outside_county*, and *home_cbg_visits* across all devices and simply dividing by the *device_count*. To allow for a relative interpretation, we apply a log transformation, yielding our second main social distancing measure log(dCBGVs).

S1.1.2 Industry Visits Shares and Shifts

We construct baseline industry visit shares for each county in the pre-COVID period by computing the number of visits by Safegraph users to a particular industry in that county as a proportion of total visits to all industries in that county during the months of January and February 2020 (visits are aggregated from 1/1/20-2/29/20). We construct relative national shifts in industry visits by computing the total number of visits to a particular industry across all of US every day from 3/1/20-4/18/20 and calculating proportional increase from baseline industry visits (i.e. total number of visits to that industry across all of US in January and February 2020). The industry visit shift-share measures S_{it} of a county i on a particular day t are computed by taking a product of the baseline industry visit shares in that county and the relative national level shifts in industry visits on that day. While analyzing ego county i's social distancing outcomes, we instrument for alter county j's social distancing outcomes by a weighted average of alter counties' industry visit shift-shares. More formally: $S_{-it} = \sum_j w_{ij} S_{jt}$ where S_{-it} is the weighted average of all alter county's industry visit shares and weights $w_{ij} = \frac{n_j SCI_{ij}}{\sum_k n_k SCI_{ik}} : k \neq i$ and n_j is the population of county j.

S1.1.3 Geographic Adjacency Matrix

In order to better understand how U.S. counties are connected due to intercounty travel, we construct a geographic adjacency matrix that captures average pre-crisis mobility patterns. As mentioned earlier, Safegraph does not provide data at the individual device level, so we are unable to determine the number of devices from county i that visit county j on a given day. However, for each day in 2020, Safegraph's "Social Distancing" dataset records how many devices from each "origin" Census block group (CBG) visit each possible "destination" CBG.[3][4]

We first create an "unsmoothed" geographic adjacency matrix. For every pair of U.S. counties i and j, we create an edge with weight equal to the total number of visit-days from devices whose home CBG is in county i to CBGs in county j across all of January and February. Thus, if a device from a CBG in county i visits two CBGs in county j on day t, we count two visit-days to county j. It is worth noting that this measure of cross-county mobility is asymmetric. We also set all self-loops to 0, i.e., even though devices may visit non-home CBGs within their own county, these visits do not contribute to the geographic adjacency matrix.

We smooth this adjacency matrix using an Empirical Bayes procedure. This procedure uses the full set of county pairs to estimate a Dirichlet multinomial posterior, $\hat{\alpha_j}$. $\hat{\alpha_j}$ will be used as an empirical prior for the percentage of visit-days to county j coming from county i. This empirical prior is conceptually similar to the average number of visit-days from each county i across all counties j. We combine this empirical prior with the observed county-level data to estimate a Dirichlet multinomial posterior distribution for every individual county. This results

[3] In the raw Safegraph data, the set of destination CBGs includes the origin CBG.

[4] As a simple example, suppose the following is true: there are 9 devices in Census block group A, and on a given day 3 of them visited Census block group B, 5 visited Census block group C, and 1 visited Census block group D. Similiarly, there are 7 devices in Census block group B, and on a given day 2 of them visited Census block group A, 3 visited Census block group C, and 0 visited Census block group D. Census block groups A and B belong to county i, and Census block groups C and D belong to county j. While we are unable to measure how many devices from county i visited county j on the day in question, because we are unable to tell how many devices visited both Census block group C and Census block group D, we *are* able to report that there were 9 visit-days to CBGs in county j from devices in county i.

in a maximum a posterior (MAP) estimate of ($\hat{\pi}_j$) for each county j. For each origin county, we let n_{ij} be the number of visit-days from devices with a home in county i to CBGs in county j. We then calculate the posterior percentage of visit days to county j coming from county i as:

$$\hat{\pi}_{ij} = \frac{n_{ij} + \hat{\alpha}_j}{N_j + K}$$

where $N_j = \sum_i n_{ij}$ and $K = \sum_j \hat{\alpha}_j$. K can be scaled up or down in order to make the empirical prior for this smoothing procedure stronger or weaker. Agresti and Hitchcock's survey of Bayesian methods for categorical data analysis (*2*) describes this smoothing procedure in greater detail.

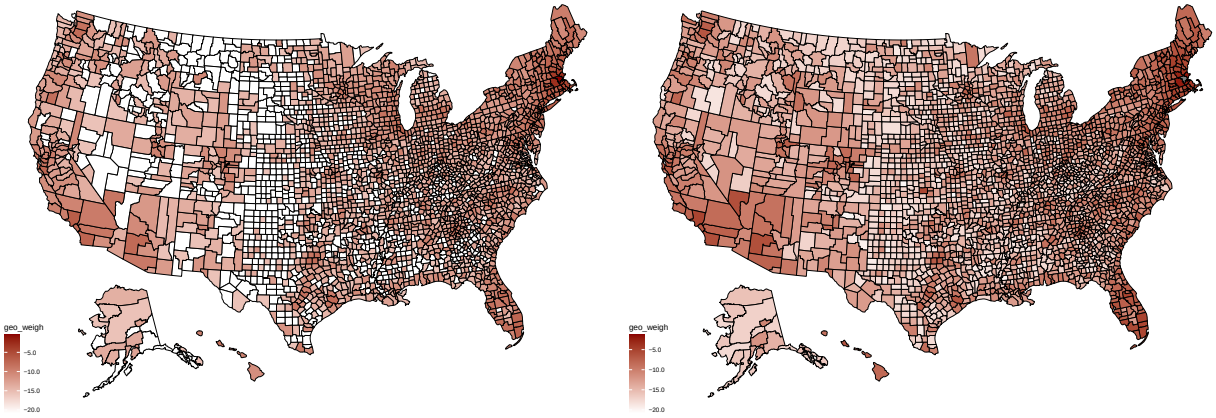

Figure S1: The spatial distribution of the raw (left) and smoothed (right) weighted geographic mobility graph of Middlesex County, MA. Smoothing does not substantially alter the geographic adjacency matrix.

S1.2 Facebook Data

Facebook's Data for Good[5] Initiative provides anonymized, aggregate privacy preserving data products for use in humanitarian organizations and universities. For this paper, we specifically make use of two main datasets: US county mobility metrics and the Social Connectedness Index.

[5] https://dataforgood.fb.com/

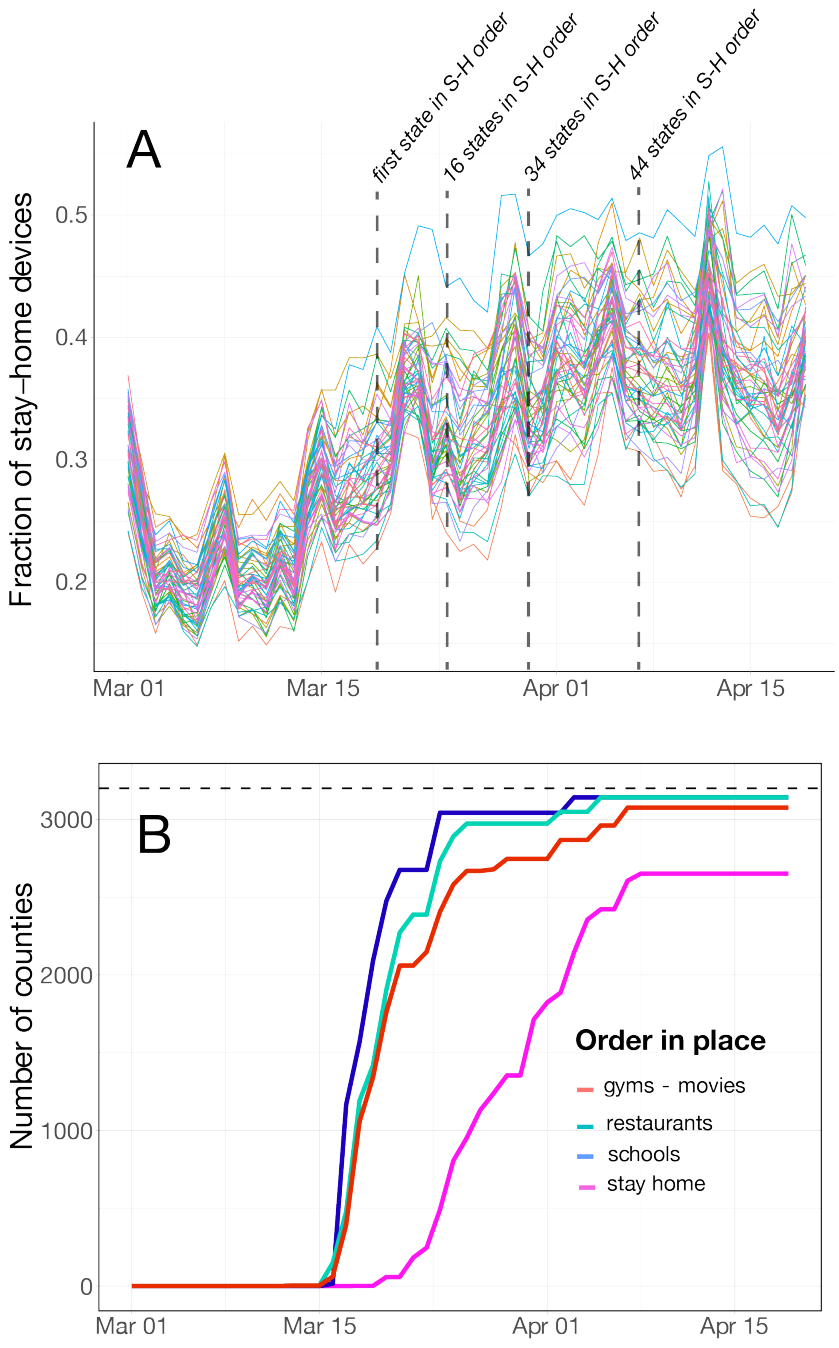

Figure S2: A. The Fraction of Devices at Home for all 50 states in the United States as a function of time. B. The number of counties with different orders-in-place as a function of time.

S1.2.1 Facebook Mobility Metrics

This data provides aggregated county-level mobility metrics built from 27 million Facebook users who have the Facebook app installed on their smartphones and have turned on location history in the United States. These metrics are are provided starting from March 1st, 2020, and continue to be updated daily. We use the data from March 1st until April 18th so that it aligns with our Safegraph data. Our last two social distancing measures come directly from this data.

To start, *ratio_single_tile_users* captures the fraction of users whose location history does not leave a single 0.6km by 0.6km block (a "level 16 Bing tile") during a 24 hour period. We subtract this measure from 1 to construct *not_single_bing_tile_users* (NSBTUs). Similar to NHDF from the Safegraph data, we also apply a inverse hyperbolic sine transformation to form our third main social distancing measure, asinh(NSBTUs). Lastly, *bing_tiles_visited_relative_change* (BTVRC) measures the relative change in the average number of Bing tiles that a Facebook user was present in during a 24-hour period compared to a the pre-crisis baseline. This baseline is computed by averaging across all of February 2020 for each specific day of the week. Since this measure already affords a relative interpretation and we do not directly have access to the raw counts, we simply use it as it is.

It is important to note that only counties with at least 300 active users in that county are tracked. Since we want to compare across both Safegraph and Facebook measures, we limit our analysis to the 2,502 counties that have complete data starting for all dates later than March 1st.

S1.2.2 Facebook Social Connectedness Index

We use the Social Connectedness Index (SCI) dataset which contains a measure of social connectedness at the U.S. county level. The SCI between two counties i and j is defined as:

$$SCI_{ij} = \frac{\text{fb\_connections}_{ij}}{\text{fb\_users}_i \times \text{fb\_users}_j} \tag{S1}$$

Here, fb_users$_i$ and fb_users$_j$ are the number of Facebook users in locations i and j, and fb_connections$_{ij}$ is the number of Facebook friendship connections between the two. SCI_{ij}, therefore, measures the relative probability of a Facebook friendship link between a given Facebook user in location i and a Facebook user in location j. Thus, if this measure is twice as large, a Facebook user in i is about twice as likely to be connected with a given Facebook user in j. Perhaps unsurprisingly, the majority of Facebook connections are geographically proximate. For the average county, 55.4% of friends live within 50 miles and 70.3% of friends live within 200 miles. More details on the SCI dataset, including more information regarding the construction of the SCI metric, are provided in (*3*). The SCI index used in this study is drawn from a snapshot of Facebook friendships as of December 31, 2019.

S1.3 Weather Data

Our weather data comes from the Global Historical Climatology Network (GHCN) database maintained by the National Oceanic and Atmospheric Association (NOAA). A detailed description of this data can be found in (*4*). The GHCN data contains daily observations of maximum temperature, minimum temperature, and precipitation for roughly 62,000 weather stations in the United States. We use the geographic coordinates of each weather station in order to derive maximum temperature and precipitation data for each county in our dataset as described below.

Starting with the daily GHCN data we restrict to weather stations within the United States that contain either maximum temperature or precipitation weather elements. Of the 3,233 counties in the United States, 243 have no weather stations and 967 have fewer than three weather stations. For each of these counties with fewer than three weather stations, we assign the nearest three stations within a 100 kilometer radius of the county centroid to this county. The above procedure assigns weather stations to 97.2% of US counties.

To obtain the weather data for each county-day pair, we average the measurements obtained

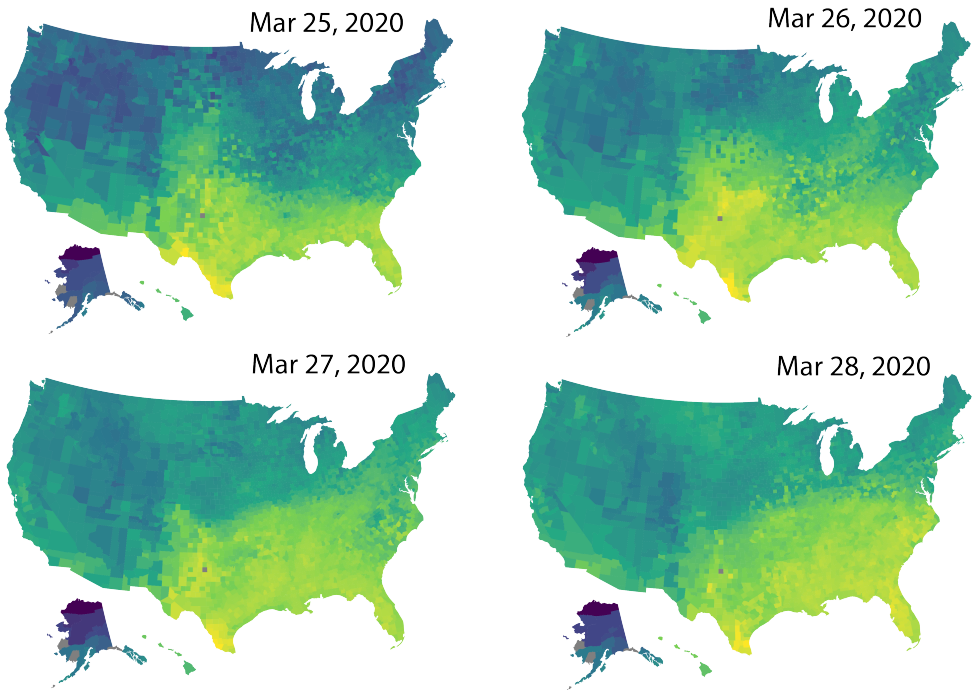

Figure S3: The maximum daily temperature (in degrees Celsius) at the county level over four consecutive days. The brighter color indicates higher maximum temperature.

by the weather stations within each county, including the weather stations assigned to the county if the county has no station within its borders as discussed above. When stations do not have either temperature or precipitation readings, we average the field excluding the missing values. In the sample, 7.3% of station days do not have precipitation readings and 64% do not have temperature readings. After aggregating to the county level, 19.9% and 4.4% of the county-day observations were still missing maximum temperature and precipitation values respectively. To fill in these missing values, we find the three nearest stations within a 100-kilometer radius that contain the missing field and take the average across the 3 measurements. This gives us a precipitation and a maximum temperature reading for 99.9% of county days in our sample. For this study, we use the weather data from January 1, 2020 through April 19, 2020.

We can visualize the maximum temperature and precipitation data across all counties in the United States over time, as shown in Figures S3 and S4, respectively. The histograms of

maximum daily county temperature across all counties in the United States over time is depicted in Fig. S5, where the histogram color darkens as time progresses. Sample histograms for specific dates are depicted in Figures S6 and S7, where the distributions of average daily county precipitation and temperature throughout the United States are respectively depicted. Interestingly, Fig. S7 depicts a bimodal distribution for the distribution of maximum temperature across counties. Sample histograms for a specific county are depicted in Figures S8 and S9, where the distributions of average daily county precipitation and temperature throughout time are respectively depicted.

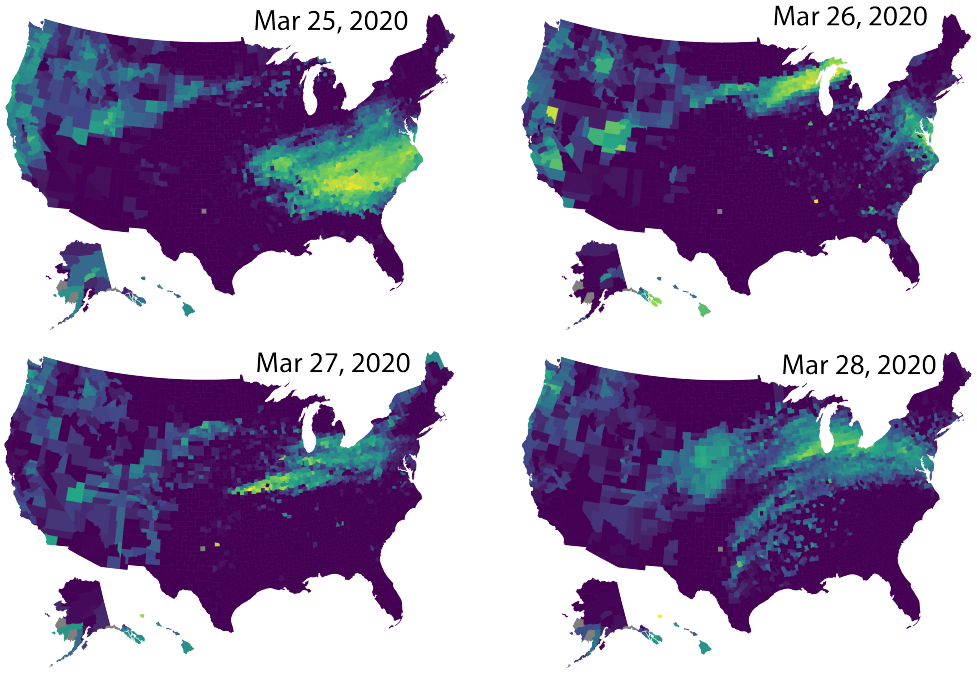

Figure S4: Daily precipitation (in millimeters) at the county level over four consecutive days. The brighter color indicates higher precipitation.

S1.4 Government Intervention Data

We use data collected by researchers at Johns Hopkins University that documents various local or federal government interventions that came into effect in various counties at various times (*5*).

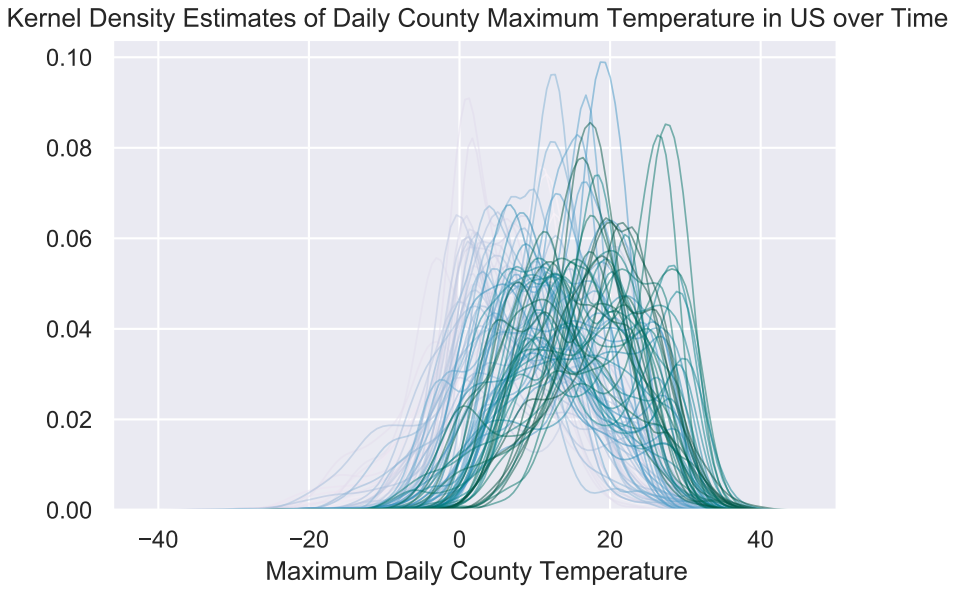

Figure S5: Kernel density estimates of county daily maximum temperature (in degrees Celsius) in the US. Each line represents the distribution of maximum temperature on a specific day, and the color of the lines gets darker as time progresses.

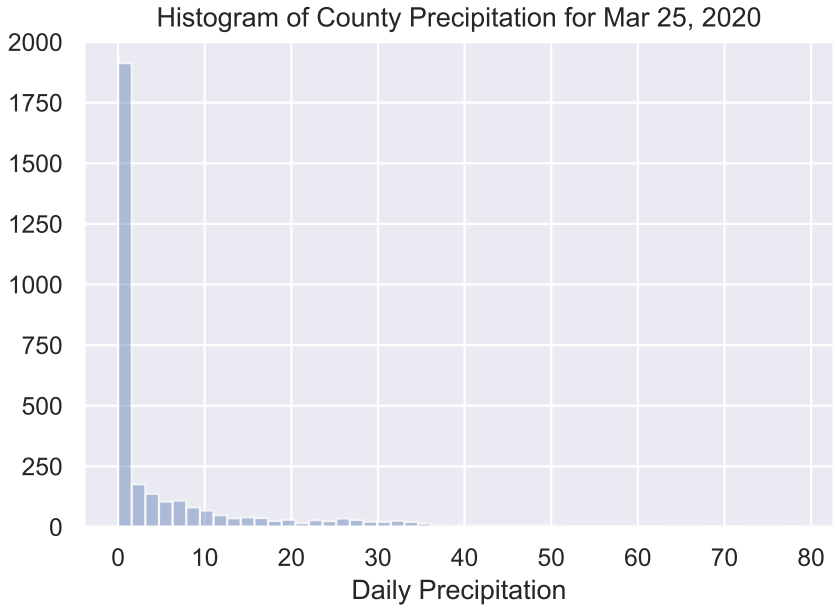

Figure S6: Histogram of precipitation (in millimeters) for one specific day in the US.

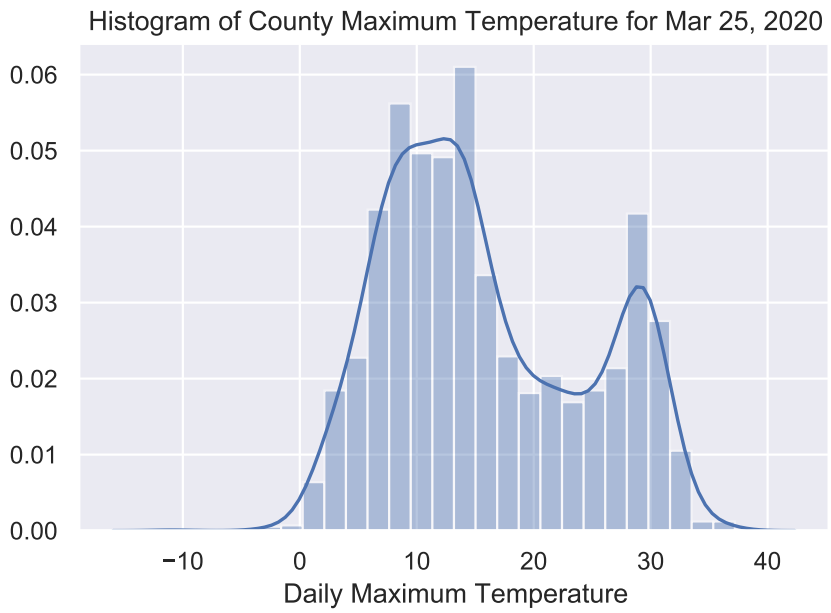

Figure S7: Histogram of daily maximum temperature (in degrees Celsius) for one specific day in the US.

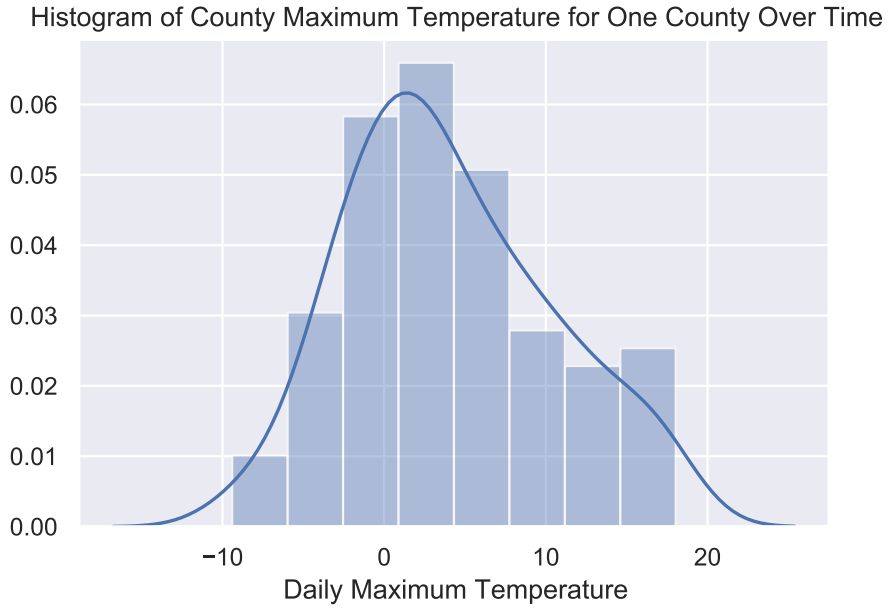

Figure S8: Histogram of daily maximum temperature (in degrees Celsius) for one specific county in the U.S..

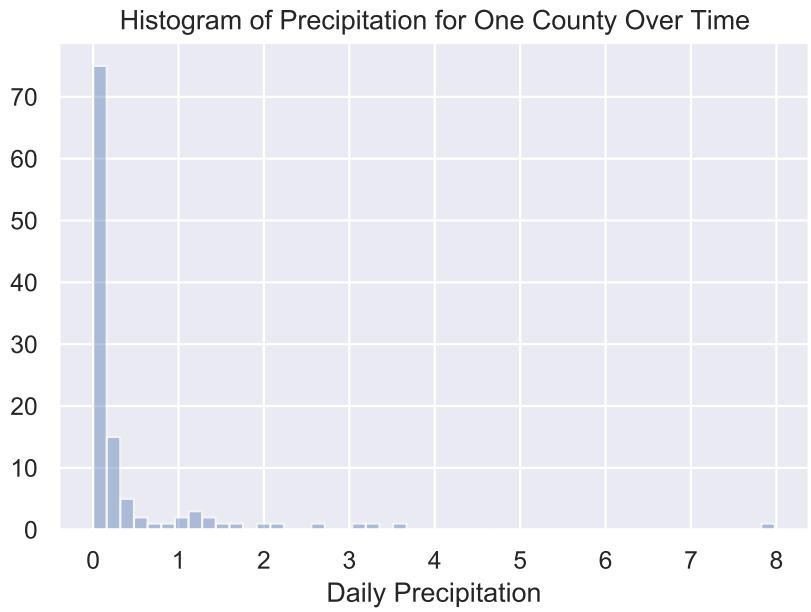

Figure S9: Histogram of precipitation (in millimeters) for one specific county in the US.

In particular, the data contains information pertaining to 1) school closures, 2) bans on gatherings of large sizes, 3) closing of restaurants, bars, and gyms, and 4) shelter-in-place orders. In the policy space considered in this paper, we use only two of these policy interventions:

- The closing of restaurants, movie theaters, and gyms (GMR).
- Shelter-in-place (SH)

In the above, we use shelter-in-place to refer not only to policies explicitly referred to as "shelter-in-place" by policymakers, but also other similar orders that are named differently (e.g., "stay at home" orders). The two policies above differ in their level of severity of curbing mobility, with shelter-in-place policy being more severe than the closing of restaurants, movie theaters, and gyms. Note that we do not include the school closure policies in our analyses because of the limitations of the Safegraph and Facebook data that we use to track movement. Many students in schools in US do not have access to smartphones. Furthermore, relative to

other policy interventions enacted across the United States, there is relatively little variation in the enactment of school closure policies. Both of these factors make it difficult to measure the impact of school closure policies on the types of outcomes we measure in this paper. We do not include bans on gatherings since most U.S. states were in this policy state for only a few days before implementing more severe restrictions, and because we suspect that the period of time during which these measures were in place corresponds with high levels of travel as people rushed to stock up on supplies, returned home, or relocated to less urban destinations. Lastly, although our analysis includes both of these policies, in the main text, we focus particularly on shelter-in-place, as both the impact of such policies and our confidence in our inferences are much higher. Fig. S2B shows the fraction of counties that had implemented one of the policies of interest as of a particular date.

S1.5 Census Data

We use county-level population estimates for 2018 from the U.S. Census Bureau to weight all county-level analyses by population. We use these same population estimates to re-weight the Facebook Social Connectedness Index, so that each edge corresponds to the fraction of county i's Facebook friendships belonging to Facebook users in county j. Finally, we use this Census data in order to calculate each state's "total influence," as reported in Fig. 3c of the main text.

S1.6 Variation in Alter Policies

Fig. S10 shows the weighted median percentage of socially connected alter counties and geographically connected alter counties implementing shelter in place polices across our sample over time, as well as the weighted interquartile range for the same measures. Most of the policy variation we observe occurs between mid-March and early April.

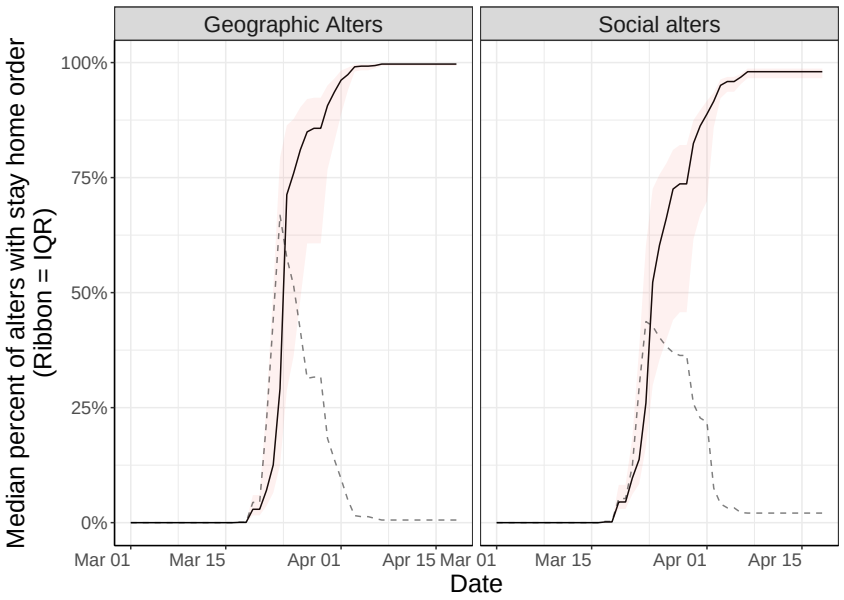

Figure S10: The weighted median percentage of geographic alter counties and social alter counties that have implemented shelter in place policies in our sample over time. Red ribbons correspond to the weighted 25th percentile and the weighted 75th percentile for the same quantities. The dashed line plots the interquartile range over time.

S1.7 Comparison of Outcome Metrics

Fig. S11 shows the weighted mean of each of our four outcome variables over time. On average, each of these outcomes begins to decline after mid-March, when social distancing generally begins across the United States. Fig. S12 shows the correlation between each of our four outcome variables. In general, all four outcome variables are quite correlated.

S2 Difference-in-differences analysis

In this section we present a series of results on the impact of policy interventions on social distancing behaviors using a number of difference-in-differences model specifications. This section is organized as follows: Section S2.1 presents a model that assumes no spillover effects, Section S2.2 presents a model that accounts for geographic spillovers, and Section S2.3 presents

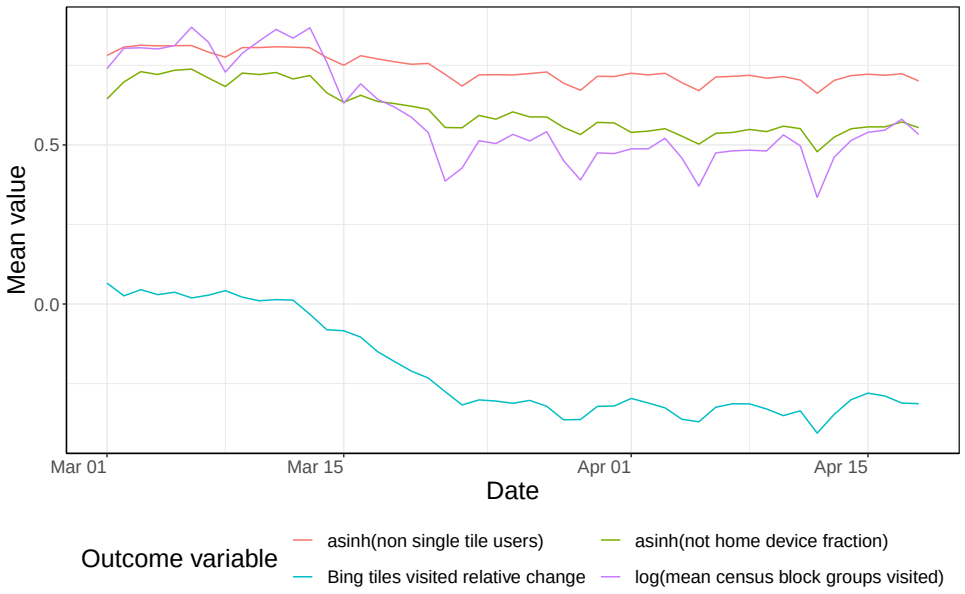

Figure S11: The weighted mean of each outcome variable for our analysis over time.

a model that accounts for both geographic and social spillovers. From there, we investigate state-level policy contrasts in Section S2.4. Finally, we directly examine geographic spillovers using a dyadic DiD in Section S2.5. We consider four outcome variables defined in Section S1.1.1 and Section S1.2.1: asinh(NSBTUs), asinh(NHDF), log(dMCGBVs), and BTVRC. All statistical inference is network-adjacency- and state-cluster-robust; see Section S4 for more details.

S2.1 Analysis without spillovers

We begin by presenting results which assume that there are no policy spillovers between counties or states. Here we estimate the following model specification:

$$Y_{it} = D_{it}\delta_1 + f(W_{it}) + \alpha_i + \tau_t + \epsilon_{it} \tag{S2}$$

where Y_{it} denotes the social distancing outcome, D_{it} is a 2-dimensional row vector of policies that denotes whether either GMR_{it} (gym, movie, and restaurant ban) or SH_{it} (shelter-in-place

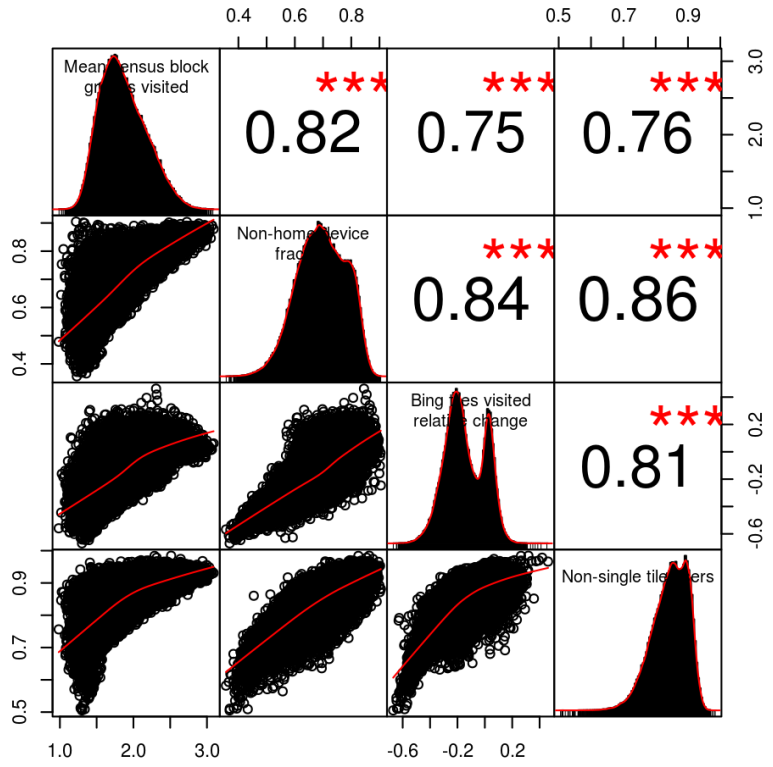

Figure S12: The correlation between different outcome variables. Diagonal tiles show histograms for each outcome. Lower tiles show scatter plots for each pair of outcomes. Upper tiles show the measured correlation between each pair of outcomes, along with p-values.

order) has been enacted in county i in time period t. The associated parameter, δ_1, is a 2-dimensional column vector of the main parameters of interest δ_1^{GMR} and δ_1^{SH} which capture the effect of each policy on county-level outcomes; that is, $D_{it} = [\text{GMR}_{it}, \text{SH}_{it}]$ and $\delta_1 = \begin{bmatrix} \delta_1^{\text{GMR}} \\ \delta_1^{\text{SH}} \end{bmatrix}$ so that $D_{it}\delta_1 = \delta_1^{\text{GMR}}\text{GMR}_{it} + \delta_1^{\text{SH}}\text{SH}_{it}$. Since we coded GMR to be zero once a county is under a shelter-in-place order, δ_1^{SH} is interpreted as the total effect of a shelter-in-place order rather than the marginal effect. $f(W_{it})$ is a term that flexibly controls for the potential non-linear impact of weather using a "double machine learning" approach. We describe this procedure in Section S3.3.1. α_i and τ_t represent a set of county and time fixed effects, and ϵ_{it} denotes the error term.

The results, found in Table S1, suggest that both the shelter-in-place and closure policies are associated with decreases in county-level mobility. Shelter-in-place orders are associated with a substantially larger reduction in mobility across all outcomes compared to closing gyms, movies theatres, and restaurants.

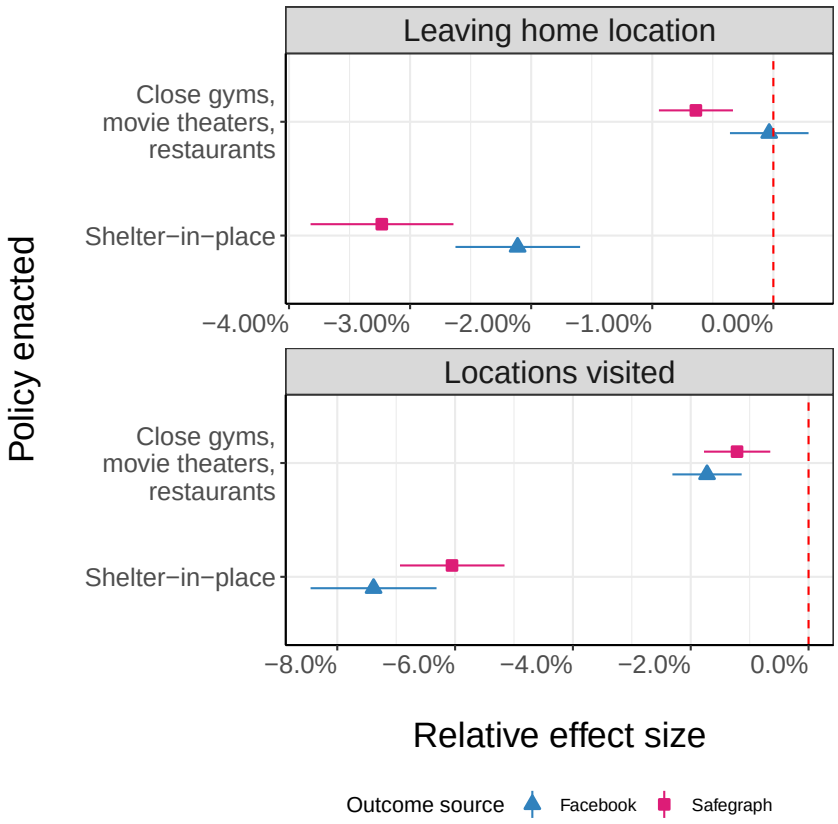

Figure S13: Difference-in-differences analysis of policy effects on mobility, neglecting any spillovers between counties.

S2.2 Analysis with Geographic Spillovers

However, as discussed in the main text, it is reasonable to think that spillovers exist (*6–11*). These spillovers may not only be biasing estimation of the direct policy effects, but are of significant interest in and of themselves. As many local officials have noted, the policies of geo-

graphically proximate counties may also be impacting a focal county. To capture such potential geographic spillovers, we estimate the following model specification:

$$Y_{it} = D_{it}\delta_1 + D^{geo}_{-it}\delta_2 + f(W_{it}) + \alpha_i + \tau_t + \epsilon_{it} \tag{S3}$$

The main difference between this specification and Equation 2 above is the inclusion of D^{geo}_{-it}, a 2-dimensional row vector where each element is the weighted average of peer county policies. Specifically, $D_{-it} = \sum_j \omega_i j * D_{jt}$ where weights $\omega_{ij} = a^{geo}_{ij} / \sum_k a^{geo}_{ik}$ and each a^{geo}_{ij} is the (i, j)th element of the geographic adjacency matrix A^g. Details of the construction of this matrix can be found in Section S1.1.3.

Results are presented in Table S2. After accounting for the impact of geographic spillovers, the analysis suggests that both shelter-in-place and closure policies are associated with increased social distancing behaviors, with more pronounced distancing behavior associated with shelter-in-place policies as expected. We also find that the shelter-in-place decisions of geographically connected counties have a causal effect on the social distancing behaviors of those in the focal county.

S2.3 Analysis with Geographic and Social Spillovers

Though many officials have openly discussed the potential for geographic spillovers, very little discourse has focused on the potential for social spillovers. In fact, such social spillovers as that are mediated by communication technology may be especially pronounced since people are actively socially distancing and spending more time connecting digitally. To investigate if social spillovers exist beyond the geographic spillovers we have already estimated, we estimate the following model specification:

$$Y_{it} = D_{it}\delta_1 + D^{geo}_{-it}\delta_2 + D^{social}_{-it}\delta_3 + f(W_{it}) + \alpha_i + \tau_t + \epsilon_{it} \tag{S4}$$

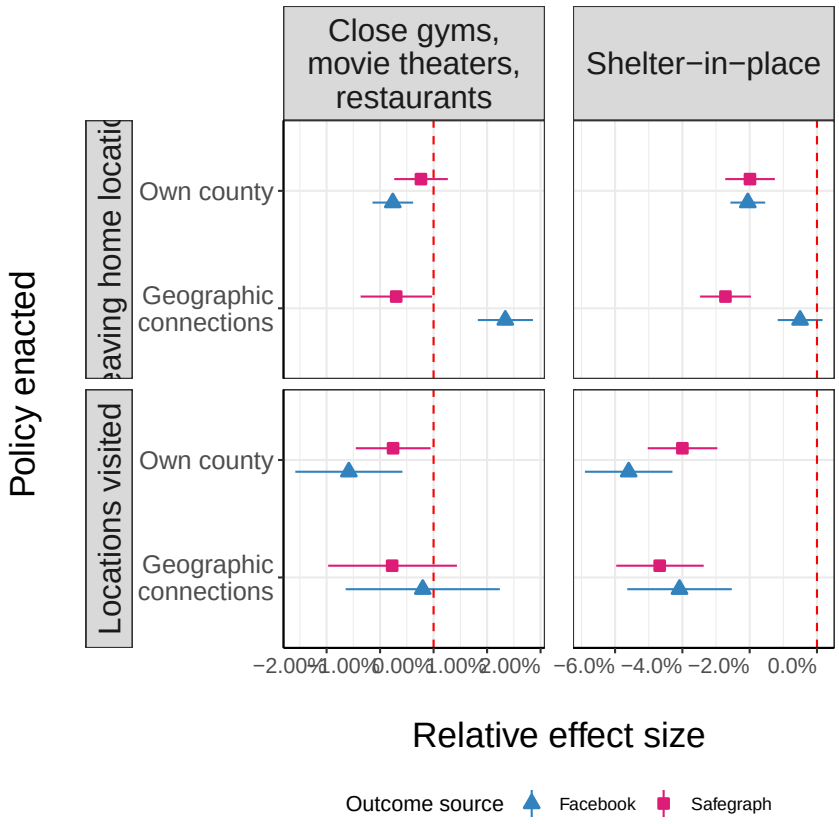

Figure S14: Difference-in-differences analysis of policy effects on mobility, allowing for spillovers via geographic connections.

This specification is the same as above, but we also estimate the effects of a weighted average of policies in alter counties, D_{-it}^{social}, where weights are based on a social adjacency matrix, A^s, that is constructed from the Facebook social connectedness index. In particular, $D_{-it}^{social} = \sum_j w_{ij} * D_{jt}$ where weights $w_{it} = a_{ij}^{social} / \sum_k a_{ik}^{social}$.

The results of this specification can be found in Table S3. When controlling for alter policies, the estimated coefficients on own-county policies are more muted, indicating a shelter in place order is associated with a 1 to 2 percent reduction in mobility across our outcome measures. Only the estimates using the Facebook mobility data are statistically significant, however.

Spillovers appear to be important in this specification as well. In particular, social-alter

policies, or the average policies enacted in other counties weighted by social connections, are more strongly associated with social distancing behavior than geo-alter policies, or policies in those localities that are geographically closer to the focal location. However, accounting for social spillovers causes the coefficient estimate for geographic spillovers to decrease in magnitude, and in some cases change sign. We believe this is a statistical artifact resulting from the correlation between the geographic and social adjacency matrices and the limitations of the linear functional form of our model specifications.

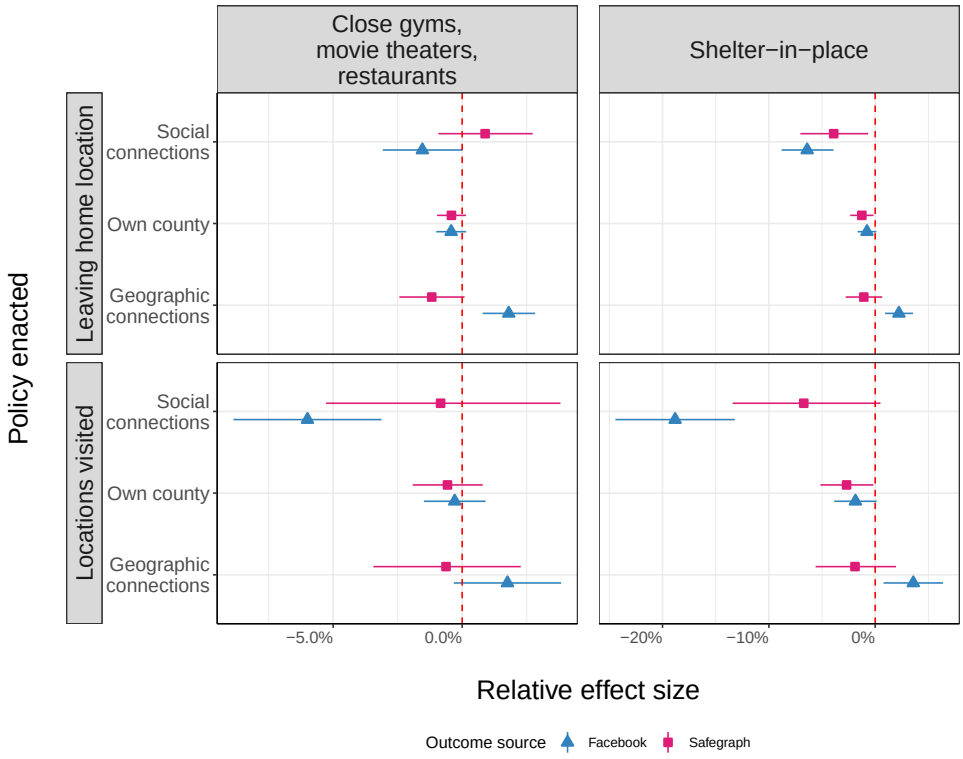

Figure S15: Difference-in-differences analysis of policy effects on mobility, accounting for spillovers via both geographic and social connections.

S2.4 State-Level Policy Contrasts

Our analysis so far has been restricted to the county level. However, in the vast majority of cases, social distancing policy is set at the state level or above. To explore counterfactuals

related to different state-level policy contrasts (i.e. one state alone implements a shelter-in-place order vs. one state alone chooses *not* to implement a shelter in place order), we fit the following model specification:

$$Y_{it} = D_{it}\delta_1 + D_{-it}^{geo,ss}\delta_2^{ss} + D_{-it}^{geo,ds}\delta_2^{ds} + D_{-it}^{social,ss}\delta_3^{ss} + D_{-it}^{social,ds}\delta_3^{ds} + f(W_{it}) + \alpha_i + \tau_t + \epsilon_{it} \quad \text{(S5)}$$

The key difference between this model specification and Equation 4 is that the D_{-it}^{geo} and D_{-it}^{social} have been decomposed into the same state (ss) and different state (ds) components. The results from estimating this model can be found in Table S5. Fig. S16 shows this model's estimate of the effects of own state's and alter states' shelter-in-place policies, after combining geo- and social-spillovers.[6] Qualitatively, these results are quite similar to those found in Table S3. We also estimate a similar model that considers only geographic spillovers, but differentiates between same state and different state alter counties. The results of that model can be found in Table S5.

Fig. S17 compares the combined own-state and alter state effects of shelter-in-place when considering no spillovers, only geographic spillovers, and both geographic and social spillovers. After accounting for social spillovers, the estimated magnitude of the effect of alter states' shelter-in-place policies increase by more than a factor of 2.

S2.5 Dyad-level Difference-in-differences

To establish the existence and relevance of geographic spillovers when enacting social distancing policies, we estimate the following two-way fixed effects difference-in-differences model, which relies on directed, dyad-level data. We create our directed, dyadic dataset by first looking at the movement from every county i to every alter county $j \neq i$ from January 1, 2020 to April 18, 2020 from the same data used to construct the geographic adjacency matrix. We proceed

[6] Another benefit of combining the two spillover effect estimates is that the combined spillover effect estimate is not impacted by the correlation of the geo- and social- adjacency matrices.

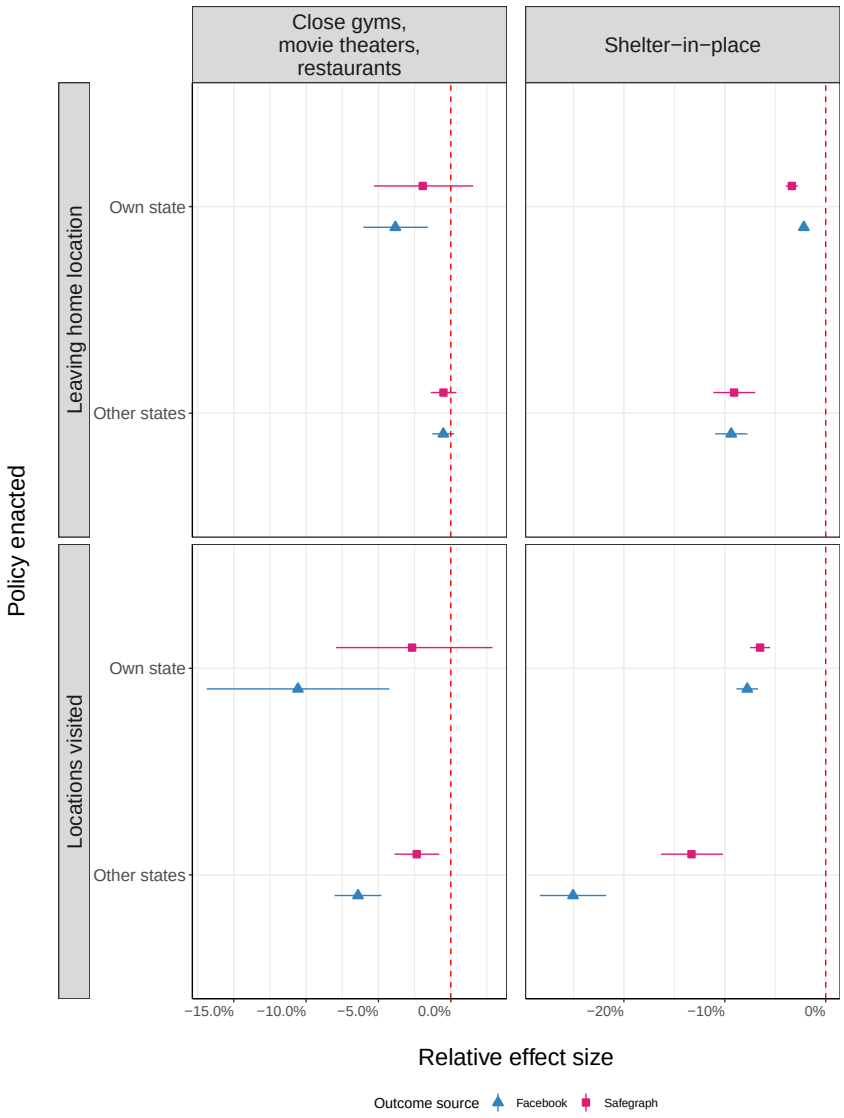

Figure S16: The estimated aggregate effect own-state and other states' shelter-in-place policies after differentiating between same state and different state spillovers. Direct effects, geo-spillovers, and social-spillovers are combined at the ego- and alter-level.

to analyze data for every ij pair for which there was at least one device moving from i to j on at least one day in January, February, March, and April, and for which movement from i to j occurred on at least 30 distinct days. This produces a set of 548,450 ij pairs. We then use data

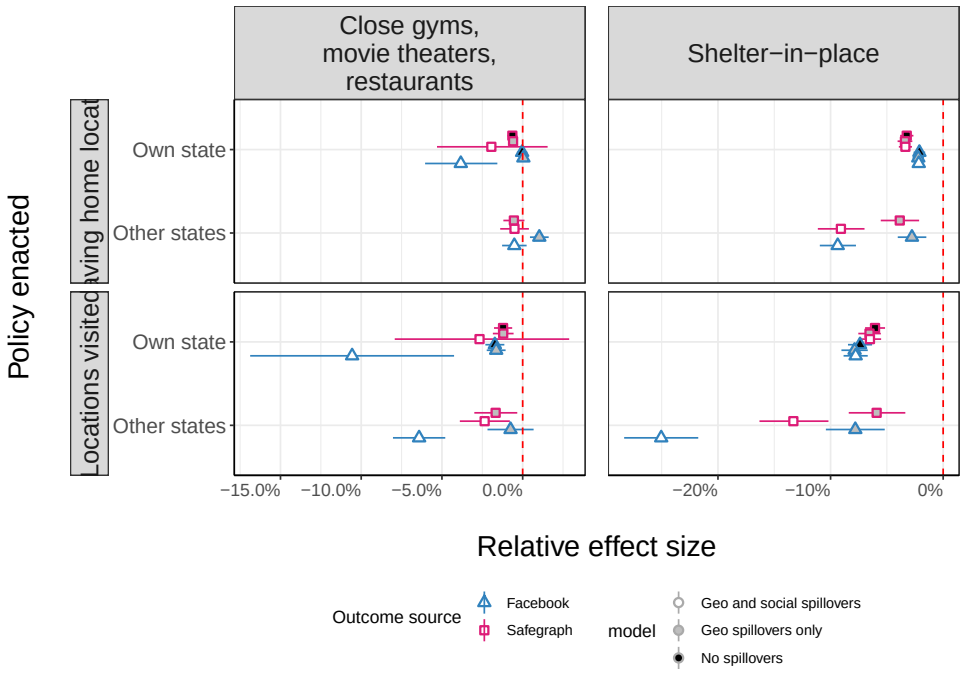

Figure S17: A comparison of the estimated impact of own-state and other states' shelter-in-place policy under three different specifications: a model with no spillovers, a model that includes only geographic spillovers (and distinguishes between within state and out of state counties), and a model that includes both geographic and social spillovers (and distinguishes between within state and out of state counties).

from March 1, 2020 to April 18, 2020 to estimate the following base model:

$$m_{i \to j,t} = D_{ijt}\delta + \alpha_{ij} + \tau_t + \epsilon_{ijt}, \tag{S6}$$

where $m_{i \to j,t}$ is a measure of the extent to which residents of county i are traveling into county j on day t. In this dyadic analysis, D_{ijt} is a measure of whether county i and county j have both enacted a given policy as of day t. Specifically, D_{ijt} is a 6-dimensional row vector that represents each possible combination of dyadic policy states, i.e., $\text{GMR}_{it} = 0$ and $\text{GMR}_{jt} = 1$, $\text{GMR}_{it} = 1$ and $\text{GMR}_{jt} = 0$, $\text{GMR}_{it} = 1$ and $\text{GMR}_{jt} = 1$, $\text{SH}_{it} = 0$ and $\text{SH}_{jt} = 1$, $\text{SH}_{it} = 1$ and $\text{SH}_{jt} = 0$, and $\text{SH}_{it} = 1$ and $\text{SH}_{jt} = 1$. The cases of $\text{GMR}_{it} = 0$ and $\text{GMR}_{jt} = 0$ and $\text{SH}_{it} = 0$ and $\text{SH}_{jt} = 0$ are omitted to prevent perfect multicollinearity and are taken as the

Table S1: The effect of policy interventions on county-level outcomes

	Dependent variable:			
	asinh(NSBTUs)	asinh(NHDF)	log(dCBGVs)	BTVRC
	(1)	(2)	(3)	(4)
Shelter-in-place	-2.13^{***}	-3.29^{***}	-6.24^{***}	-7.38^{***}
	(0.27)	(0.31)	(0.48)	(0.55)
Close gyms/movies/restaurants	-0.03	-0.64^{***}	-1.22^{***}	-1.72^{***}
	(0.17)	(0.16)	(0.29)	(0.30)
Conley s.e.	Yes	Yes	Yes	Yes
County fixed effect	Yes	Yes	Yes	Yes
Day fixed effect	Yes	Yes	Yes	Yes
Observations	122,598	122,598	122,598	122,598
R^2	0.43	0.42	0.39	0.38
Adjusted R^2	0.41	0.41	0.38	0.37
Residual Std. Error (df = 120046)	6.36	8.54	14.85	15.81

Notes: NSBTU refers to the fraction of Facebook users in a given county who visit multiple Bing tiles on a given day. NHDF refers to the fraction of devices that are not completely at home, as measured by Safegraph. dCBGVs refers to the mean number of Census block groups devices from a given county visit, as measured by Safegraph. BTVRC refers to the relative change in the number of Bing tiles users visit, as measured by Facebook. All values multiplied by 100. *p<0.1; **p<0.05; ***p<0.01

base cases. α_{ij} is a dyad-level fixed effect, and τ_t is a day-level fixed effect. As with before, when a given county has enacted shelter-in-place, we also set the closing of restaurants, movie theaters, and gyms policy to 0. Observations are weighted by the number of devices observed in county i on day t, and we report cluster-robust standard errors that are clustered at the level of county i's U.S. state and county j's U.S. state.

Column 1 of Table S6 reports our estimates, where $m_{i \to j,t}$ is the arcsinh transform of the average number of Census block groups in county j devices from county i visit on day t. Column 1 of Table S7 repeats this analysis after restricting our sample to county dyads that are physically adjacent. Our results are qualitatively similar.

Table S2: The effect of policy interventions on county-level outcomes, including geo-spillovers.

	Dependent variable:			
	asinh(NSBTUs)	asinh(NHDF)	log(dCBGVs)	BTVRC
	(1)	(2)	(3)	(4)
Ego shelter-in-place	-2.08^{***}	-2.01^{***}	-4.08^{***}	-5.60^{***}
	(0.27)	(0.38)	(0.55)	(0.66)
Geo-alter shelter-in-place	-0.50	-2.76^{***}	-4.79^{***}	-4.09^{***}
	(0.34)	(0.40)	(0.70)	(0.79)
Ego close gyms/movies/restaurants	-0.77^{***}	-0.24	-0.76^{**}	-1.58^{***}
	(0.20)	(0.26)	(0.36)	(0.51)
Geo-alter close gyms/movies/restaurants	1.33^{***}	-0.70^{**}	-0.78	-0.20
	(0.26)	(0.34)	(0.62)	(0.74)
Conley s.e.	Yes	Yes	Yes	Yes
County fixed effect	Yes	Yes	Yes	Yes
Day fixed effect	Yes	Yes	Yes	Yes
Observations	122,598	122,598	122,598	122,598
R^2	0.43	0.42	0.40	0.38
Adjusted R^2	0.42	0.41	0.39	0.37
Residual Std. Error (df = 120044)	6.33	8.51	14.79	15.76

Notes: NSBTU refers to the fraction of Facebook users in a given county who visit multiple Bing tiles on a given day. NHDF refers to the fraction of devices that are not completely at home, as measured by Safegraph. dCBGVs refers to the mean number of Census block groups devices from a given county visit, as measured by Safegraph. BTVRC refers to the relative change in the number of Bing tiles users visit, as measured by Facebook. All values multiplied by 100. *p<0.1; **p<0.05; ***p<0.01

Table S3: The effect of policy interventions on county-level outcomes, including geo- and social-spillovers.

	Dependent variable:			
	asinh(SBTUs)	asinh(HDF)	log(dCBGVs)	BTVRC
	(1)	(2)	(3)	(4)
Ego shelter-in-place	-0.78^{*}	-1.26^{**}	-2.73^{**}	-1.86^{*}
	(0.46)	(0.57)	(1.31)	(1.02)
Social-alter shelter-in-place	-6.61^{***}	-3.98^{**}	-6.95^{*}	-18.82^{***}
	(1.33)	(1.69)	(3.80)	(2.86)
Geo-alter shelter-in-place	2.21^{***}	-1.08	-1.92	3.58^{**}
	(0.66)	(0.88)	(1.97)	(1.43)
Ego close gyms/movies/restaurants	-0.43	-0.42	-0.57	-0.29
	(0.30)	(0.29)	(0.70)	(0.61)
Social-alter close gyms/movies/restaurants	-1.55^{*}	0.89	-0.84	-5.99^{***}
	(0.80)	(0.93)	(2.34)	(1.46)
Geo-alter close gyms/movies/restaurants	1.79^{***}	-1.19^{*}	-0.63	1.76^{*}
	(0.51)	(0.65)	(1.47)	(1.06)
Conley s.e.	Yes	Yes	Yes	Yes
County fixed effect	Yes	Yes	Yes	Yes
Day fixed effect	Yes	Yes	Yes	Yes
Observations	122,598	122,598	122,598	122,598
R^2	0.44	0.43	0.40	0.39
Adjusted R^2	0.43	0.41	0.39	0.38
Residual Std. Error (df = 120042)	6.29	8.48	14.76	15.64

Notes: SBTU refers to the fraction of Facebook users in a given county who visit only one Bing tile on a given day. HDF refers to the fraction of devices completely at home, as measured by Safegraph. dCBGVs refers to the mean number of Census block groups devices from a given county visit, as measured by Safegraph. BTVRC refers to the relative change in the number of Bing tiles users visit, as measured by Facebook. All values multiplied by 100. *p<0.1; **p<0.05; ***p<0.01

Table S4: The effect of policy interventions on county-level outcomes, including in-state and out-of-state geo-spillovers.

	Dependent variable:			
	asinh(NSBTUs)	asinh(NHDF)	log(dCBGVs)	BTVRC
	(1)	(2)	(3)	(4)
Ego shelter-in-place	-1.92^{***}	-2.71^{***}	-3.67^{***}	-5.37^{***}
	(0.71)	(0.57)	(0.95)	(1.65)
In-state geo-alter shelter-in-place	-0.32	-0.75	-3.10^{***}	-2.50
	(0.62)	(0.52)	(0.97)	(1.61)
Out-of-state geo-alter shelter-in-place	-2.81^{***}	-3.93^{***}	-6.09^{***}	-7.80^{***}
	(0.67)	(0.90)	(1.36)	(1.33)
Ego close gym/movies/restaurants	-0.53	-0.86	-0.42	-1.72
	(0.69)	(0.53)	(0.86)	(1.60)
In-state geo-alter close gyms/movies/restaurants	0.57	0.26	-0.79	0.08
	(0.63)	(0.51)	(0.91)	(1.57)
Out-of-state geo-alter close gyms/movies/restaurants	1.02^{***}	-0.55	-1.69^{**}	-0.75
	(0.29)	(0.33)	(0.69)	(0.73)
Conley s.e.	Yes	Yes	Yes	Yes
County fixed effect	Yes	Yes	Yes	Yes
Day fixed effect	Yes	Yes	Yes	Yes
Observations	122,598	122,598	122,598	122,598
R^2	0.44	0.43	0.40	0.39
Adjusted R^2	0.43	0.41	0.39	0.38
Residual Std. Error (df = 120042)	6.29	8.49	14.78	15.69

Notes: SBTU refers to the fraction of Facebook users in a given county who visit only one Bing tile on a given day. HDF refers to the fraction of devices completely at home, as measured by Safegraph. dCBGVs refers to the mean number of Census block groups devices from a given county visit, as measured by Safegraph. BTVRC refers to the relative change in the number of Bing tiles users visit, as measured by Facebook. All values multiplied by 100.
*p<0.1; **p<0.05; ***p<0.01

Table S5: The effect of policy interventions on county-level outcomes, including in-state and out-of-state geo- and social-spillovers.

	Dependent variable:			
	asinh(NSBTUs)	asinh(NHDF)	log(dCBGVs)	BTVRC
	(1)	(2)	(3)	(4)
Ego shelter-in-place	-0.32	-1.77^{***}	-2.70^{**}	-1.62
	(0.45)	(0.66)	(1.23)	(1.37)
In-state geo-alter shelter-in-place	1.38^{**}	-0.20	-2.76^{**}	1.74
	(0.62)	(0.78)	(1.23)	(1.31)
Out-of-state geo-alter shelter-in-place	3.43^{***}	0.61	0.82	8.15^{***}
	(0.66)	(1.03)	(1.61)	(1.93)
In-state social-alter shelter-in-place	-3.25^{***}	-1.44	-1.26	-7.89^{***}
	(0.97)	(1.37)	(2.21)	(2.32)
Out-of-state social-alter shelter-in-place	-13.26^{***}	-10.14^{***}	-15.09^{***}	-33.19^{***}
	(1.00)	(1.32)	(2.06)	(2.37)
Ego close gym/movies/restaurants	0.61^{*}	-0.17	0.10	0.83
	(0.32)	(0.50)	(0.80)	(1.18)
In-state geo-alter close gyms/movies/restaurants	1.63^{***}	0.67^{*}	-0.92	2.43^{***}
	(0.29)	(0.41)	(0.71)	(0.69)
Out-of-state geo-alter close gyms/movies/restaurants	1.74^{***}	-1.70^{***}	-2.37^{**}	3.55^{***}
	(0.40)	(0.52)	(0.99)	(1.17)
In-state social-alter close gyms/movies/restaurants	-2.29^{***}	-1.19	-0.53	-5.09^{***}
	(0.51)	(0.82)	(1.15)	(1.47)
Out-of-state social-alter close gyms/movies/restaurants	-2.26^{***}	1.19^{*}	-0.01	-9.96^{***}
	(0.50)	(0.70)	(1.11)	(1.32)
Conley s.e.	Yes	Yes	Yes	Yes
County fixed effect	Yes	Yes	Yes	Yes
Day fixed effect	Yes	Yes	Yes	Yes
Observations	122,598	122,598	122,598	122,598
R^2	0.46	0.44	0.41	0.41
Adjusted R^2	0.45	0.43	0.39	0.40
Residual Std. Error (df = 120038)	6.16	8.40	14.69	15.46

Notes: SBTU refers to the fraction of Facebook users in a given county who visit only one Bing tile on a given day. HDF refers to the fraction of devices completely at home, as measured by Safegraph. dCBGVs refers to the mean number of Census block groups devices from a given county visit, as measured by Safegraph. BTVRC refers to the relative change in the number of Bing tiles users visit, as measured by Facebook. All values multiplied by 100.
*p<0.1; **p<0.05; ***p<0.01

Table S6: The effect of policy interventions on dyadic movement patterns (all connected counties)

	Dependent variable:		
	asinh(dCBGs/device)	log(dCBGs/device + 1)	dCBGs/device
	(1)	(2)	(3)
Or. = 0, Dest. = 1 (G/M/R)	-5.38^{***}	-5.21^{***}	-5.35^{***}
	(1.71)	(1.62)	(1.73)
Or. = 1, Dest. = 0 (G/M/R)	-0.04	-0.32	0.08
	(3.21)	(2.97)	(3.27)
Or. = 1, Dest. = 1 (G/M/R)	-3.43^{***}	-3.34^{***}	-3.37^{***}
	(1.11)	(1.10)	(1.13)
Or. = 0, Dest. = 1 (shelter-in-place)	0.73	0.41	0.85
	(2.99)	(2.70)	(3.05)
Or. = 1, Dest. = 0 (shelter-in-place)	15.03^{***}	13.64^{***}	15.38^{***}
	(2.58)	(2.35)	(2.67)
Or. = 1, Dest. = 1 (shelter-in-place)	-6.23^{***}	-6.44^{***}	-5.96^{***}
	(1.68)	(1.58)	(1.75)
Clustered s.e.	Yes	Yes	Yes
Dyad fixed effect	Yes	Yes	Yes
Day fixed effect	Yes	Yes	Yes
Observations	26,874,050	26,874,050	26,874,050
R^2	0.89	0.90	0.89
Adjusted R^2	0.89	0.90	0.88
Residual Std. Error (df = 26325546)	0.36	0.31	0.37

Notes: Or. refers to the origin county, Dest. refers to the destination county. 1 indicates that a given policy is in place, whereas 0 indicates it is not. G/M/R refers to closing gyms, movie theaters, and/or restaurants, and shelter-in-place refers to a shelter in place order. dCBGs/device refers to the average number of Census block groups in the destination county that devices from the origin county visited on a given day. All coefficients and standard errors are multiplied by 10,000. *p<0.1; **p<0.05; ***p<0.01

Table S7: The effect of policy interventions on dyadic movement patterns (adjacent counties)

	Dependent variable:		
	asinh(dCBGs/device)	log(dCBGs/device + 1)	dCBGs/device
	(1)	(2)	(3)
Or. = 0, Dest. = 1 (G/M/R)	−3.60	−2.90	−3.84
	(2.67)	(2.13)	(2.89)
Or. = 1, Dest. = 0 (G/M/R)	0.42	0.33	0.48
	(1.34)	(1.15)	(1.40)
Or. = 1, Dest. = 1 (G/M/R)	−0.98*	−0.91*	−0.94*
	(0.50)	(0.47)	(0.51)
Or. = 0, Dest. = 1 (shelter-in-place)	−11.63***	−9.91***	−11.99***
	(3.95)	(3.31)	(4.10)
Or. = 1, Dest. = 0 (shelter-in-place)	5.51**	4.25**	5.97**
	(2.74)	(2.13)	(3.01)
Or. = 1, Dest. = 1 (shelter-in-place)	−5.13***	−4.97***	−5.00***
	(1.20)	(1.08)	(1.23)
Clustered s.e.	Yes	Yes	Yes
Dyad fixed effect	Yes	Yes	Yes
Day fixed effect	Yes	Yes	Yes
Observations	878,472	878,472	878,472
R^2	0.89	0.90	0.88
Adjusted R^2	0.89	0.90	0.88
Residual Std. Error (df = 860490)	1.67	1.41	1.74

Notes: Or. refers to the origin county, Dest. refers to the destination county. 1 indicates that a given policy is in place, whereas 0 indicates it is not. G/M/R refers to closing gyms, movie theaters, and/or restaurants, and shelter-in-place refers to a shelter in place order. dCBGs/device refers to the average number of Census block groups in the destination county that devices from the origin county visited on a given day. All coefficients and standard errors are multiplied by 1,000. *p<0.1; **p<0.05; ***p<0.01

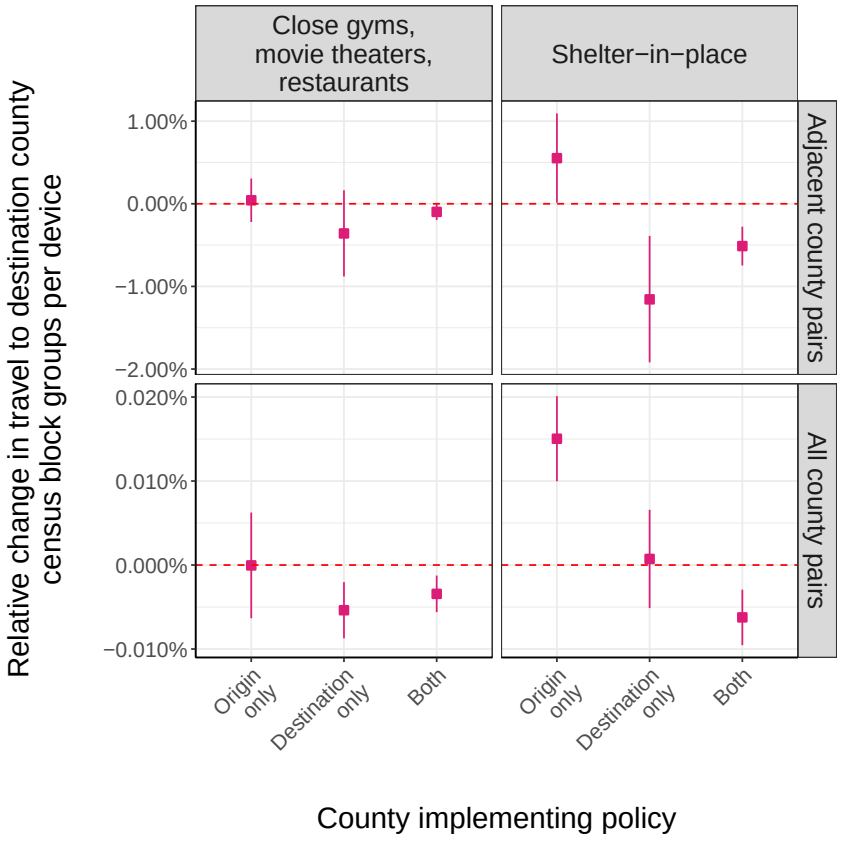

Figure S18: Difference-in-differences analysis of the geographic spillovers arising from different dyad-level policy vectors in response to COVID-19.

S2.6 Robustness Checks

S2.6.1 Leading and Lagging Policy Effects

As a supplement to the parallel pre-trends plot of Fig. 1b in the main text, we consider other evidence to test the parallel trends assumption for difference-in-differences that counties with different policy introduction timing have parallel potential outcomes (*12, 13*).

In particular, given that one may reasonably posit that, e.g., shelter-in-place policies may be anticipated by the public, such that some effects could precede the policy introduction, we analyze leads and lags of the policy variable as a test of anticipation effects and other dynamics. Figures S19, S20 and S21 display marginal effects of leads and lags of the policies we analyze.

Note that each estimate is of the marginal effect of the leaded (or lagged) policy, so the increases after the introduction of the policy need not indicate a reversal of effects.

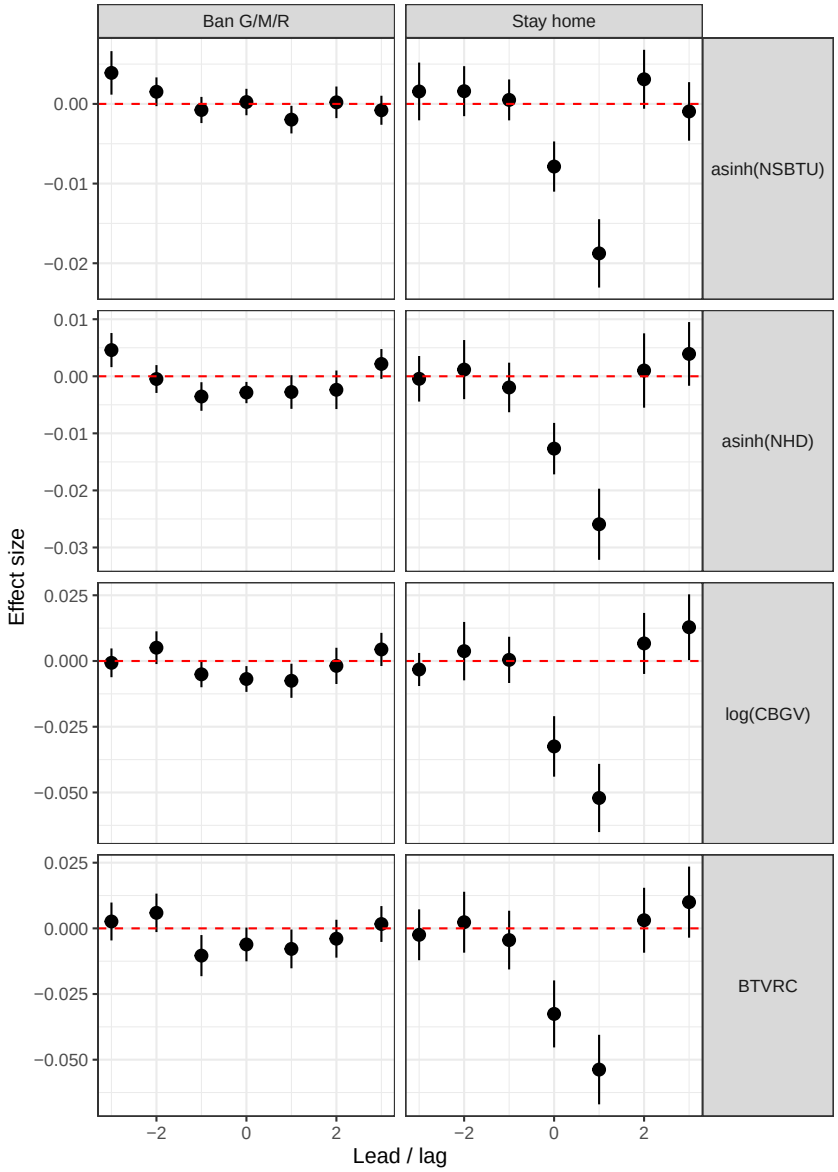

Figure S19: The leads and lags of the own-state policy decisions to close gyms, movie theaters, and restaurants and implement shelter-in-place. For shelter-in-place / "stay home" policies, there is no evidence of anticipatory effects. Note that each estimate is of the marginal effect of that lead or lag policy, so an estimate of zero after the policy corresponds to a constant effect, not a reversal.

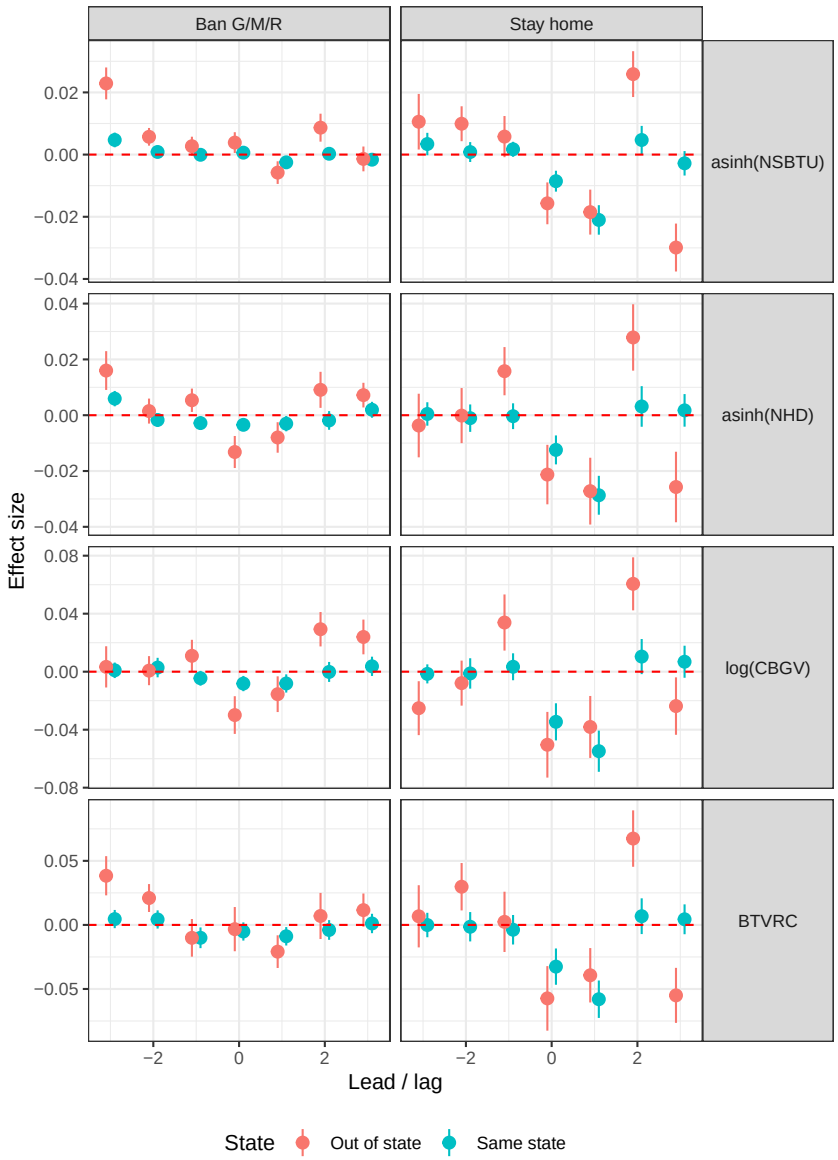

Figure S20: The leads and lags of the own-state, as well as alter-state policy decisions to close gyms, movie theaters, and restaurants and implement shelter-in-place / "stay home" (geographic alters only). Note that each estimate is of the marginal effect of that lead or lag policy, so an estimate of zero after the policy corresponds to a constant effect, not a reversal.

S2.6.2 Randomization Inference

As a robustness check, we use a Fisherian randomization inference (FRI) procedure to estimate the null distributions of the spillover effects (*14–17*). In a simple experimental context, this is

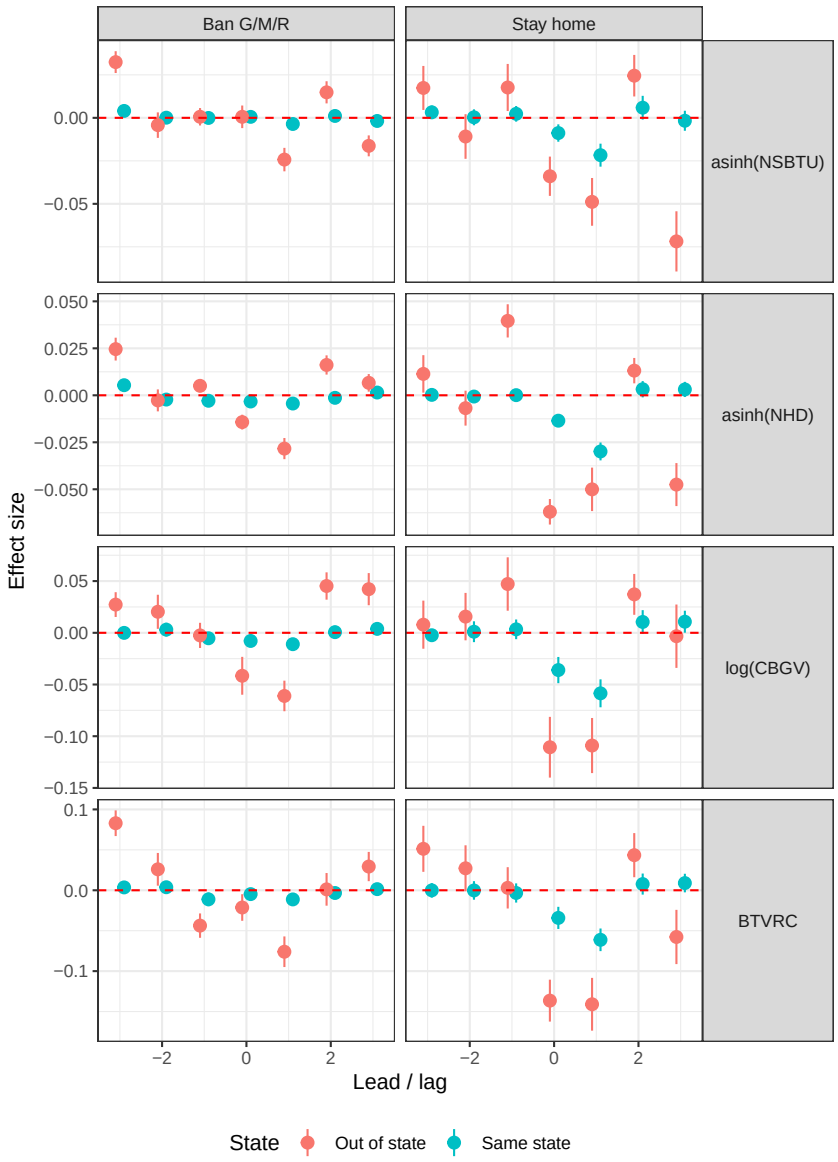

Figure S21: The leads and lags of the own-state, as well as alter-state policy decisions to close gyms, movie theaters, and restaurants and implement shelter-in-place / "stay home" (geographic and social alters). Note that each estimate is of the marginal effect of that lead or lag policy, so an estimate of zero after the policy corresponds to a constant effect, not a reversal.

accomplished by re-drawing (typically, permuting) the treatment assignment vector R times and then estimating the treatment effect under the permuted treatment assignment vectors. However, in our context, there is an important structure to the policy assignment vectors that our re-

sampling procedure needs to account for. Specifically, the main concern is that county-level policy is often set at the state level. though there are some exceptions. Furthermore, we need to make sure that once a county is "treated" it remains "treated."

We record the the fraction of counties that have implemented a gym/movie theater/restaurant ban or a shelter-in-place order for each state–date pair. To create a re-sampled policy draw, we start by randomly assigning each state a "replacement" state whose policy assignments take the place of the "original" state. For example, suppose that Oklahoma (a state without a shelter-in-place order) was assigned the replacement state of California. The re-sampled data would then include a policy vector for Oklahoma in which a shelter-in-place order was applied across all counties starting on March 19, 2020, the date of California's shelter-in-place order.

Since there is some deviation between county-level policy and state-level policy, and the number of counties varies from state to state, we ensure that the fraction of "treated" counties in the original state approximately matches the expected fraction of "treated" counties in the replacement state on each day. Specifically, for each original state, we randomly create an order in which that state's counties implement a policy. This is done to insure consistency in policy assignment at the county level (i.e. if county 10 is treated on date 1, then we would want it to stay treated since our data does not cover lifting of policies). For example, suppose that a state has 41 counties and the replacement state has 50% of its counties treated on date 1 and all of its counties treated on date 2. In this case, our procedure would have the first 50% (rounded up to the nearest whole county) of counties in the original state treated on date 1 according to the randomly drawn order, and 100% of counties treated on date 2.

We repeat this policy replacement procedure for every single state (without replacement, so replacement states are only used once each draw) to form a single re-sampled policy draw. We then perform the DML residualization procedure on the the re-sampled policy vector and then re-estimate Equation 4. We then run a Wald test of the full unrestricted model against re-

stricted models where the social and geographic spillovers sum to 0 for each policy and record the resulting state-cluster-robust F-statistics. This test statistic is selected based on the expectation that it will be approximately pivotal and result in valid tests of a weak (non-sharp) null hypothesis (*16, 18, 19*).

Overall, we repeat the policy resampling procedure $R = 1{,}000$ times to form the Fisherian null distributions. Fig. S22 shows the null distribution of spillover F-statistics for our model specification obtained through random permutation, as well as our observed F-statistic values, for both closing gyms/movies theaters/restaurants and implementing shelter-in-place for the Safegraph outcomes asinh(NDHF). The p-values for spillovers from closing gyms, movie theaters, and restaurants are 0.763 (for leaving home, asinh(NDHF)) and 0.706 (for locations visited, log(dCBGVs)), whereas the p-values for spillovers from implementing shelter-in-place are $< .001$. Fig. S23 shows the policy spillover F-statistic null distributions and observed F-statistic values for the Facebook outcomes. We find that the p-values for spillovers from closing gyms, movie theaters, and restaurants are 0.346 (for leaving home, asinh(NSTUs)) and $< .001$ (for locations visited, BTVRC), whereas the p-values for spillovers from implementing shelter-in-place are $< .001$. These results are generally qualitatively consistent with the results presented in Table S5 using the adjacency and clustering-robust standard errors.

S3 Identifying Endogenous Social Effects

In prior sections, we had primarily focused on measuring the effect of alter county policy spillovers on a focal county's social distancing behavior. In this section we focus directly on measuring the endogenous peer (or social) effects in social distancing behavior amongst peers. However, inference for endogenous peer effects must overcome issues of simultaneity (i.e., reflection (*20*)), confounding with homophily, and other processes that produce correlations in

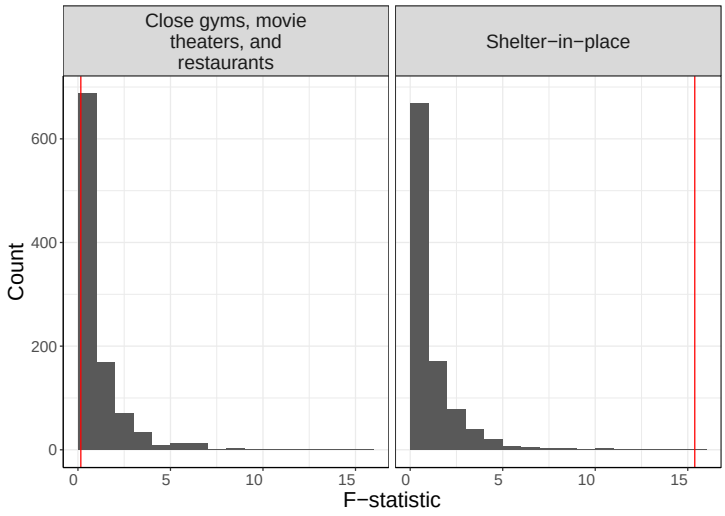

(a) asinh(NHD)

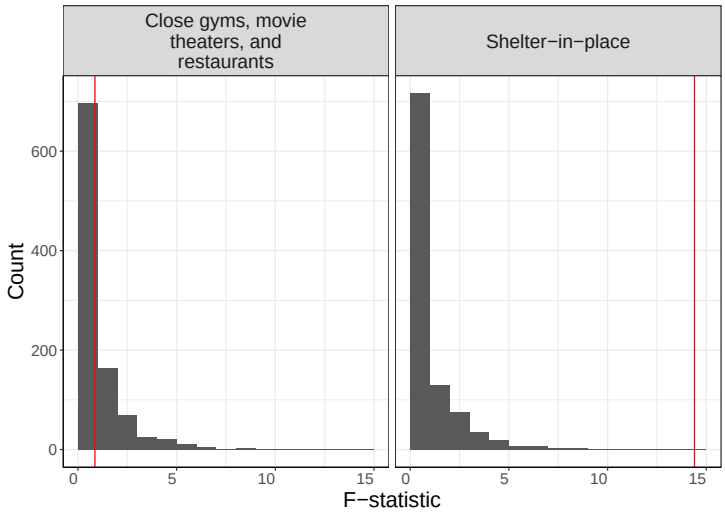

(b) log(dCBGVs)

Figure S22: Null distribution (grey) and observed value (red line) of the test statistic (F-statistic) for spillovers of alter county policies on the Safegraph outcomes.

behaviors among connected units. We address this problem through an instrumental variable[7] (IV) approach that leverages fluctuations in peer weather and geographic variation in shares of industries that are deferentially affected during this period as shocks to peer behavior. The

[7]As long as the standard IV assumptions (relevancy and exclusion) are met, (*21, 22*) prove that IV approaches are able to consistently estimate endogenous peer effects.

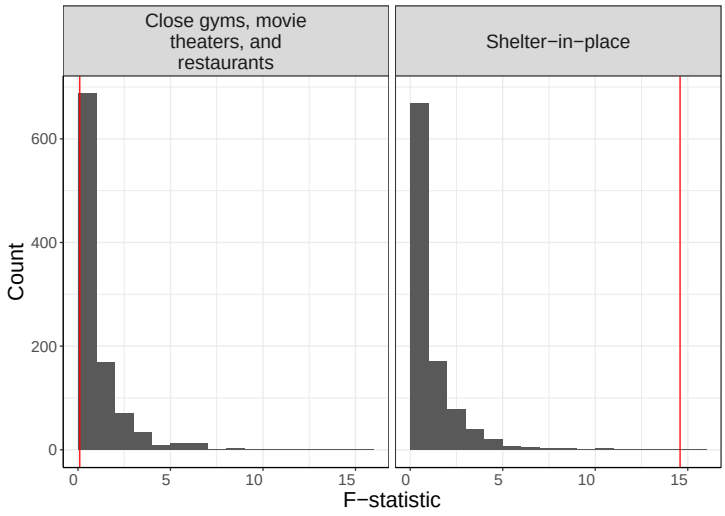

(a) asinh(NSTU)

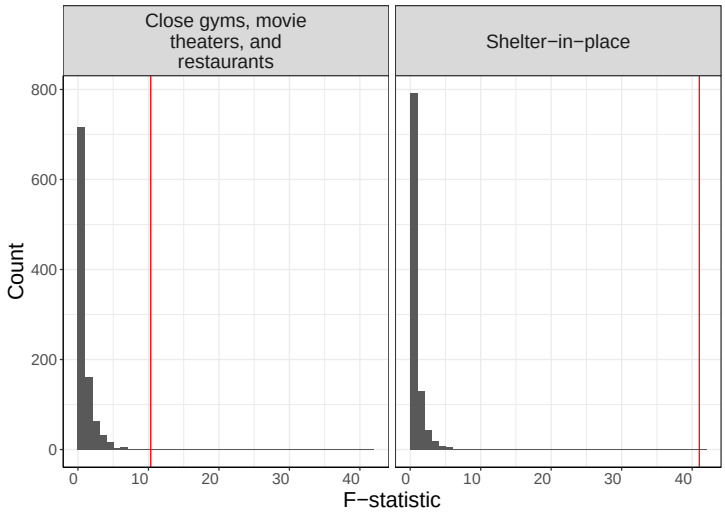

(b) BTVRC

Figure S23: Null distribution (grey) and observed value (red line) of the test statistic (F-statistic) for spillovers of alter county policies on the Facebook outcomes.

rest of this section is organized as follows: Sections S3.1 and S3.2 describe and motivate our 2 sets of instruments, Section S3.3 details our model specifications, Section S3.4 reports our estimation results, and Section S3.5 presents a series of robustness checks.

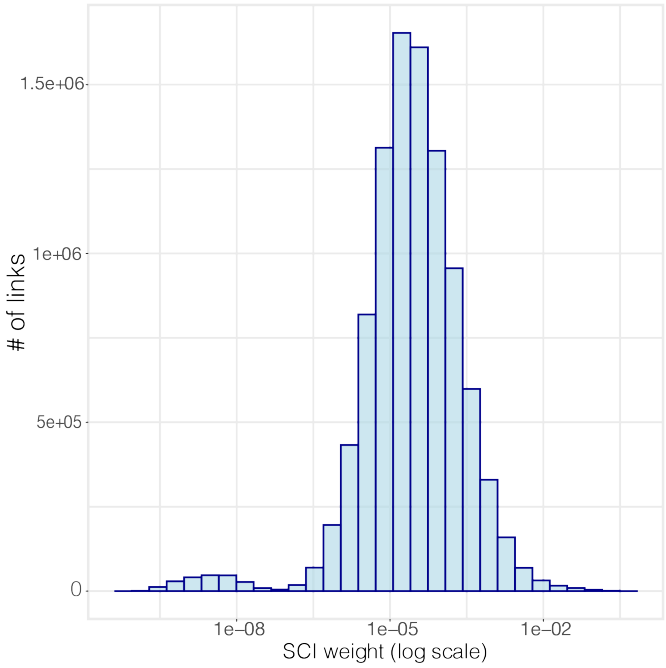

Figure S24: The distribution of the SCI weights between counties.

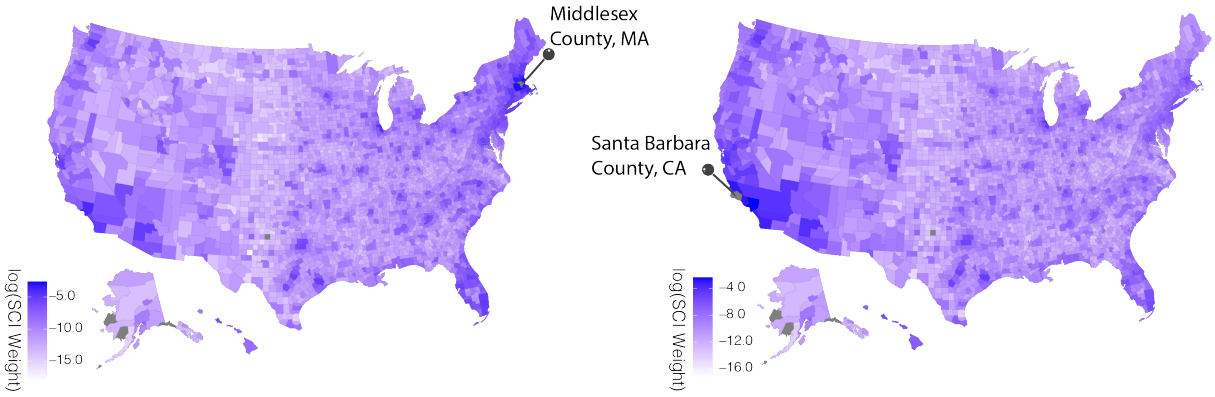

Figure S25: The spatial distribution of the weighted social graph of Middlesex County, MA (left) and Santa Barbara County, CA (Right) as extracted by the SCI.

S3.1 Weather Instruments

One set of the instrumental variables we use is constructed from the weather,[8] which has been used to identify social spillovers across a variety of different contexts (*23–25*). The intuition

[8]These instruments are built from the county-level precipitation and max temperature measures computed via a procedure described in Section S1.3. The exact details on how the specific weather instruments are constructed are detailed in Section S3.3.2 below.

is that while the variations in my weather will have a meaningful impact on my behavior, they shouldn't directly affect the behavior of a friend who lives in a different area. This relationship between weather and movement seems to be true for social distancing behaviors: even if a shelter-in-place order is in effect, there are considerably more people out and about when the weather is nice. We observe this effect in our data as shown in Fig. 1c in the main text of the paper.

As such, we should be able to use peer weather to instrument for peer behavior, thereby allowing us obtain a consistent estimate of endogenous peer effects. However, one major concern with weather instruments is the fact that geographically proximate locations tend to have similar weather. Theoretically, this should not create an identification issue: even if "peer weather" is highly correlated with "own weather," it should be conditionally ignorable so long as the effects of "own weather" are properly accounted for. However, properly accounting for "own weather" can be difficult as the impact of weather can be very nonlinear. For example, going from 0mm to 1mm of rain is going to elicit a much larger behavioral response than going from 20mm to 21mm. Furthermore, there are likely significant interactions between different weather measures to consider: the effect of rain when it is a comfortable 22C outside is going to be very different than the effect of rain when it is a frigid 2C outside. These complexities may result in a "practical" violation of conditional ignorability since peer weather may be providing additional information about own weather that is not fully captured by a linear model specification. To address this potential concern, we adopt a modeling approach that can flexibly account for these issues that we detail in Section S3.3.1.

S3.2 Industry Visits Shift-Share Instrument

Our second set of instruments are Bartik-like (*26*) industry shift-share variables that measure the industry visit composition of alter counties[9]. At a high level, these instruments typically average national employment growth across industries using local industry employment shares as weights to produce a measure of local labor demand that is unrelated to changes in local labor supply. Bartik popularized the application of the variable of interest in empirical regression analyses as a shift-share, whereby the variable of interest comprises a regression of related variables weighted by their relative importance (where, in Bartik's case, labor market outcomes were predicted based on regional sectoral shocks).

Industries have been impacted differently by governmental policies depending on whether they are deemed essential or non-essential. Moreover, some industries have more jobs that can be done at home than other industries (*27*). Industries with more jobs that can be done at home are more likely to close their physical locations and ask their employees to work from home either on their own volition or due to government regulations. Furthermore, counties have different industry compositions. Fig. S26 plots the total visits in the Safegraph data in the US for a sample of 8 industries. There is substantial heterogeneity in changes in number of visits across industries. The exact construction of these variables is described above in Section S1.1.2.

S3.3 Model Specification

We estimate the following set of semi-parametric model specifications:

$$Y_{it} = Y_{-it}\beta + D_{it}\delta_1 + D_{-it}^{geo}\delta_2 + D_{-it}^{social}\delta_3 + S_{it}\psi + f(W_{it}) + \alpha_i + \tau_t + \epsilon_{it} \tag{S7}$$

$$Y_{-it} = D_{it}\gamma_1 + D_{-it}^{geo}\gamma_2 + D_{-it}^{social}\gamma_3 + S_{it}\pi + g(W_{it}) + h(W_{-it}, S_{-it}, D_{-it}^{social}) + \alpha_{-i} + \tau_t + \nu_{-it} \tag{S8}$$

[9]These instruments are constructed from the Safegraph visits data, in a process described in Section S1.1.2 above.

Figure S26: Log(total visits) to a sample of 8 industries from March 1, 2020 to April 18, 2020 across the U.S.

Equation 7 refers to our main model specification. Our dependent variable Y_{it} denotes the social distancing behavior of individuals in county i on date t (as measured by our 4 main social distance outcomes: asinh(NSBTUs), asinh(NHDF), log(dCBGVs), and BTVRC). Y_{-it}, denotes the weighted average of social distancing behaviors of individuals in other counties where weights are determined by the social connectedness between county i and other counties j. More formally: $Y_{-it} = \sum_j w_{ij} * Y_{jt}$ where $w_{ij} = \frac{n_j * a_{ij}^{social}}{\sum_k n_k * a_{ik}^{social}} : k \neq i$ and n_j is the population

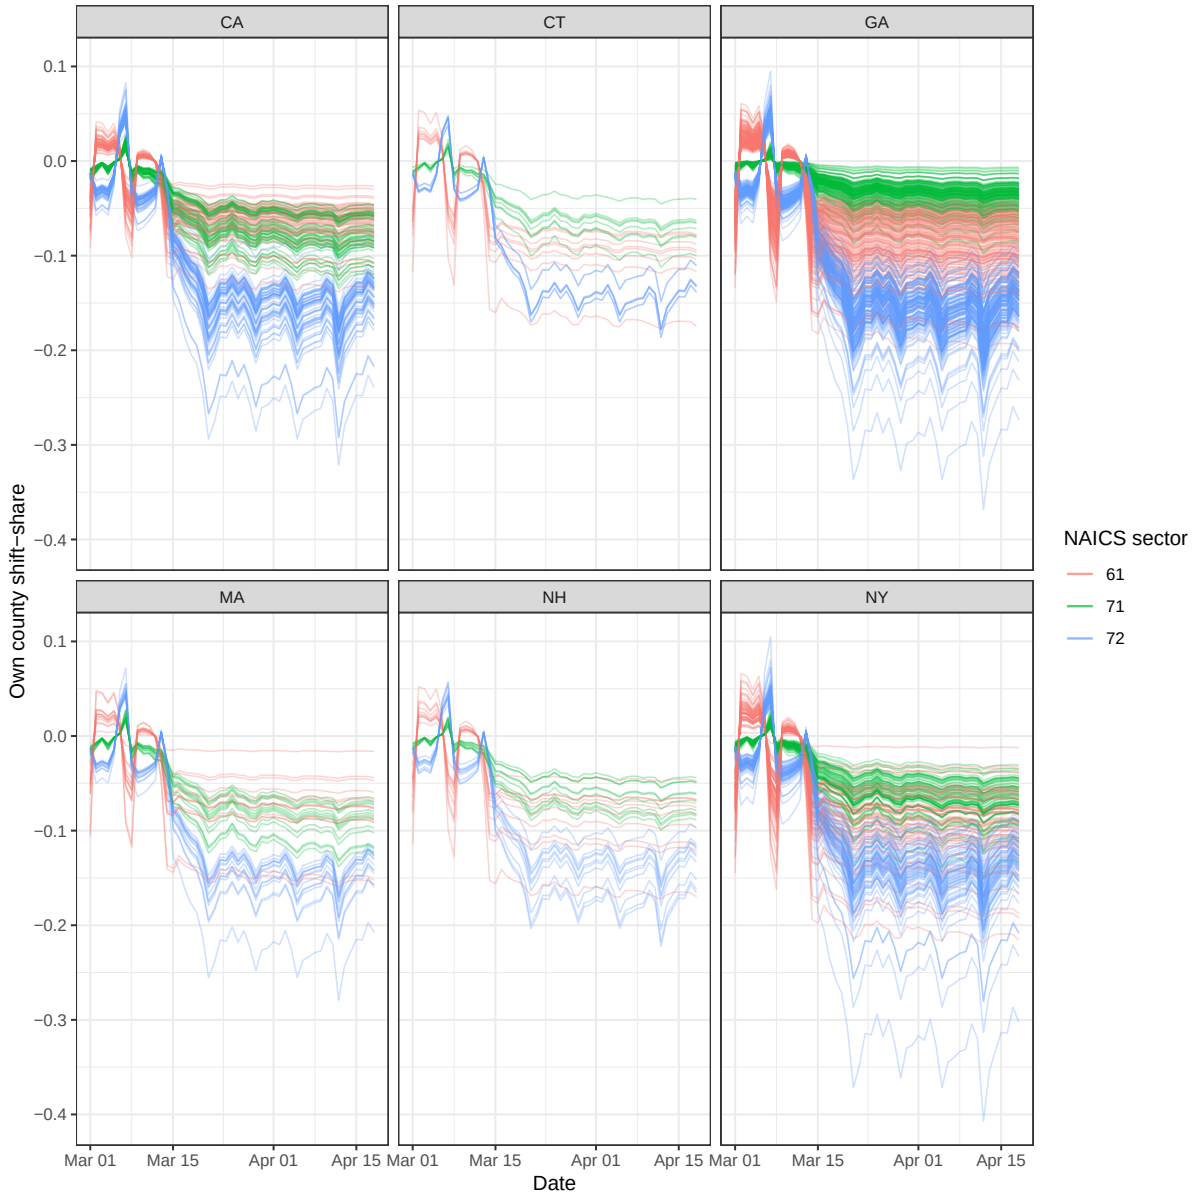

Figure S27: Illustration of county-level shift-shares for three NAICS sectors — 61 (Educational Services), 71 (Arts, Entertainment, and Recreation), and 72 (Accommodation and Food Services) — for six states. For a given NAICS sector, the curves for different counties differ only in scale due to baseline differences in visits to those locations, as all are multiplied by the common, national changes in visitation.

of county j[10]. β is the main parameter of interest, representing the endogenous peer effect of

[10]Since 3 of our outcomes have been nonlinearly transformed, we first compute the weighted average of the untransformed variable and then apply the appropriate transformation on the weighted average.

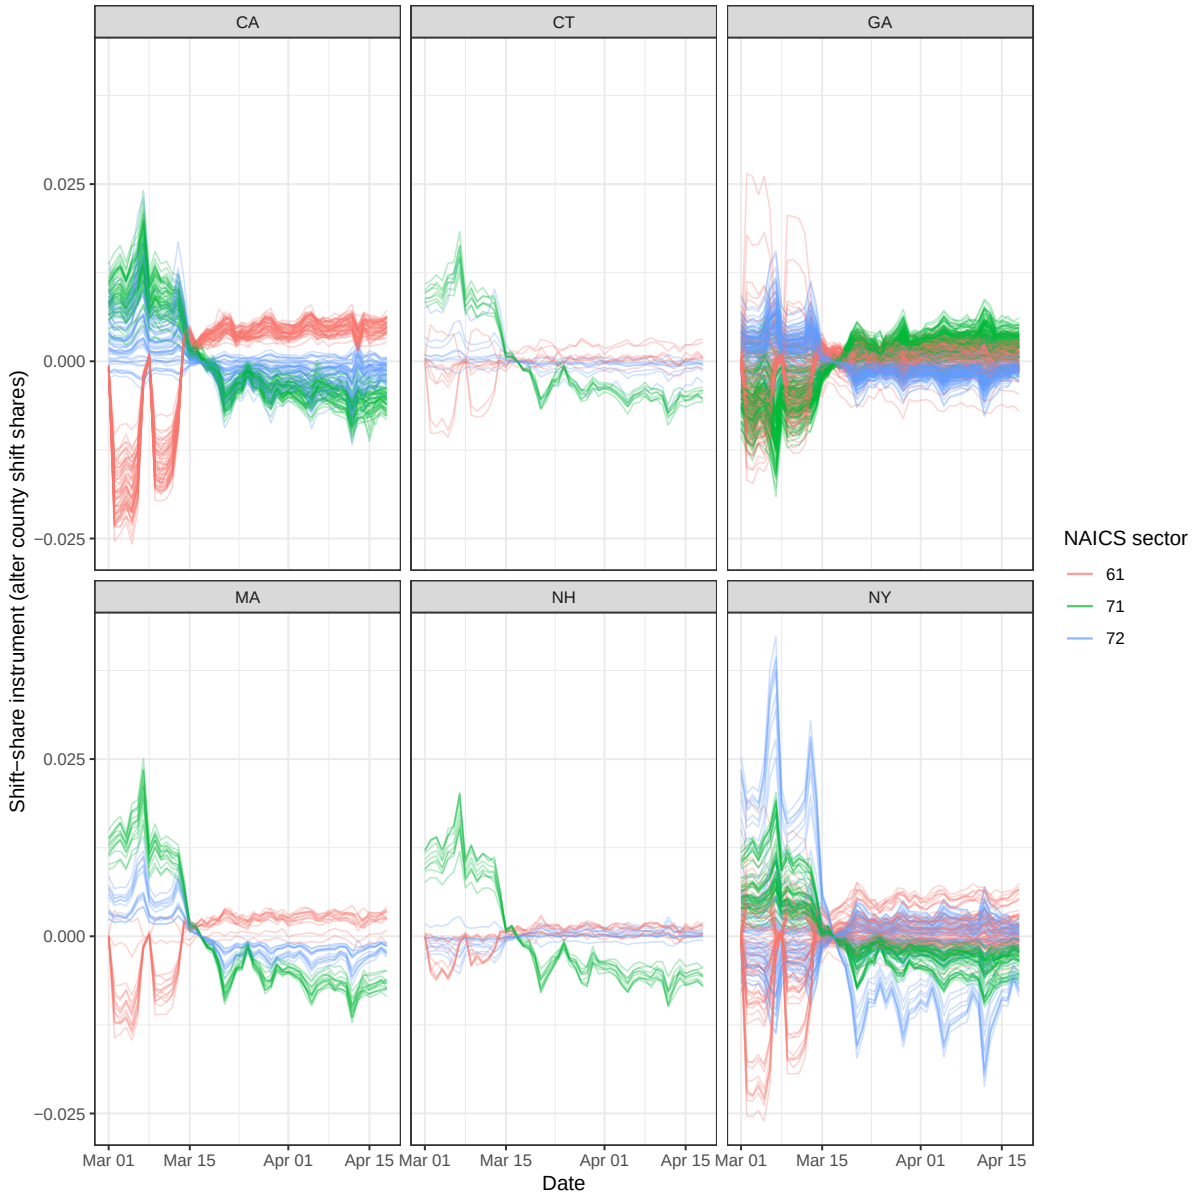

Figure S28: Illustration of shift-share instruments for three NAICS sectors — 61 (Educational Services), 71 (Arts, Entertainment, and Recreation), and 72 (Accommodation and Food Services) — for six states. For a given NAICS sector, the curve for a focal county is a weighted combination of the shift-shares of its alter counties as determined by SCI.

social distancing behavior. D_{it} denotes whether any county-level[11] social distancing policies

[11] State-wide policies are automatically propagated down to the county level.

are in effect in county i on date t, while δ_1 captures the effects of policy on social distancing behavior. Both D_{-it}^{social} and D_{-it}^{geo} are weighted averages of other counties' social distancing policies on date t, where the weights used to construct D_{-it}^{social} are identical to the weights used to construct Y_{-it} while the weights used to construct D_{-it}^{geo} are constructed using the pre-crisis county-to-county movement as discussed in Section S1.1. The associated parameters, δ_2 and δ_3 capture the potential spillover effects of other counties' policies that may not be mediated by endogenous peer behavior. As mentioned above, each of D_{it}, D_{it}^{geo}, and D_{it}^{social} are 2-dimensional row vectors to both account for both gym, movies, and restaurant bans and shelter-in-place orders. Likewise, δ_1, δ_2, and δ_3 are 2-dimensional column vectors. S_{it} is a

W_{it} reflects the weather (measured as precipitation and maximum temperature) in county i and date t and $f(.)$ is a non parametric function that can capture nonlinear responses of Y_{it} to fluctuations in W_{it}. We will more formally describe how $f(.)$ is estimated below in Section S3.3.1. α_i and τ_t denote a set of county and date fixed effects respectively. Lastly, the ϵ_{it} represents the error term.

In the first-stage (Equation 8), Y_{-it} takes its place as the dependent variable. γ_1, γ_2, γ_3, and $g(.)$[12] capture the effects D_{it}, D_{it}^{geo}, and W_{it} on Y_{-it} respectively. W_{-it} and I_{-it} are weighted averages of weather and industry shares of other counties (not i), again weighted by social connectedness. These variables and their interactions with D_{-it} will form the set of potential candidates for our instruments. The associated function $h(.)$ will select a much smaller subset that we will actually use in our estimation. We will describe this selection procedure in Section S3.3.2 below. Again, α_i and τ_t represent fixed effects for county and day, and ν_{-it} denotes the first stage error term.

[12]Like $f(.)$, $g(.)$ is a function that can capture the nonlinear effects of W_{it}.

S3.3.1 Flexibly Controlling for the Effect of Weather

To flexibly control for the impact of weather, we employ a "double machine learning" (DML) procedure (*28*). This approach is designed to estimate and draw inferences on a low-dimensional parameter in the presence of high-dimensional nuisance parameters. Consider the following "canonical example" from (*28*) which we reproduce here:

$$Y = D\theta_0 + g_0(Z) + U, \quad \mathbb{E}[U|D, Z] = 0$$
$$D = m_0(Z) + V, \quad \mathbb{E}[V|Z] = 0$$

Y denotes the outcome, D is a policy or treatment variable, θ_0 is the low-dimensional parameter of interest, Z is a high-dimensional vector of covariates ($g_0(Z)$ can be considered to be the high-dimensional nuisance parameter), and U and V are the errors. The basic intuition behind DML is that $g_0(.)$ and $m_0(.)$ can be estimated using non-parametric statistical methods (aka machine learning) and then "partialed out" (*29*) from both Y and D. Then one simply regresses the residuals of the dependent variable on the residuals of the treatment variable in order to estimate θ_0. In order to provide guarantees that key moment conditions are satisfied, the machine learning predictions needs to be orthogonalized which can be achieved via sample splitting. As such, the general double ML algorithm is as follows:

1. Split the dataset into K equal size partitions or "folds." Let $F_k, F_k^c : k \in 1, ..., K$ denote each fold and its complement.[13]

2. Estimate g_0 and m_0 with some non-parametric statistical model of choice using only the observations in B_1^c

3. Form residuals $\tilde{Y} := Y - \hat{g}_0(Z)$ and $\tilde{D} := D - \hat{m}_0(Z)$ only on observations in F_1.

[13]Suppose a dataset has 100 observations and is split into 5 block. B_1 consists of observations 1-20 and $F_1^c := F_2, F_3, F_4, F_5$ consists of the remaining observations 21-80.

4. Regress $\tilde{Y}$ on $\tilde{D}$ to obtain an estimate of θ_0. Overall, this estimate can be thought of as function of F_1 and F_1^c: $\hat{\theta}_0(F_1, F_1^c)$.

5. Repeat steps 2-4 for the the remaining $K-1$ folds

6. Form the final estimate of θ_0 by averaging across all estimates: $\hat{\theta}_0^* = \frac{1}{K}\sum_k \hat{\theta}_0(F_k, F_k^c)$

In our case, we consider a county's own weather to be the high-dimensional nuisance parameter, as we are not principally interested in identifying the effect of own weather on social distancing behavior. We use gradient boosted decision trees via XGBoost (*30*), a state-of-art machine learning algorithm, to estimate $f(.)$ in Equation 7, $g(.)$ in Equation 8, as well as the effect of weather on any of the other variables included in our models. XGBoost is an ensemble method that works by fitting a series of forward stage-wise decision trees aimed to minimizing a specified loss function. To give a general idea of the basic procedure:

1. Fit an initial decision tree T_1 that minimizes $E\left[(Y - T_1(X))^2\right]$, where Y is the outcome and X are the covariates or features.

2. Each successive tree is then fitted on the residuals of the previous state[14]:
   $T_n = \arg\min_T \mathbb{E}[(Y - \sum_{n-1}^{i=1} T_i(X) - T(X))^2]$

In order to prevent overfitting, this iterative process is stopped once out-of-sample predictive performance starts to decline.

As with many other machine learning algorithms, there are a number of hyperparameters that control this estimation procedure of XGBoost. In particular, we adjust:

- `tree_depth`: Controls the depth that each tree-based model is allowed to grow to. The deeper the tree, the more complex the model.

[14]To be more precise, the degree to which each successive tree contributes to the ensemble can be controlled via tuning hyperparameter called a learning rate. We provide a little bit more detail on this below.

- `eta`: Controls the "learning rate" or step size of each model. One way to think of this parameter is as a form of regularization on each model step in order to prevent overfitting.

- `nrounds`: The maximum number of stages the fitting process is allowed to continue on for.

We fix `tree_depth` to 2 and `eta` = 0.5, but allow `nrounds` to run up to a maximum of 100. Then, for each individual variable, the optimal number of rounds (given our choice of `tree_depth` and `eta`) is determined via a cross-validation procedure for each variable individually[15]. Once the optimal `nrounds` is determined, we form the residuals for all our dependent variables and covariates by first partialing out the set of fixed effects and then following the DML approach described above. This then leads to the following set of partialed out specifications:

$$\tilde{Y}_{it} = \tilde{Y}_{-it}\beta + \tilde{D}_{it}\delta_1 + \tilde{D}^{geo}_{-it}\delta_2 + \tilde{D}^{social}_{-it}\delta_3 + \tilde{S}_{it}\psi + \epsilon_{it} \tag{S9}$$

$$\tilde{Y}_{-it} = \tilde{D}_{it}\gamma_1 + \tilde{D}^{geo}_{-it}\gamma_2 + \tilde{D}^{social}_{-it}\gamma_3 + \tilde{S}_{it}\pi h(\tilde{W}_{-it}, \tilde{S}_{-it}, \tilde{D}^{social}_{-it}) + \nu_{-it} \tag{S10}$$

S3.3.2 Sparse Instrumental Variable Selection Procedure

We use sparse linear model (*31, 32*) to estimate $h(.)$. In order to try to maximize the power of our first stage, we opt to include a large set of instruments formed by interacting W_{-it}, S_{-it}, and D^{social}_{-it}.

To allow for nonlinearity in the impact of weather, we first construct a sequence of county-level indicator variables that take a value of 1 if the amount of rainfall in county i on date t falls within or exceeds a specific precipitation ventile [16] (20-quantiles). We generate a similar

[15]We note that it would be more optimal to do an exhaustive grid search across the entire hyperparameter space for each individual variable that needs to have the effect of weather partialed out. However, such a grid search would be extremely computationally expensive and would only yield very minor improvements in predictive accuracy.

[16]For example $V^{\text{prcp},1}_{it} = \mathbf{1}(\text{prcp}_{it} > q) : q = \arg_x Pr(\text{prcp}_{it} \geq x) = 0$, $V^{\text{prcp},2}_{it} = \mathbf{1}(\text{prcp}_{it} \geq q) : q =$

sequence of indicator variables for maximum temperature as well. To avoid perfect multicolinearity, we remove the first max temperature ventile and the first 9 precipitation ventiles since these indicators always took a value of 1[17]. The remaining county-level level covariates are then multiplied by the social-weighting matrix[18] to form the peer weather covariates ($W_{-it} = V_{-it}^{\text{prcp},10}, ..., V_{-it}^{\text{prcp},20}, V_{-it}^{\text{prcp},2}, ..., V_{-it}^{\text{prcp},20}$).

These peer weather covariates are then interacted with the entire set of peer shift-share variables S_{-it} (also constructed by multiplying the county-level shift-shares with the social-weighting matrix) as well as D_{-it}^{social} (specifically only the shelter-in-place component, since there's relatively less variance in the gyms, movies, and restaurants ban) to form a set of 1610 potential instruments. We partialed out all the fixed effects and then used the DML procedure described above to remove the effect of "own" county weather. This large number of possibly weak instruments could produce an highly biased or asymptotically inconsistent estimator (*33–35*). To address this issue and to systematically select instruments, we use a post-LASSO IV procedure (*32*), in which an L1-penalized regression (*31*) is first used to select the set of variables and then the selected set are used in standard unregularized model. We select the largest penalty λ that is within within one standard error of the minimum MSE in 3-fold cross-validation. This procedure selected a total of 377 instruments that we use in our primary estimation results.

S3.4 Results

The primary results of this instrumental variables analysis are shown in Table S8. Across all four of our main outcomes, we find extremely strong and highly consistent evidence of ex-

$\arg_x Pr(\text{prcp}_{it} \geq x) = 0.5$, etc. It is also worth noting that this construction means that if $V_{it}^{\text{prcp},k} = 1$, then $V_{it}^{\text{prcp},j} = 1 : j < k$.

[17] In the case of the first ventile, all values will be greater than or equal to the minimum value, meaning that the first indicators will always take a value of 1 based on how they are constructed. In the case of the other precipitation ventiles, there is a large mass at 0 since rain doesn't occur that frequently in most counties.

[18] More formally: $V_{-it}^{\text{prcp},k} = \sum_j w_i j * V_{jt}^{\text{prcp},k}$ and $V_{-it}^{\text{tmax},k} = \sum_j w_i j * V_{jt}^{\text{tmax},k}$.

tremely strong positive endogenous social effects in social distancing. If we take the average across our the parameter estimates of our four outcome variables, our results suggest that a 1% increase in social distancing behavior of all other counties will lead to a 2% increase in social distancing behavior within a focal county. While this estimate may seem extreme, the COVID-19 pandemic certainly does count as extraordinary circumstances: a single family member or friend, let alone all of them, has the capability to drastically influence an individuals' social distancing behavior.[19]

Beyond just the peer effect, several other clear trends emerge. First, we find that county policy does have a statistically significant direct effect on social distancing behaviors. Moreover, largely consistent with the results of Sections S2.2 and S2.3, this direct effect is notably less impactful than what the naive DiD estimates in S2.1 would suggest. Furthermore, our results strongly suggest that the peer policy spillovers a largely mediated by endogenous peer behavior. Comparing the total policy spillover effects in Section S2.3 with those seen here, we see that the effect is generally much closer to generally quite close to 0.

S3.4.1 Visualization of the Instrumental Variables

It is common to visualize instrumental variables analyses with binned scatterplots that show how the outcome varies with the instrument, sometimes expressed in terms of first-stage fitted values (*36*). Often these plots partial out other covariates, as we do here. Fig. S29 show how two illustrative instruments are associated with one outcome, the measure of Facebook devices leaving their home location, asinh(NSBTUs). The precipitation instrument (Fig. S29 top left) is the first selected by the LASSO regularization path. Percentiles of this measure of rain in alter counties are associated with mobility in the focal county. Percentiles of the shift-share instrument (Fig. S29 bottom left) for the arts, entertainment, and recreation sector, which is

[19] https://www.nytimes.com/2020/03/31/opinion/coronavirus-fox-news.html

one of the first shift-share instruments selected by the LASSO regularization path, measure exposure of the county-day to national shifts in visitation of those locations. Bins formed by the intersection of deciles of these two instruments (Fig. S29 right) are highly associated with both residuals of mobility in the peer counties and residuals of mobility in the focal county.

S3.5 Robustness Checks

S3.5.1 Varying the Number of Selected instruments

Fig. S30 shows the estimate results of adding number of instruments based on their importance as obtained by LASSO selection. The first instrument selected is the most important instrument with the higher entry λ, followed by the second instruments; that is, this shows how the post-LASSO regularization path translates into our second-stage estimates of endogenous peer effects and effects of policies. When there is a very small number of instruments, the estimates are similar to what we have when we have a large number of instruments, with some variation as a few more instruments are added. This observation is consistent with what we observe when conducting randomization inference with a single instrument.

S3.5.2 Only Weather Instruments

Table S9 presents results from an instrumental variable regression analysis using only the weather instruments. Although the estimated peer effect coefficient is consistently smaller compared to Table S8, the magnitudes are still quite comparable. Moreover, the remainder of the results are largely both qualitatively and quantitatively consistent with our main results.

S3.5.3 Only Shift-share Instruments

Table S10 presents results from an instrumental variable regression analysis using only the shift-share instruments. Much like our analysis using only weather instruments, these results again confirm our main results. We find strong evidence of large, positive endogenous social effects

Table S8: Effects of policy interventions and endogenous alter behaviors (Full set of instruments)

	Dependent variable:			
	BTVRC	asinh(NSBTUs)	asinh(NHDF)	log(dCBGVs)
	(1)	(2)	(3)	(4)
Close gyms, movie theaters, restaurants	-0.011^{**}	-0.003	-0.003	-0.007
	(0.005)	(0.002)	(0.002)	(0.006)
Close gyms, movie theaters, restaurants (social alters)	0.004	-0.001	0.007	0.002
	(0.012)	(0.005)	(0.007)	(0.020)
Close gyms, movie theaters, restaurants (geo alters)	0.011	0.003	-0.003	0.006
	(0.009)	(0.003)	(0.005)	(0.012)
Shelter-in-place	-0.033^{***}	-0.010^{***}	-0.015^{***}	-0.030^{***}
	(0.009)	(0.003)	(0.004)	(0.011)
Shelter-in-place (social alters)	0.051^{***}	0.020^{***}	0.042^{***}	0.066^{**}
	(0.019)	(0.007)	(0.012)	(0.029)
Shelter-in-place (geo alters)	-0.005	-0.004	-0.016^{**}	-0.020
	(0.010)	(0.004)	(0.007)	(0.016)
Social alters' BTVRC	1.650^{***}			
	(0.039)			
Social alters' asinh(NSBTUs)		1.620^{***}		
		(0.054)		
Social alters' asinh(NHDF)			1.630^{***}	
			(0.052)	
Social alters' log(dCBGVs)				1.910^{***}
				(0.052)
Conley s.e.	Yes	Yes	Yes	Yes
County fixed effect	Yes	Yes	Yes	Yes
Day fixed effect	Yes	Yes	Yes	Yes
Observations	122,598	122,598	122,598	122,598
R^2	0.824	0.876	0.816	0.818
Adjusted R^2	0.821	0.874	0.812	0.814
Residual Std. Error (df = 120016)	8.420	2.950	4.800	8.140

Notes: NSBTU refers to the fraction of FB users in a given county who visit multiple Bing tiles on a given day. NHDF refers to the fraction of devices that are not completely at home, as measured by Safegraph. dCBGVs refers to the mean number of Census block groups devices from a given county visit, as measured by Safegraph. BTVRC refers to the relative change in the number of Bing tiles users visit, as measured by FB. *p<0.1; **p<0.05; ***p<0.01

Table S9: Weather instruments only

	Dependent variable:			
	asinh(NSBTUs)	BTVRC	asinh(NHDF)	log(dCBGVs)
Close GMR	−0.003	−0.009	−0.004	−0.007
	(0.003)	(0.006)	(0.003)	(0.007)
Close GMR (Social Alters)	−0.003	−0.006	0.005	0.001
	(0.008)	(0.015)	(0.009)	(0.024)
Close GMR (Geo Alters)	0.005	0.009	−0.002	0.004
	(0.005)	(0.010)	(0.006)	(0.015)
Shelter-in-place	−0.010**	−0.030***	−0.015**	−0.030**
	(0.005)	(0.010)	(0.006)	(0.013)
Shelter-in-place (Social Alters)	0.006	0.013	0.025	0.045
	(0.015)	(0.031)	(0.017)	(0.039)
Shelter-in-place (Geo Alters)	−0.0003	−0.001	−0.011	−0.020
	(0.007)	(0.015)	(0.009)	(0.020)
Peer Effect	1.393***	1.362***	1.515***	1.477***
	(0.114)	(0.079)	(0.084)	(0.080)

*p<0.1; **p<0.05; ***p<0.01

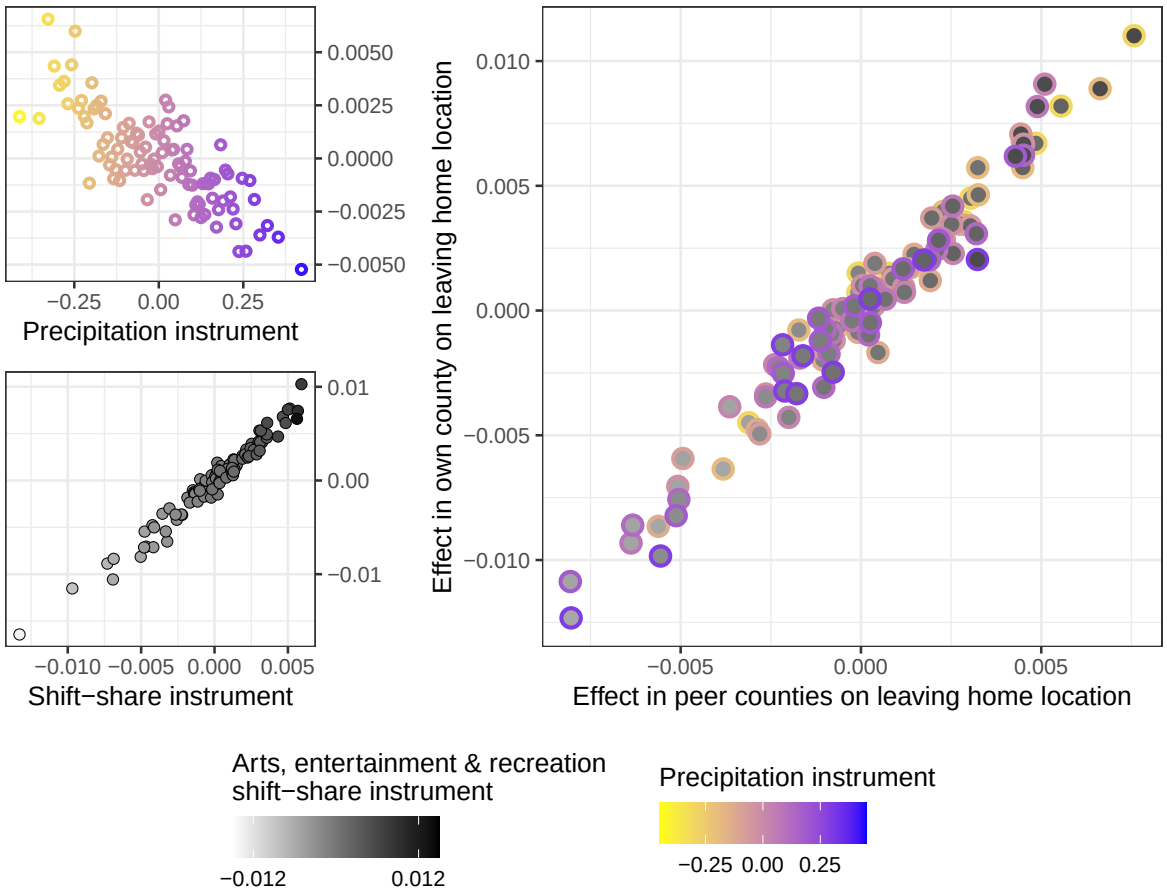

Figure S29: Visual instrumental variables plot for a selected weather instrument and a selected shift-share instrument and Facebook devices leaving their home location, asinh(NSBTUs). (left) Binned scatter plots of instrument residuals and outcome residuals. (top left) As precipitation in more alter counties exceeds a low threshold (0.06cm), more devices remain in their home location in the focal (ego) county. (bottom left) For ego counties with alter counties that are more exposed to national, negative shifts in people going to arts, entertainment, and recreation locations (NAICS sector 71), more devices remain in their home location in the ego county. (right) Forming bins by the deciles of each instrument, we see that their combined effect on mobility in alter counties translates to effects on mobility in the ego county.

in social distancing behavior. Moreover, we again see that these endogenous social effects are largely mediating the effect of peer county policy spillovers.

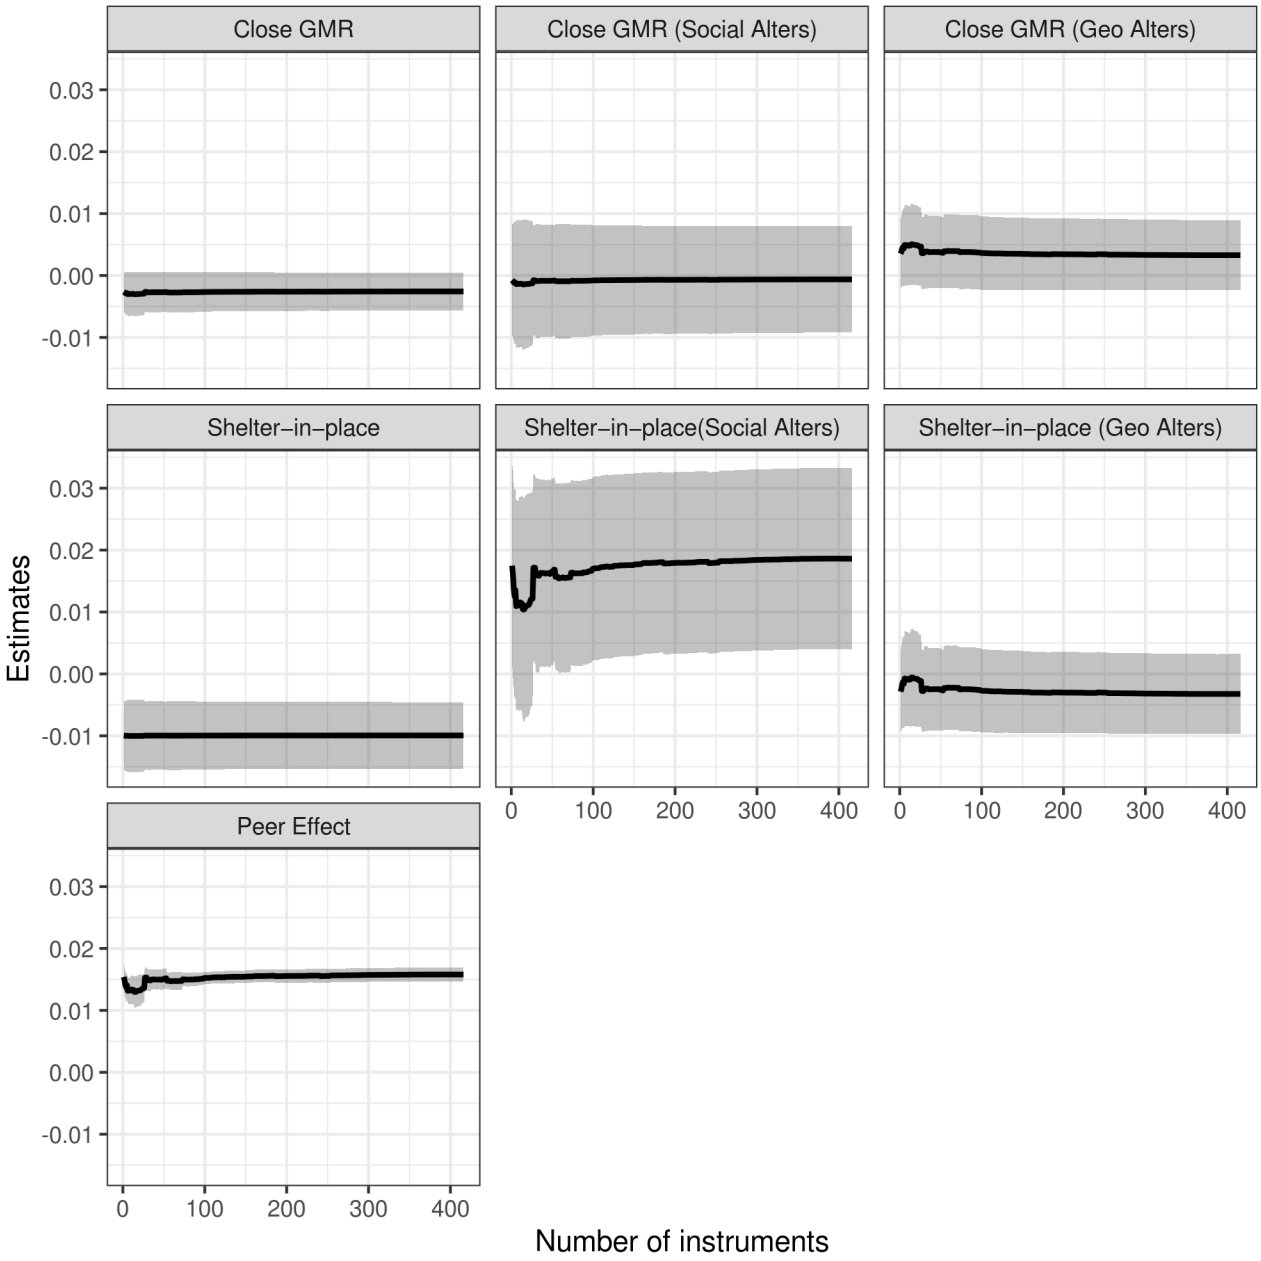

Figure S30: Varying the number of instruments, where instruments are added according to the LASSO regularization path.

Table S10: Shift-share instruments only

	Dependent variable:			
	asinh(NSBTUs)	BTVRC	asinh(NHDF)	log(dCBGVs)
Close GMR	−0.003	−0.011	−0.003	−0.007
	(0.003)	(0.006)	(0.003)	(0.008)
Close GMR (Social Alters)	−0.001	0.005	0.007	0.005
	(0.007)	(0.016)	(0.009)	(0.025)
Close GMR (Geo Alters)	0.004	0.008	−0.003	0.004
	(0.005)	(0.012)	(0.006)	(0.016)
Shelter-in-place	−0.010***	−0.032***	−0.015***	−0.031**
	(0.004)	(0.010)	(0.005)	(0.013)
Shelter-in-place (Social Alters)	0.017*	0.055**	0.042***	0.079**
	(0.009)	(0.025)	(0.015)	(0.037)
Shelter-in-place (Geo Alters)	−0.003	−0.008	−0.016*	−0.025
	(0.006)	(0.015)	(0.010)	(0.022)
Peer Effect	1.528***	1.602***	1.616***	1.694***
	(0.067)	(0.073)	(0.082)	(0.092)

*p<0.1; **p<0.05; ***p<0.01

S3.5.4 Randomization Inference

While our primary asymptotic inference described above was selected to be robust to network and within-state dependence, we use Fisherian randomization inference (FRI) to further examine the robustness of our instrumental variable results to dependence in key regressors — in particular, the weather instruments (*37, 38*). Let $\tilde{Y}_{it}$ be the outcome, $\tilde{Y}_{-it}$ the endogenous peer behavior, and $\tilde{W}_{-it}$ the vector of instrumental variables. Under the maintained exclusion restriction and the sharp null hypothesis that the potential outcomes are given by $\tilde{Y}_{it}(d) = \tilde{Y}_{-it}(0) + \beta_0 \tilde{Y}_{-it}$, then $\tilde{Y}_{it} - \beta_0 \tilde{Y}_{-it}$ is invariant to changes to $\tilde{W}_{-it}$. For the special case of $\beta_0 = 0$, this null hypothesis implies that $\tilde{Y}_{it}$ is invariant to changes in $\tilde{W}_{-it}$.

This kind of "reduced form" test with instrumental variables (*37*) is known to have low power when overidentified (*39*), so we focus on the case of a single weather instrument. We draw a permuted weather instrument W_{-it} from a historical weather distribution that preserves the correlation across consecutive days and geographic locations. In particular, we draw historical weather within a 28 days time window before and after the focal days of the year from dates in 2015–2019 for all counties. This gives us a total of $28 \times 2 \times 4 = 224$ counterfactual weather conditions. We compute alters' weather as in the primary analysis, partialing out i's weather and fixed effects as with the observed data to create $\tilde{W}_{-it}$.

For each of these permuted sets of weather instruments, we compute a test statistic measuring the association of $Y_{it} - \beta Y_{-it}$ with W_{-it}. Here we use the F-statistic for a single weather instrument that measures precipitation in alter counties. For each of the 4 primary movement outcomes we look at (BTVRC, asinh(NSBTUs), log(dCBGVs), and asinh(nhd)), we plot this null distribution of F-statistics and the observed F-statistic in Fig. S31. The red vertical lines correspond to the F statistics under current weather, the p-value (fraction of times the F-statistic is equal to or bigger than the F-statistic corresponding to the current weather) is reported under each test. Even with only this single instrument, we can reject all null hypotheses at the 5%

level except for the outcome asinh(NSBTUs). This provides some evidence for our results for endogenous peer effects that does not rely on asymptotic approximations.

S3.5.5 Excluding Counties with Correlated Weather

As a robustness check on the conditional ignorability assumption of our weather instruments, we conduct a sensitivity analysis where we exclude counties with sufficiently correlated weather to a focal county i from our analysis. Our model specifications in this section only make use of the weather instruments. We run 3 styles of analysis following this selection procedure:

1. Renormalized: the weights used to construct Y_{-it} and W_{-it} are renormalized to sum to 1. This procedure changes the interpretation of β to be the endogenous social effect of all peer counties with sufficently uncorrelated weather.

2. Non-renormalized: the weights used to construct Y_{-it} and W_{-it} are not renormalized. As such the maximum sum of all the weights equals the (weighted) fraction of non-excluded peer counties.

3. Full Endogenous Variable: regardless of the threshold, no counties are excluded from the construction of Y_{-it}. However, non-renormalized weights are used to compute W_{-it}.

We report the results using both 2SLS and LIML estimation for each of our 4 outcome variables with thresholds starting at 0.05 until 1 in 0.05 increments. The corresponding figures and outcomes are as follows:

- Fig. S32: asinh(NSBTUs)
- Fig. S33: asinh(NHD)
- Fig. S34: log(dCBGVs)

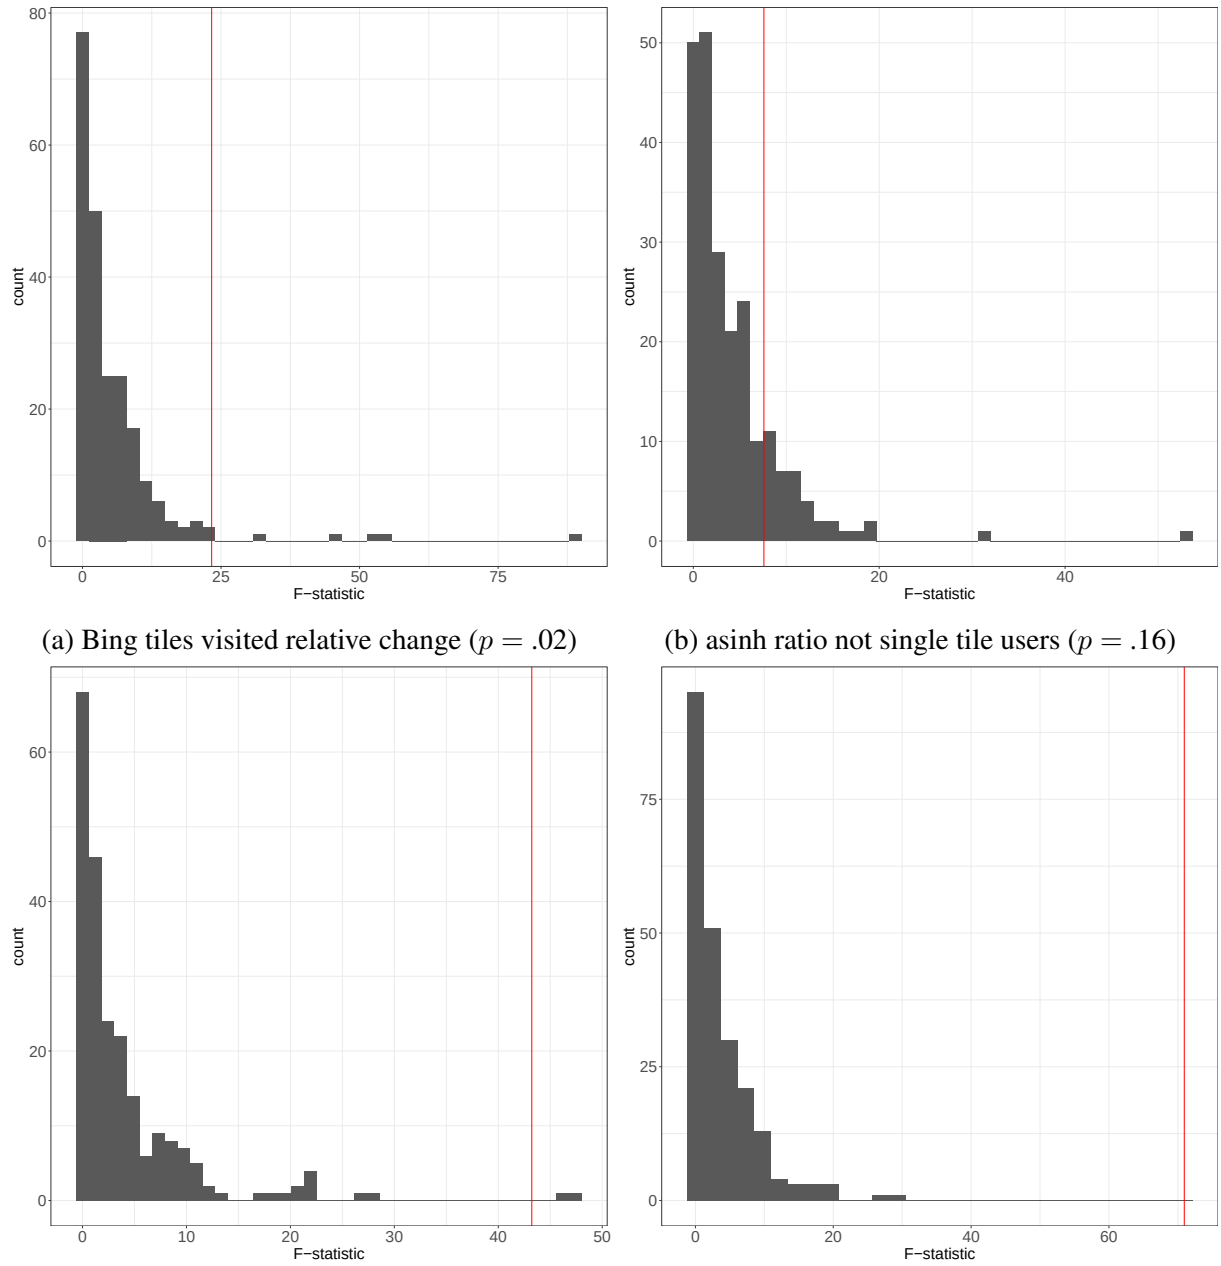

(a) Bing tiles visited relative change ($p = .02$)

(b) asinh ratio not single tile users ($p = .16$)

(c) log mean census block groups visited ($p = .01$)

(d) asinh not home devices ($p < .01$)

Figure S31: Null distribution (grey) and observed value (red line) of test statistic (F-statistic) for association between a single weather instrument and the outcomes from Fisherian randomization inference using historical weather data.

- Fig. S35: BTVRC

A couple of major patterns emerge from our results. The direct policy effects stay quite consistent across outcomes, estimation methods, thresholds, and model types. We also see that generally, as we decrease the threshold, the estimated impact of peer policy becomes more significant, at least for both the renormalized and non-renormalized model types. This is consistent with the idea that these endogenous peer mobility behaviors are causally mediating many of the effects of peer policies.

Beyond this, we also see that as the threshold decreases, the peer effect for the renormalized models drops to 0. This also makes sense, since as we exclude more and more peer counties, their total effect is becoming less and less. Looking at the non-renormalized models, we see that there is some fluctuation as the threshold decreases, however it remains relatively stable across a wide range of possible thresholds. The full endogenous variable results are even more consistent, with LIML estimates that are nearly flat line across the entire range of thresholds.

Overall, we see these results as extremely consistent with our main results. As such, we believe the DML approach is effectively controlling for the impact of own weather, and giving us confidence that our weather instruments are indeed conditionally ignorable.

S4 Asymptotic Inference with Adjacency- and Cluster-robust Standard Errors

Spatial and network econometricians have previously made use of estimators for variance-covariance matrices that are consistent in the presence of spatially correlated errors (*40, 41*). These estimators are Huber–White "sandwich" estimators with the functional form:

$$\widehat{\text{Var}}\left[\hat{\beta}\right] = (\mathbf{X}'\mathbf{X})^{-1}\,\mathbf{X}\,(\hat{u}\hat{u}'\odot\mathbf{B})\,\mathbf{X}\,(\mathbf{X}'\mathbf{X})^{-1}\,, \tag{S11}$$

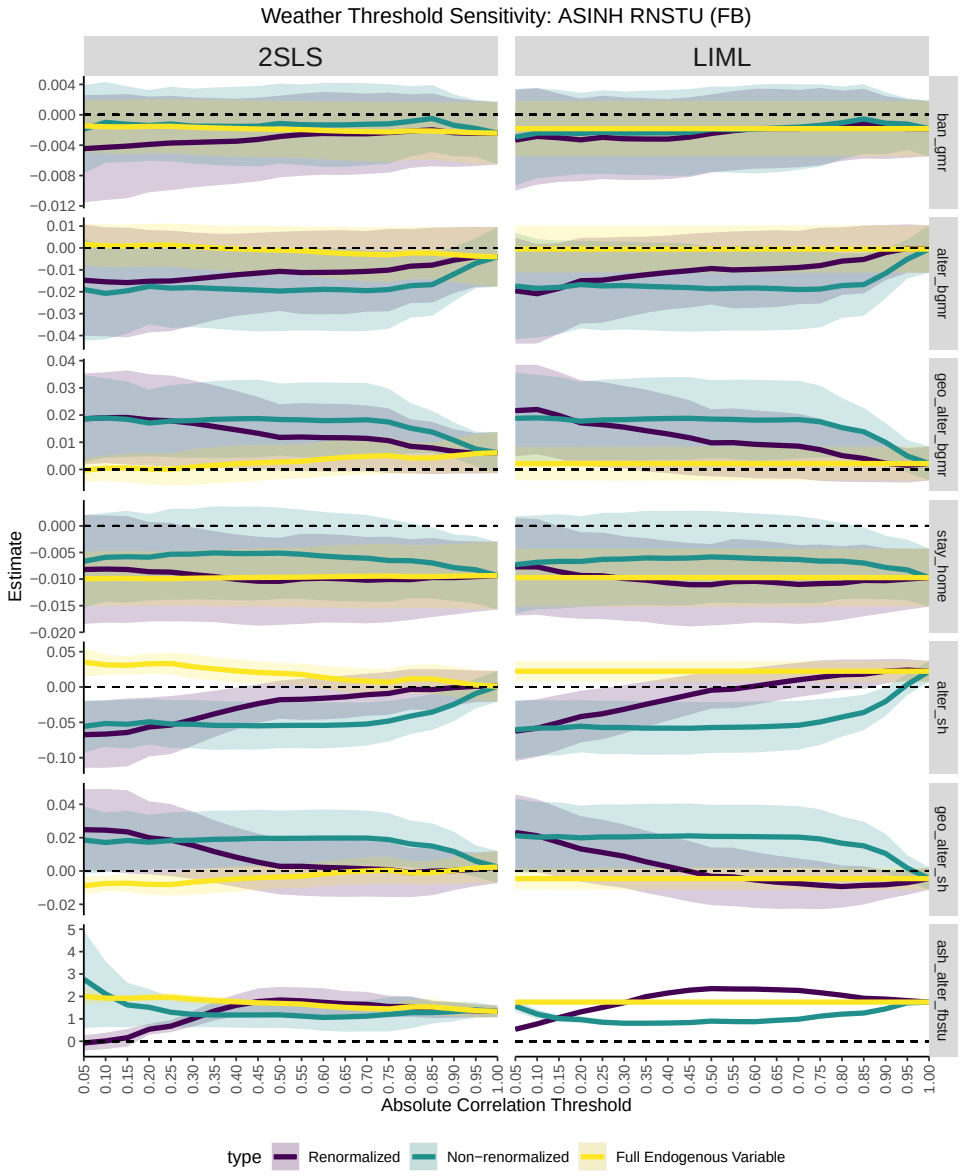

Figure S32: Coefficient plots showing how the estimated parameters $\hat{\beta}$, $\hat{\delta}_1$, $\hat{\delta}_2$, and $\hat{\delta}_3$ change as counties with correlated weather are excluded for the outcome variable of asinh(NSBTUs). Starting at the right where all counties are included, we move in 0.05 increments until we reach the left side, where only counties with less than 0.05 absolute precipitation correlation are included. 2 types of models are estimated: Renormalized (purple), Non-renormalized (green), and Full Endogenous Variable (Yellow). Coefficient estimates produced using both 2SLS (left) and LIML (right) are included. Due to the high computational cost of adjacency-robust SEs, state-clustered standard are reported here instead.

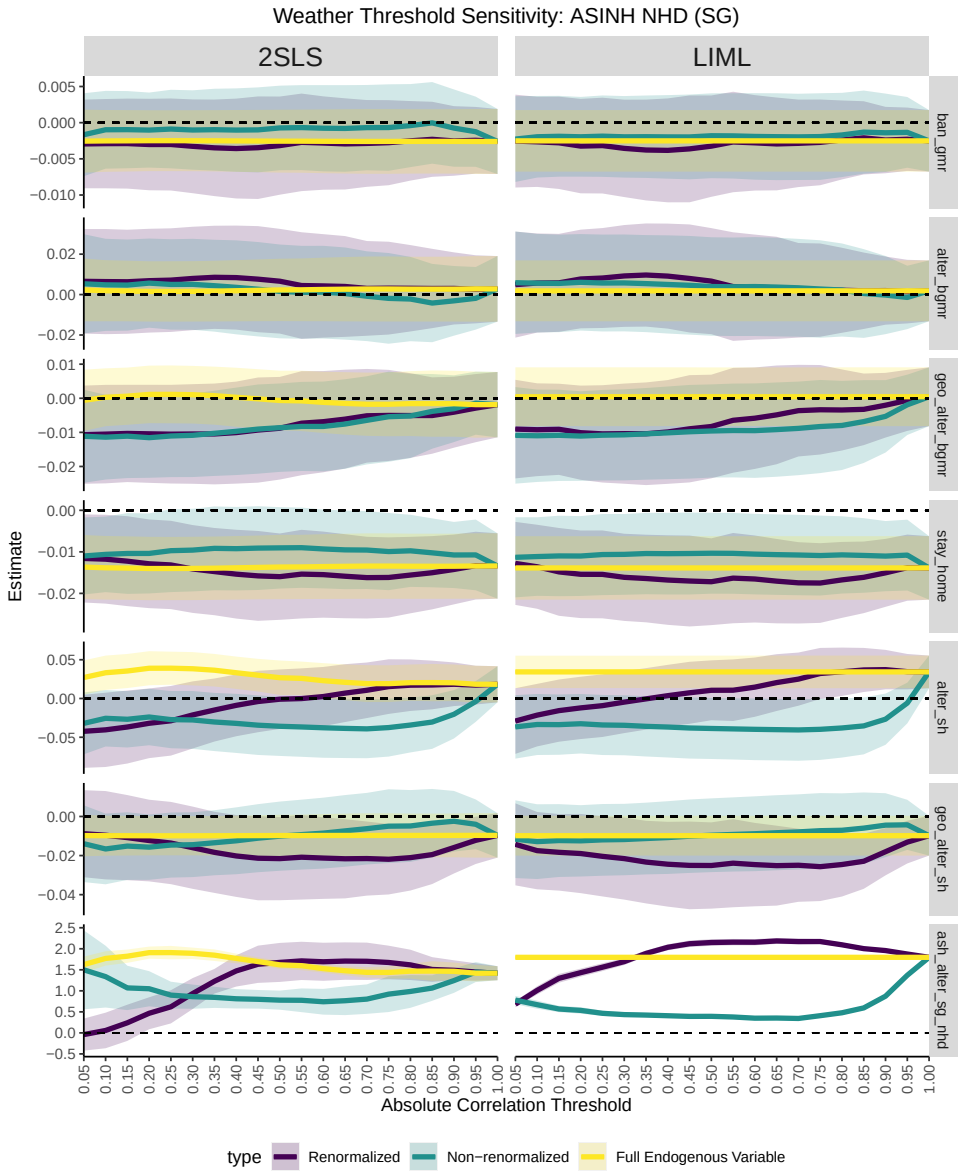

Figure S33: Coefficient plots showing how the estimated parameters $\hat{\beta}$, $\hat{\delta}_1$, $\hat{\delta}_2$, and $\hat{\delta}_3$ change as counties with correlated weather are excluded for the outcome variable of asinh(NHDF). Starting at the right where all counties are included, we move in 0.05 increments until we reach the left side, where only counties with less than 0.05 absolute precipitation correlation are included. 2 types of models are estimated: Renormalized (purple), Non-renormalized (green), and Full Endogenous Variable (Yellow). Coefficient esimates produced using both 2SLS (left) and LIML (right) are included. Due to the high computational cost of adjacency-robust SEs, state-clustered standard are reported here instead.

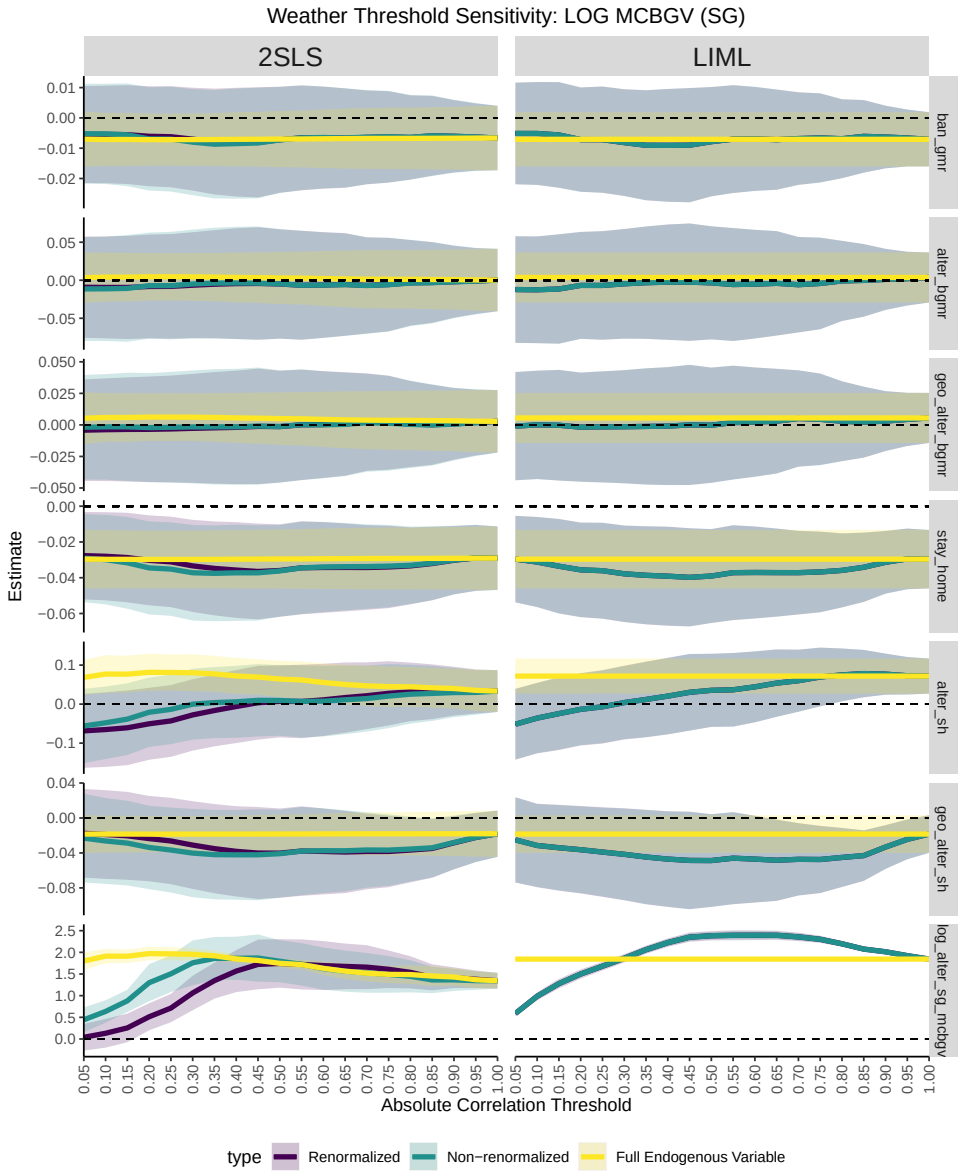

Figure S34: Coefficient plots showing how the estimated parameters $\hat{\beta}$, $\hat{\delta}_1$, $\hat{\delta}_2$, and $\hat{\delta}_3$ change as counties with correlated weather are excluded for the outcome variable of log(dMCBGVs). Starting at the right where all counties are included, we move in 0.05 increments until we reach the left side, where only counties with less than 0.05 absolute precipitation correlation are included. 2 types of models are estimated: Renormalized (purple), Non-renormalized (green), and Full Endogenous Variable (Yellow). Coefficient estimates produced using both 2SLS (left) and LIML (right) are included. Due to the high computational cost of adjacency-robust SEs, state-clustered standard are reported here instead.

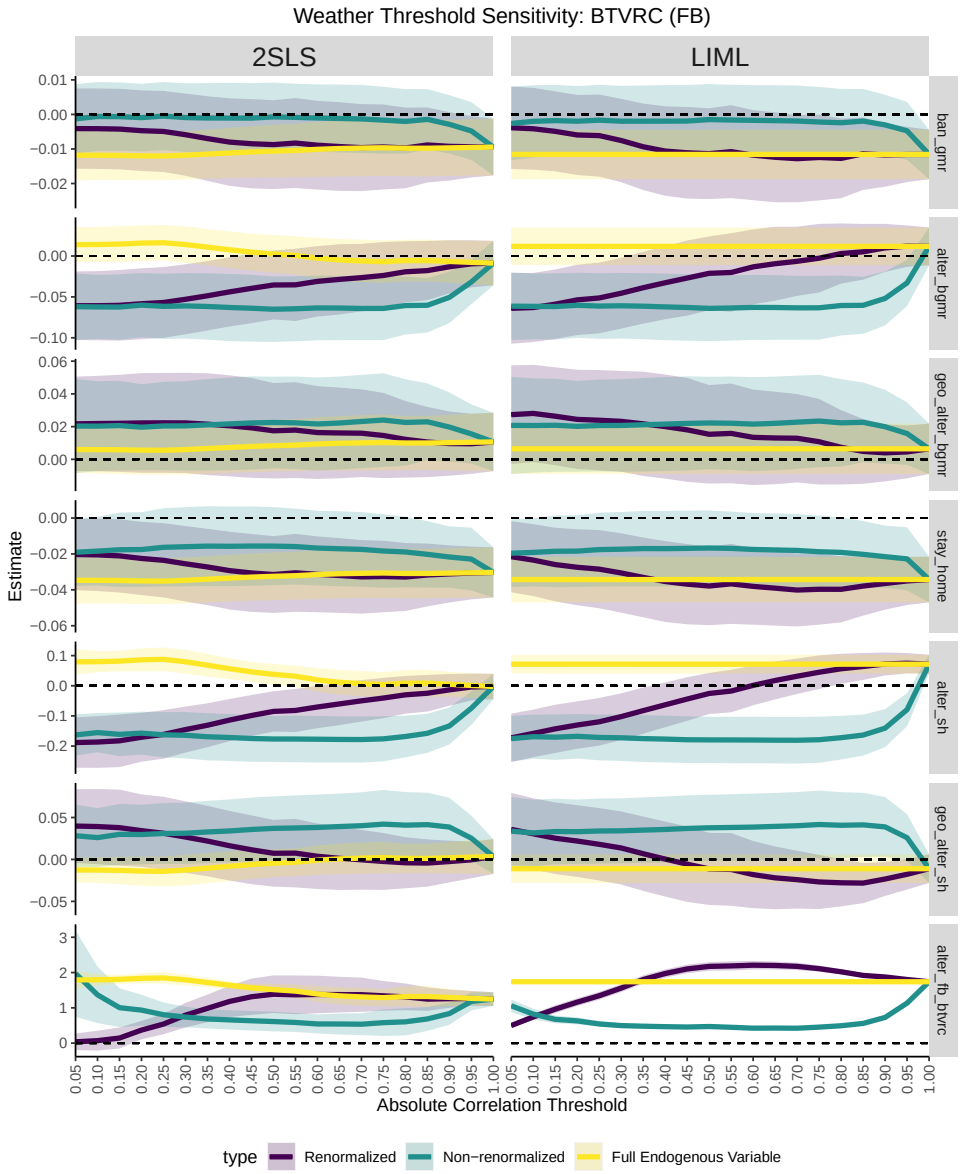

Figure S35: Coefficient plots showing how the estimated parameters $\hat{\beta}$, $\hat{\delta}_1$, $\hat{\delta}_2$, and $\hat{\delta}_3$ change as counties with correlated weather are excluded for the outcome variable of BTVRC. Starting at the right where all counties are included, we move in 0.05 increments until we reach the left side, where only counties with less than 0.05 absolute precipitation correlation are included. 2 types of models are estimated: Renormalized (purple), Non-renormalized (green), and Full Endogenous Variable (Yellow). Coefficient estimates produced using both 2SLS (left) and LIML (right) are included. Due to the high computational cost of adjacency-robust SEs, state-clustered standard are reported here instead.

where $\hat{u}$ is the vector of residuals, $\odot$ is element-wise multiplication, and $\mathbf{B}$ is an $n \times n$ matrix that selects and/or weights pairs of observation. In versions of this estimator that are "adjacency-robust," $\mathbf{B}$ is related to a given network's adjacency matrix, whereas in versions of this estimator that are robust to one-way clustering, $\mathbf{B}$ is a block diagonal matrix with $B_{ij} = 1$ if and only if i and j are in the same cluster.

In our setting, we believe that there is the potential for errors to be correlated for counties that are connected according to our social and/or geographic connectedness networks, as well as counties that are located in the same U.S. state (since the enactment of many COVID-19 related policies was made at the state-level, as opposed to the county-level). In order to account for this, we use a variance-covariance estimator that combines both adjacency- and cluster/state-robust estimators with the functional form provided in Eq. 11. The adjacency-robust estimator consists of Eq. 11, with the "selector" matrix $\mathbf{B}$ equal to the elementwise maximum of the social adjacency matrix used in our analyses and the state clustering matrix, i.e., $B_{ij}^{both} = \max(B_{ij}^{\text{ad}}, B_{ij}^{\text{clu}})$ (*41–44*). Our difference-in-differences and instrumental variables analyses make use of this combined adjacency-and cluster-robust sandwich estimator.

S5 Ranking States' Influence

In order to rank the extent to which each individual U.S. state's shelter-in-place policy decision impacts each other U.S. state, we combined the different-state coefficient estimates obtained in our difference-in-differences models (with asinh(nhd) as the outcome of interest) with the social adjacency and geographic adjacency matrices that we use throughout our analysis. In other words, we define state i's influence on state j as:

$$\text{Inf}_{ij} = \delta_2^{ds} A_{ij}^{geo} + \delta_3^{ds} A_{ij}^{social}. \tag{S12}$$

Fig. S36 and Fig. S37 show the 20 alter states that would cause the largest reduction in

the focal state by implementing their own shelter-in-place policy. We next define the spillover between alter state i and ego state j as the ratio of the indirect effect of i's shelter-in-place policy on mobility levels in j to the direct effect of j's shelter in place policy on mobility levels in j. Fig. S38 shows the empirical CDF for the set of state pairs consisting of each ego state's most influential alter state. The min spillover is 0.22, the maximum spillover is 0.83, and the median spillover is 0.42.

For each state j, we can also estimate its "total influence" by taking the weighted sum of our estimated state-level influence across all states. In other words,

$$TI_j = \sum_i n_i \left(\delta_2^{ds} A_{ij}^{geo} + \delta_3^{ds} A_{ij}^{sci} \right), \tag{S13}$$

where n_i is the population of state i. TI_j captures the total reduction in mobility that j would trigger across the rest of the U.S. by implementing their own shelter-in-place policy (assuming no other states have implemented their own). Fig. S39 shows the relationship between total state influence and state population. In general, more populous states have greater capacity to impact other states' mobility levels.

Finally, we can estimate the total influence that each region of the U.S. has on each focal state by summing individual "influences" at the region-level, i.e.,

$$I_{ir} = \sum_{j \in r} \left(\delta_2^{ds} A_{ij}^{geo} + \delta_3^{ds} A_{ij}^{sci} \right). \tag{S14}$$

S6 Analytical Model of the Loss From Anarchy

To illustrate the "loss from anarchy", that is, states failing to cooperate over network effects, we introduce a simple game-theoretic model. This model is inspired by both the classic literature on inefficiencies arising from unpriced externalities and non-cooperative game theory (*45*) as

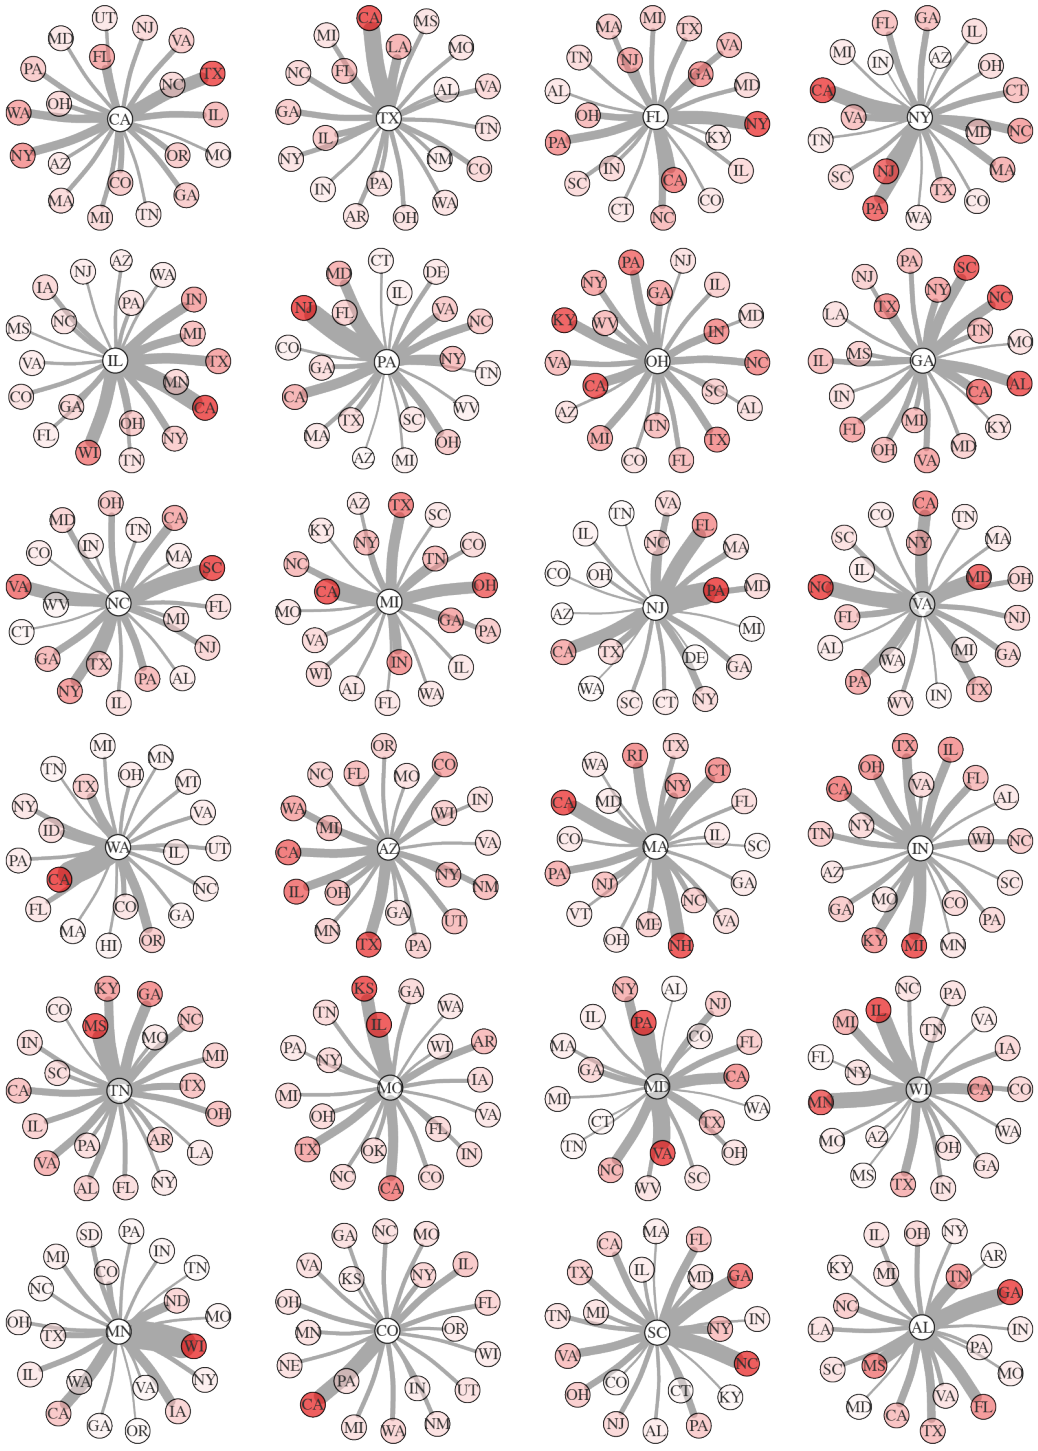

Figure S36: Each ego network shows the alter states whose own shelter-in-place policies would cause the greatest mobility reductions in the ego state. Edge thickness and alter node color both correspond to the log of this mobility reduction. Estimates are obtained by combining point estimates from our difference-in-differences model, our social adjacency matrix, and our geo-adjacency matrix. From left-top to right-bottom focal states are sorted by their population.

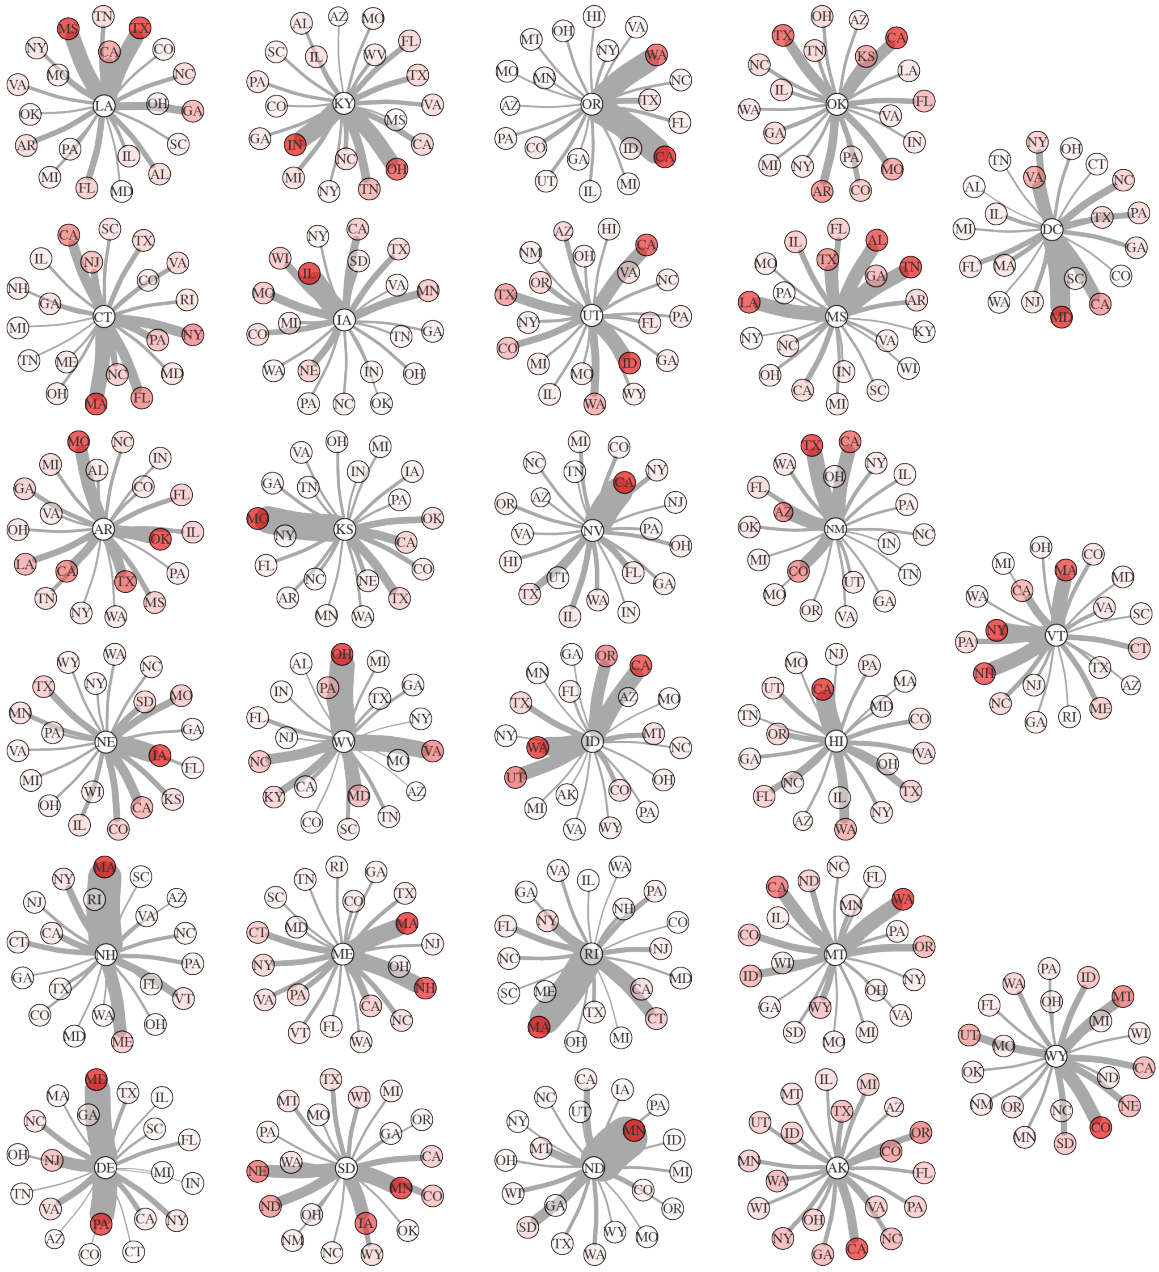

Figure S37: Each ego network shows the alter states whose own shelter-in-place policies would cause the greatest mobility reductions in the ego state. Edge thickness and alter node color both correspond to the log of this mobility reduction. Estimates are obtained by combining point estimates from our difference-in-differences model, our social adjacency matrix, and our geo-adjacency matrix. From left-top to right-bottom focal states are sorted by their population.

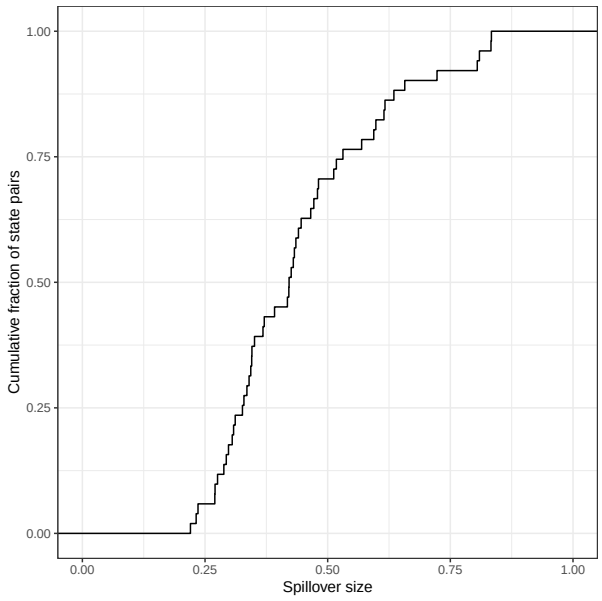

Figure S38: The empirical CDF of the spillover between state i and j, defined as the ratio of the indirect effect of i's shelter-in-place policy on j's mobility to the direct effect of j's shelter-in-place policy on j's mobility. Sample limited to the most influential alter state for each ego state ($n = 51$).

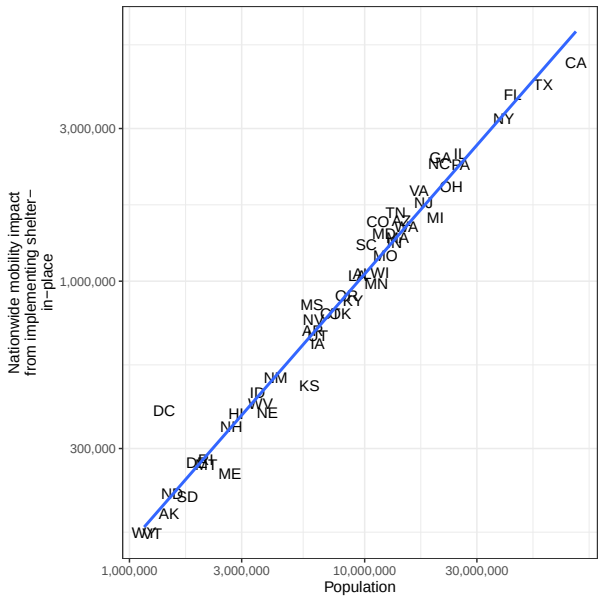

Figure S39: The relationship between a state's "total influence", i.e. how much nationwide out-of-state mobility would reduce after that state implemented shelter-in-place, and the state's population. The two are highly correlated.

well as the more modern reinterpratation of these losses as a "price of anarchy" (*46*). We refer to a loss from anarchy because in the main text we focus on a difference in utilities rather than the ratio which is more common to the price of anarchy literature.

The agents in our model are a set of states, indexed by s, each attempting to achieve a target mobility level, denoted by $\hat{r}_s$ by choosing an intensity of distancing policy $p_s \in \mathbb{R}$ which has a linear cost c. Restricting attention to two states i and j, the policy decisions of state i spillover to state j linearly (but potentially asymmetrically) according to parameter $0 \leq z_{ij} \leq 1$:

$$r_i = \bar{r}_i - p_i - p_j z_{ji}, \tag{S15}$$

$$r_j = \bar{r}_j - p_j - p_i z_{ij}, \tag{S16}$$

where $\bar{r}_s \geq 0$ is the initial mobility level in state s.

We assume that each state's cost of adopting a policy level p_s is linear at c per unit. We further assume that their cost of missing their target mobility level is quadratic. Their utility function can then be written as

$$U_s(r_s, p_s) = -(r_s - \hat{r}_s)^2 - cp_s. \tag{S17}$$

In the following sections, we solve for the Nash equilibrium outcome of the model, as well as the choice of a social planner who equally weighs states. We solve for the two state problem first, and then for more complex settings.

This simple model generates a loss from lack of cooperation so long as spillovers exist (i.e. $z_{ij} > 0, \forall_{i,j}$), changing policies is not costless ($c > 0$), and transfer payments between states are not allowed. States with opposed targets will wastefully pull in different directions. Even when states are identical, there is a free rider problem. In Equation (S47) we show that the degree of divergence between the social planner and Nash outcomes is increasing in the size of spillovers. To overcome this Coasian challenge, interstate compacts (potentially with monetary or resource transfers to redistribute the surplus from cooperation) are desirable.

S6.1 Nash Equilibrium Solution

Let $\delta_i = \hat{r}_i - \bar{r}_i$ and without loss of generality assume that $\Delta := \delta_j - \delta_i \geq 0$. Replacing (S15) in (S17), the optimal policies, p_i^{NE} and p_j^{NE}, for the two states satisfy:

$$p_i^{NE} = -\frac{c}{2} - \delta_i - p_j^{NE} z_{ji} \tag{S18}$$

$$p_j^{NE} = -\frac{c}{2} - \delta_j - p_i^{NE} z_{ij} \tag{S19}$$

When the spillovers are perfect (i.e. $z_{j,i} = z_{i,j} = 1$) there are no pure strategy equilibria unless $\delta_i = \delta_j = \delta$, in which case $p_i^1 = p_j^1 = -c/4 - \delta/2$, where we use p_s^1 to denote the equilibrium state policies with maximum spillover.

We use p_s^0 to denote the equilibrium state policies with no spillovers. From (S18) we have:[20]

$$p_i^0 = -\frac{c}{2} - \delta_i, \tag{S20}$$

$$p_j^0 = -\frac{c}{2} - \delta_j, \tag{S21}$$

with the corresponding utilities

$$U_i^0 = \frac{c^2}{4} + c\delta_i, \tag{S22}$$

$$U_j^0 = \frac{c^2}{4} + c\delta_j, \tag{S23}$$

and social welfare

$$U^0 = U_i^0 + U_j^0 = \frac{c^2}{2} + c(\delta_i + \delta_j). \tag{S24}$$

Note that with $\delta_j \geq \delta_i$, we have $p_j^0 \leq p_i^0$: the state with the lowest target has to make the most effort and take the highest policy.

[20] We drop "NE" from the notation, as without spillovers the Nash Equilibrium and Socially Optimal strategies are identical.

When $0 \leq z_{ij} < 1$ and $0 \leq z_{ji} < 1$ we can solve (S18) to get the following equilibrium:

$$p_i^{NE} = -\frac{c(1-z_{ji})}{2(1-z_{ji}z_{ij})} - \delta_i \frac{1}{1-z_{ji}z_{ij}} + \delta_j \frac{z_{ji}}{1-z_{ji}z_{ij}} \tag{S25}$$

$$= p_i^0 \frac{1-z_{ji}}{1-z_{ji}z_{ij}} + \Delta \frac{z_{ji}}{1-z_{ji}z_{ij}} \tag{S26}$$

$$= \alpha_{ji} p_i^0 + \Delta_{ji}, \tag{S27}$$

$$p_j^{NE} = -\frac{c(1-z_{ij})}{2(1-z_{ji}z_{ij})} - \delta_j \frac{1}{1-z_{ji}z_{ij}} + \delta_i \frac{z_{ij}}{1-z_{ji}z_{ij}} \tag{S28}$$

$$= p_j^0 \frac{1-z_{ij}}{1-z_{ji}z_{ij}} - \Delta \frac{z_{ij}}{1-z_{ji}z_{ij}} \tag{S29}$$

$$= \alpha_{ij} p_j^0 - \Delta_{ij} \leq p_j^0, \tag{S30}$$

where $0 < \alpha_{ji} = (1-z_{ji})/(1-z_{ji}z_{ij}) < 1$ is the reduction in the effort by state i to use the spillover from the policy of state j, and vice versa for α_{ij}. On the other hand, $\Delta_{ji} = \Delta z_{ji}/(1-z_{ji}z_{ij})$ is the correction made by state i to account for the spillovers from the policy of state j with a differing desired target output level ($\Delta > 0$), and vice versa for Δ_{ij}. Note that state j with the higher target δ_j always reduces her policy as a result of spillovers: $p_j^{NE} \leq p_j^0$. Alternatively, state i who has the lower target increases her policy in face of spillovers only if $\Delta > (1-z_{ij})p_i^0$.

In fact, the gap between the two policies increases unbounded, with the increasing spillover size. With $z_{ij} = z_{ji} = z$ we have:

$$p_i^{NE} = \frac{p_i^0}{1+z} + \frac{z\Delta}{1-z^2},$$
$$p_j^{NE} = \frac{p_j^0}{1+z} - \frac{z\Delta}{1-z^2},$$

and

$$p_i^{NE} - p_j^{NE} = \frac{\Delta}{1+z} + \frac{2z\Delta}{1-z^2} = \Delta \frac{1+z}{1-z^2} \to \infty, \text{ as } z \to 1. \tag{S31}$$

The individual utilities and welfare in this equilibrium are as follows:

$$U_i^{NE} = \frac{c^2}{4}(2\alpha_{ji} - 1) + c\alpha_{ji}\delta_i - c\Delta_{ji}, \tag{S32}$$

$$U_j^{NE} = \frac{c^2}{4}(2\alpha_{ij} - 1) + c\alpha_{ij}\delta_j + c\Delta_{ij}, \tag{S33}$$

$$U^{NE} = U_i^{NE} + U_j^{NE} = -\frac{c^2}{2}(1 - \alpha_{ji} - \alpha_{ij}) + c\delta_i\alpha_{ij} + c\delta_j\alpha_{ji} \tag{S34}$$

$$= -\frac{c^2}{2}(1 - \alpha_{ji} - \alpha_{ij}) + c\alpha_{ji}\delta_i + c\alpha_{ij}\delta_j + c(\Delta_{ij} - c\Delta_{ji}) \tag{S35}$$

$$\leq U^0 + c\Delta\frac{z_{ij} - z_{ji}}{1 - z_{ji}z_{ij}} \tag{S36}$$

Hence, if $z_{ij} < z_{ji}$, i.e. the state with the higher target outcome has a higher spillover, then the equilibrium welfare in the presence of spillovers decreases by at least $c\Delta(z_{ji} - z_{ji})/(1 - z_{ji}z_{ij})$. On the other hand, if $0 < z_{ij} = z_{ji} = z$, then

$$U^{NE} = \frac{c^2}{2}\left(\frac{1-z}{1+z}\right) + \frac{c(\delta_i + \delta_j)}{1+z} = \frac{U^0}{1+z} - \frac{c^2 z}{2(1+z)} \leq \frac{U^0}{1+z}. \tag{S37}$$

For c large enough we have $0 < U^{NE} \leq U^0$, in which case the Nash equilibrium welfare in the presence of spillovers decreases by a factor of at least $1/(1+z)$.

S6.2 Socially Optimal Solution

We now shift attention to a social planner's problem who chooses p_i and p_j to maximize the social welfare $U_i(r_i, p_i) + U_j(r_j, p_j)$. The socially optimal policies p_i^{SO} and p_j^{SO} satisfy the following:

$$p_i^{SO}(1 + z_{ij}^2) + p_j^{SO}(z_{ij} + z_{ji}) = -\frac{c}{2} - \delta_i - z_{ij}\delta_j, \tag{S38}$$

$$p_j^{SO}(1 + z_{ij}^2) + p_i^{SO}(z_{ij} + z_{ji}) = -\frac{c}{2} - \delta_j - z_{ji}\delta_i. \tag{S39}$$

Solving for the socially optimal policies we get:

$$p_i^{SO} = -\frac{c(1 + z_{ji}^2 - z_{ij} - z_{ji})}{2(1 - z_{ji}z_{ij})^2} - \delta_i\frac{1}{1 - z_{ji}z_{ij}} + \delta_j\frac{z_{ji}}{1 - z_{ji}z_{ij}} \leq p_i^{NE}, \tag{S40}$$

$$p_j^{SO} = -\frac{c(1 + z_{ij}^2 - z_{ij} - z_{ji})}{2(1 - z_{ji}z_{ij})^2} - \delta_j\frac{1}{1 - z_{ji}z_{ij}} + \delta_i\frac{z_{ij}}{1 - z_{ji}z_{ij}} \leq p_j^{NE}, \tag{S41}$$

The subsequent individual utilities and social welfare under the above policies are as follows:

$$U_i^{SO} = \frac{c^2(1+2z_{ji}^2 - z_{ij}^2 - 2z_{ji})}{4(1-z_{ji}z_{ij})^2} + c\delta_i \frac{1}{1-z_{ji}z_{ij}} - c\delta_j \frac{z_{ji}}{1-z_{ji}z_{ij}}, \tag{S42}$$

$$U_j^{SO} = \frac{c^2(1+2z_{ij}^2 - z_{ji}^2 - 2z_{ij})}{4(1-z_{ji}z_{ij})^2} + c\delta_j \frac{1}{1-z_{ji}z_{ij}} - c\delta_i \frac{z_{ij}}{1-z_{ji}z_{ij}}, \tag{S43}$$

$$U^{SO} = U_i^{SO} + U_j^{SO} = \frac{c^2}{4}(\alpha_{ij}^2 + \alpha_{ji}^2) + c\delta_i\alpha_{ij} + c\delta_j\alpha_{ji}, \tag{S44}$$

Substituting $z_{ij} = z_{ji} = z$ in (S44) we get:

$$U^{SO} = \frac{c^2}{2}\left(\frac{1}{1+z}\right)^2 + \frac{c\delta_i + c\delta_j}{1+z} = \frac{U^0}{z+1} - \frac{c^2 z}{2(1+z)^2} \leq \frac{U^0}{z+1}. \tag{S45}$$

S6.3 Loss from Anarchy

To understand the inefficiency, consider the following cost of anarchy given by the difference between the socially optimal welfare U^{SO} in (S44) and the equilibrium welfare U^{NE} in (S34):

$$U^{SO} - U^{NE} = \frac{c^2}{4}(2 - 2\alpha_{ji} - 2\alpha_{ij} + \alpha_{ij}^2 + \alpha_{ji}^2) = \frac{c^2}{4}((1-\alpha_{ji})^2 + (1-\alpha_{ij})^2) \geq 0. \tag{S46}$$

Using (S37) and (S45), for $z_{ij} = z_{ji} = z > 0$, the disutility from lack of cooperation, ∇_z, is given by:

$$\nabla_z := U^{SO} - U^{NE} = \frac{c^2}{2}\left(\frac{z}{1+z}\right)^2 = \frac{c^2}{2}\left(\frac{1}{1+1/z}\right)^2, \tag{S47}$$

which is strictly increasing in z: *with no spillovers, there is no loss in efficiency but as spillovers increase, the lack of cooperation is a bigger and bigger problem*. As $c \to \infty$, the utility values are dominated by the c^2 terms and they become non-negative. In this regime, we can quantify the inefficiency of the Nash equilibrium in terms of its Price of Anarchy (PoA) as follows:

$$PoA = \frac{U^{SO}}{U^{NE}} \sim 1 + \frac{z^2}{1-z^2}, \text{ as } c \to \infty. \tag{S48}$$

It is instructive to briefly consider the case where policies are cost-less. Setting $c = 0$ in (S20), (S25), and (S34) we get $p_j^0 = -\delta_j \leq p_i^0 = -\delta_i$ and:

$$p_j^{NE} = p_j^0 - \frac{z}{1+z}(\Delta \frac{1}{1-z} + p_j^0) \leq p_j^0, \tag{S49}$$

$$p_i^{NE} = p_i^0 + \frac{z}{1+z}(\Delta \frac{1}{1-z} - p_i^0). \tag{S50}$$

In face of spillovers, state j with the higher mobility target always reduces her policy, while state i that has the lower mobility target reduces her policy only if $\delta_i/\delta_j < 2 - z$.

The case of n symmetric states. To understand the role of spillovers beyond the case of two interacting states, it is useful to take a brief look at the case with n symmetric and identical agents: $z_{ij} = z$ and $\delta_i = \delta$ for all $i, j = 1, .., n$, $i \neq j$. The Nash equilibrium policy and social welfare are as follows:

$$p_i^{NE} = \frac{-c}{2(1+z(n-1))} - \frac{\delta}{1+z(n-1)}, \forall_i \tag{S51}$$

$$U_{(n)}^{NE} = \sum_{i=1}^{n} U_i^{NE} = n\left(\frac{1-z(n-1)}{1+z(n-1)}\frac{c^2}{4} + \frac{c\delta}{1+z(n-1)}\right) \sim -n\frac{c^2}{4}, \tag{S52}$$

where $\sim$ denotes asymptotic equality: $f_n \sim g_n \leftrightarrow \lim_{n\to\infty} f_n/g_n = 1$. On the other hand, the socially optimal policy and welfare are given by:

$$p_i^{SO} = \frac{-c}{2(1+z(n-1))^2} - \frac{\delta}{1+z(n-1)}, \forall_i, \tag{S53}$$

$$U_{(n)}^{SO} = \sum_{i=1}^{n} U_i^{SO} = n\left(\frac{c^2}{4(1+z(n-1))^2} + \frac{c\delta}{1+z(n-1)}\right) \sim \frac{c\delta}{z}. \tag{S54}$$

The disutility from lack of cooperation among n states grows linearly worse as n increases:

$$\nabla_z^{(n)} := U_{(n)}^{SO} - U_{(n)}^{NE} \sim n\frac{c^2}{4}. \tag{S55}$$

Next suppose that the net spillover from other states is bounded; in particular, assume that

$zn \sim z'$ for some fixed $0 < z' < 1$. Then (S51) to (S55) become:

$$p_i^{NE} \sim \frac{-c}{2(1+z')} - \frac{\delta}{1+z'}, \forall_i, \tag{S56}$$

$$U_{(n)}^{NE} \sim n\left(\frac{1-z'}{1+z'}\frac{c^2}{4} + \frac{c\delta}{1+z'}\right), \tag{S57}$$

$$p_i^{SO} \sim \frac{-c}{2(1+z')^2} - \frac{\delta}{1+z'}, \forall_i, \tag{S58}$$

$$U_{(n)}^{SO} \sim n\left(\frac{c^2}{4(1+z')^2} + \frac{c\delta}{1+z'}\right), \tag{S59}$$

$$\nabla_{z'}^{(n)} \sim n\frac{c^2}{4}\left(\frac{z'}{1+z'}\right)^2. \tag{S60}$$

The disutility from lack of cooperation, $\nabla_z^{(n)}$, increases with increasing n, c, and z'. Similarly to (S48), we also have:

$$PoA = \frac{U_{(n)}^{SO}}{U_{(n)}^{NE}} \sim 1 + \frac{z'^2}{1-z'^2}, \text{ as } c \to \infty. \tag{S61}$$

It is worth noting that the mobility outcome, $r_i = \bar{r}_i - (1 + (n-1))zp_i$, at the Nash Equilibrium, is given by $r_i^{NE} = \hat{r}_i + c/2$, which overshoots the mobility target by $c/2$. At the socially optimal solution we get $r_i^{SO} = \hat{r}_i + c/2(1 + z(n-1))$, which, compared to the Nash equilibrium, overshoots the mobility target by a lesser amount: $c/2(1 + z(n-1))$. Due to the nonzero cost of implementing policies $c > 0$, both solution concepts predict mobility outcomes that are above the desired level $\hat{r}_i$; however, $r_i^{SO} \sim \hat{r}_i$, whereas $r_i^{NE} = \hat{r}_i + c/2$, independently of n or z.

The case of two blocks with geographic and social spillovers. Finally, to understand the effect of heterogeneity, we consider the case of n states that are organized in two geographic blocks, I and J, of sizes n_I and n_J, $n_I + n_J = n$. The states in block I all have the same mobility target, δ_I, and states in block J, all have the same mobility target δ_J. Each state $i \in I$ is subject to geographic spillovers (short ties) from other states within her own block, as well as social spillovers from the states in the other block, due to their social connectedness and long

ties. We assume that all the states within the same block are subject to the same net spillover effects. In particular, for all $i \in I$, let $0 \leq z_{II} = \sum_{i' \in I, i' \neq i} z_{i'i} \leq 1$ be the net spillovers from within block I and $0 \leq z_{JI} = \sum_{j \in J} z_{ji} \leq 1$ be the net spillovers from outside of block I. Similarly, for any $j \in J$, denote the net spillovers by $0 \leq z_{JJ} = \sum_{j' \in J, j' \neq j} z_{j'j} \leq 1$ and $0 \leq z_{IJ} = \sum_{i \in I} z_{ij} \leq 1$.

By symmetry all states within the same block adopt the same equilibrium policies. We denote the Nash equilibrium policies by p_I^{NE} and p_J^{NE} The equilibrium policies for each block are given by:

$$\begin{aligned}
p_I^{NE} &= \frac{-c}{2} \frac{1 + z_{JJ} - z_{JI}}{(1 + z_{II})(1 + z_{JJ}) - z_{JI} z_{IJ}} \\
&\quad - \delta_I \frac{1 + z_{JJ}}{(1 + z_{II})(1 + z_{JJ}) - z_{JI} z_{IJ}} + \delta_J \frac{z_{JI}}{(1 + z_{II})(1 + z_{JJ}) - z_{JI} z_{IJ}} \\
&= \frac{1 + z_{JJ}}{(1 + z_{II})(1 + z_{JJ}) - z_{JI} z_{IJ}} p_I^0 - \frac{z_{JI}}{(1 + z_{II})(1 + z_{JJ}) - z_{JI} z_{IJ}} p_J^0 \\
p_J^{NE} &= \frac{-c}{2} \frac{1 + z_{II} - z_{IJ}}{(1 + z_{II})(1 + z_{JJ}) - z_{JI} z_{IJ}} \\
&\quad - \delta_J \frac{1 + z_{II}}{(1 + z_{II})(1 + z_{JJ}) - z_{JI} z_{IJ}} + \delta_I \frac{z_{IJ}}{(1 + z_{II})(1 + z_{JJ}) - z_{JI} z_{IJ}} \\
&= \frac{1 + z_{II}}{(1 + z_{II})(1 + z_{JJ}) - z_{JI} z_{IJ}} p_J^0 - \frac{z_{IJ}}{(1 + z_{II})(1 + z_{JJ}) - z_{JI} z_{IJ}} p_I^0,
\end{aligned}$$

where $p_I^0 = -c/2 - \delta_I$ and $p_J^0 = -c/2 - \delta_J$ are the optimal policies in the absence of any spillovers — Equation (S20). Assuming further symmetries: $n_I = n_J = n/2$, $z_{II} = z_{JJ} = z_g$,

and $z_{IJ} = z_{JI} = z_s$, the Nash equilibrium simplifies as follows:

$$
\begin{aligned}
p_I^{NE} &= \frac{-c/2}{1+z_g+z_s} - \frac{1+z_g}{(1+z_g)^2 - z_s^2}\delta_I + \frac{z_s}{(1+z_g)^2 - z_s^2}\delta_J, \\
p_J^{NE} &= \frac{-c/2}{1+z_g+z_s} - \frac{1+z_g}{(1+z_g)^2 - z_s^2}\delta_J + \frac{z_s}{(1+z_g)^2 - z_s^2}\delta_I, \\
U_I^{NE} &= \frac{c^2}{4}\left(\frac{1-z_g-z_s}{1+z_g+z_s}\right) + c\delta_I\frac{1+z_g}{(1+z_g)^2 - z_s^2} - c\delta_J\frac{z_s}{(1+z_g)^2 - z_s^2}, \\
U_J^{NE} &= \frac{c^2}{4}\left(\frac{1-z_g-z_s}{1+z_g+z_s}\right) + c\delta_J\frac{1+z_g}{(1+z_g)^2 - z_s^2} - c\delta_I\frac{z_s}{(1+z_g)^2 - z_s^2}, \\
U_{g,s}^{NE} &= \frac{n}{2}(U_I^{NE} + U_J^{NE}) = n\frac{c^2}{4}\left(\frac{1-z_g-z_s}{1+z_g+z_s}\right) + \frac{nc(\delta_I+\delta_J)/2}{1+z_g+z_s}.
\end{aligned}
$$

The social optimum in this case is given by:

$$
\begin{aligned}
p_I^{SO} &= \frac{-c/2}{(1+z_g+z_s)^2} - \frac{1+z_g}{(1+z_g)^2 - z_s^2}\delta_I + \frac{z_s}{(1+z_g)^2 - z_s^2}\delta_J, \\
p_J^{SO} &= \frac{-c/2}{(1+z_g+z_s)^2} - \frac{1+z_g}{(1+z_g)^2 - z_s^2}\delta_J + \frac{z_s}{(1+z_g)^2 - z_s^2}\delta_I, \\
U_I^{SO} &= \frac{c^2/4}{(1+z_g+z_s)^2} + c\delta_I\frac{1+z_g}{(1+z_g)^2 - z_s^2} - c\delta_J\frac{z_s}{(1+z_g)^2 - z_s^2}, \\
U_J^{SO} &= \frac{c^2/4}{(1+z_g+z_s)^2} + c\delta_J\frac{1+z_g}{(1+z_g)^2 - z_s^2} - c\delta_I\frac{z_s}{(1+z_g)^2 - z_s^2}, \\
U_{g,s}^{SO} &= \frac{n}{2}(U_I^{SO} + U_J^{SO}) = n\frac{c^2}{4}\frac{1}{(1+z_g+z_s)^2} + \frac{nc(\delta_I+\delta_J)/2}{1+z_g+z_s}
\end{aligned}
$$

The loss from anarchy in this case is given by:

$$
\nabla_{z_g,z_s}^{(n)} := U_{g,s}^{SO} - U_{g,s}^{NE} = n\frac{c^2}{4}\left(\frac{z_g+z_s}{1+z_g+z_s}\right)^2, \tag{S62}
$$

which is the same as $\nabla_{z'}^{(n)}$ in (S60) with $z' = z_g + z_s$.

S6.3.1 Model Calibration

Fig. 4 in the main text reports the equilibrium policy choices p, resulting change in mobility $\hat{r} - \bar{r}$, and utility U for a pair of states with similar but not identical target mobility reduction targets $\delta_1 = 1.05$ and $\delta_2 = .95$. The pair of states have symmetrical spillovers z plotted on the

horizontal in both panels. The cost of mobility reduction policies c is set to 1 in panel A, and is plotted on the vertical axis.

But what is a realistic value for z, and what can that value tell us about how large the price of anarchy will be? Fig. S38 plots a CDF of the spillover sizes for states in the US. The plot is restricted to 51 state pairs: for each state only the spillovers from their most influential outside state is plotted. Spillovers, z, in this figure are calculated to be compatible with the parameter of the model. They are defined as the ratio "Total change in alter state's social distancing due to ego state implementing shelter-in-place/Total change in ego state's social distancing due to ego state implementing shelter-in-place." We transform the spillover sizes for state pairs in Fig. S38, under the asymptotic ($c \to \infty$) PoA equation in (S48), to obtain the cumulative distribution of PoAs, plotted in Fig. S40. The minimum, median, mean, and maximum values are reported in Tables S11.

Table S11: Spillover sizes and PoAs for state pairs

	minimum	median	mean	maximum
z	0.220	0.421	0.446	0.834
PoA	1.051	1.215	1.248	3.285

S7 Software

Data processing, analysis, and plotting was conducted in R (*47*) and Python (*48*). `pandas` (*49*), `jsonlite` (*50*), and various `tidyverse` libraries (*51*)—`dplyr`, `lubridate`, `readr`, `stringr`, `tidyr`, etc.—were used to process and prepare the data for analysis. Regression analysis was performed using `lfe` (*52*). Other statistical analysis were conducted using `xgboost` (*53*), `glmnet` (*54*), `DirichletMultinomial` (*55*), `aod` (*56*), and `Hmisc` (*57*). `doMC` (*58*) and `iterators` were used to parallelize computation. Tables were created using

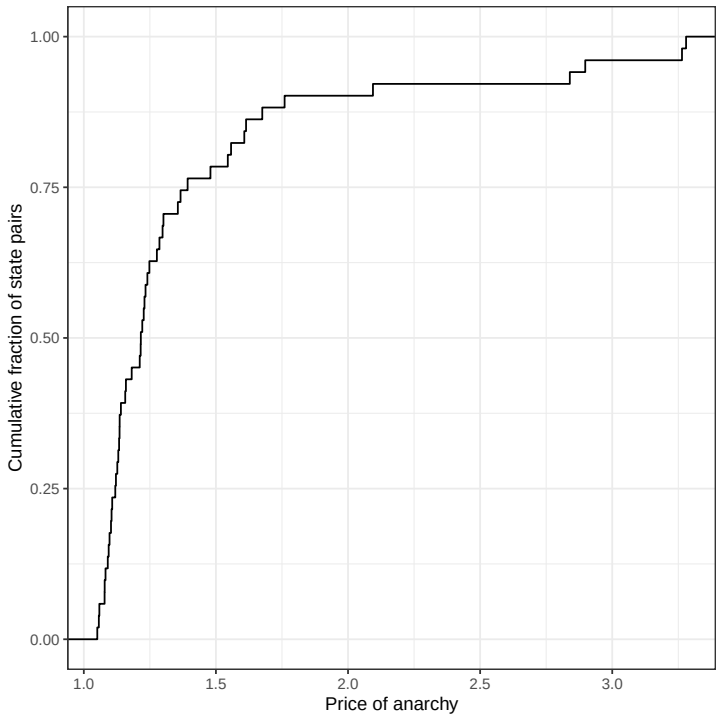

Figure S40: The CDF of spillover sizes (z) in Fig. S38 transformed using the asymptotic PoA equation in (S48), $PoA(z) = 1/(1 - z^2)$.

the `stargazer` package (*59*). Plots were generated using `ggplot2` (*60*), `matplotlib` (*61*), `cowplot` (*62*), `urbnmapr`, and `performanceAnalytics` (*63*).

References

1. J. B. Burbidge, L. Magee, A. L. Robb, *Journal of the American Statistical Association* **83**, 123 (1988).

2. A. Agresti, D. B. Hitchcock, *Statistical Methods & Applications* **14**, 297 (2005).

3. M. Bailey, R. Cao, T. Kuchler, J. Stroebel, A. Wong, *Journal of Economic Perspectives* **32**, 259 (2018).

4. M. J. Menne, I. Durre, R. S. Vose, B. E. Gleason, T. G. Houston, *Journal of atmospheric and oceanic technology* **29**, 897 (2012).

5. B. D. Killeen, *et al.* (2020).

6. D. R. Cox, *Planning of Experiments* (Wiley, 1958).

7. M. G. Hudgens, M. E. Halloran, *Journal of the American Statistical Association* **103**, 832 (2008).

8. P. M. Aronow, C. Samii, *et al.*, *The Annals of Applied Statistics* **11**, 1912 (2017).

9. P. M. Aronow, D. Eckles, C. Samii, S. Zonszein, *arXiv preprint arXiv:2001.05444* (2020).

10. S. Athey, D. Eckles, G. W. Imbens, *Journal of the American Statistical Association* **113**, 230 (2018).

11. L. Keele, R. Titiunik, J. R. Zubizarreta, *Journal of the Royal Statistical Society: Series A (Statistics in Society)* **178**, 223 (2015).

12. O. Ashenfelter, D. Card, *The Review of Economics and Statistics* **67**, 648 (1985).

13. J. D. Angrist, J.-S. Pischke, *Mostly Harmless Econometrics: An Empiricist's Companion* (Princeton University Press, 2008).

14. R. A. Fisher, *The Design of Experiments* (Oliver and Boyd, Edinburgh, 1935).

15. G. W. Imbens, D. B. Rubin, *Causal Inference in Statistics, Social, and Biomedical Sciences* (Cambridge University Press, 2015).

16. J. G. MacKinnon, M. D. Webb, *Journal of Econometrics* (2020).

17. J. Bowers, M. M. Fredrickson, C. Panagopoulos, *Political Analysis* **21**, 97 (2013).

18. E. Chung, J. P. Romano, *et al.*, *The Annals of Statistics* **41**, 484 (2013).

19. J. Wu, P. Ding, *arXiv preprint arXiv:1809.07419* (2018).

20. C. F. Manski, *The Review of Economic Studies* **60**, 531 (1993).

21. W. A. Brock, S. N. Durlauf (Elsevier, 2001), vol. 5 of *Handbook of Econometrics*, pp. 3297 – 3380.

22. S. N. Durlauf, H. Tanaka, *Economic Inquiry* **46**, 25 (2008).

23. L. Coviello, *et al.*, *PloS one* **9** (2014).

24. S. Aral, C. Nicolaides, *Nature communications* **8**, 1 (2017).

25. S. Aral, M. Zhao, *Available at SSRN 3328864* (2019).

26. T. J. Bartik (1991).

27. J. I. Dingel, B. Neiman, How many jobs can be done at home?, *Tech. rep.*, National Bureau of Economic Research (2020).

28. V. Chernozhukov, *et al.*, Double/debiased machine learning for treatment and structural parameters (2018).

29. R. Frisch, F. V. Waugh, *Econometrica: Journal of the Econometric Society* pp. 387–401 (1933).

30. T. Chen, C. Guestrin, *Proceedings of the 22nd acm sigkdd international conference on knowledge discovery and data mining* (2016), pp. 785–794.

31. R. Tibshirani, *Journal of the Royal Statistical Society: Series B (Methodological)* **58**, 267 (1996).

32. A. Belloni, D. Chen, V. Chernozhukov, C. Hansen, *Econometrica* **80**, 2369 (2012).

33. J. C. Chao, N. R. Swanson, *Econometrica* **73**, 1673 (2005).

34. C. Hansen, J. Hausman, W. Newey, *Journal of Business & Economic Statistics* **26**, 398 (2008).

35. J. A. Hausman, W. K. Newey, T. Woutersen, J. C. Chao, N. R. Swanson, *Quantitative Economics* **3**, 211 (2012).

36. J. D. Angrist, A. B. Krueger, Estimating the payoff to schooling using the vietnam-era draft lottery, *Tech. rep.*, National bureau of economic research (1992).

37. G. W. Imbens, P. R. Rosenbaum, *Journal of the Royal Statistical Society: Series A (Statistics in Society)* **168**, 109 (2005).

38. A. D. Cooperman, *Political Analysis* **25**, 277 (2017).

39. V. Chernozhukov, C. Hansen, *Economics Letters* **100**, 68 (2008).

40. T. G. Conley, *Journal of econometrics* **92**, 1 (1999).

41. D. Eckles, R. F. Kizilcec, E. Bakshy, *Proceedings of the National Academy of Sciences* **113**, 7316 (2016).

42. T. G. Conley, F. Molinari, *Journal of Econometrics* **140**, 76 (2007).

43. H. H. Kelejian, I. R. Prucha, *Journal of Econometrics* **140**, 131 (2007).

44. A. C. Cameron, J. B. Gelbach, D. L. Miller, *Journal of Business & Economic Statistics* **29**, 238 (2011).

45. R. H. Coase, *Classic papers in natural resource economics* (Springer, 1960), pp. 87–137.

46. E. Koutsoupias, C. Papadimitriou, *Annual Symposium on Theoretical Aspects of Computer Science* (Springer, 1999), pp. 404–413.

47. R. C. Team, *et al.* (2013).

48. G. Rossum (1995).

49. W. McKinney, *et al.*, *Proceedings of the 9th Python in Science Conference* (Austin, TX, 2010), vol. 445, pp. 51–56.

50. J. Ooms, *arXiv preprint arXiv:1403.2805* (2014).

51. H. Wickham, *et al.*, *Journal of Open Source Software* **4**, 1686 (2019).

52. S. Gaure, *The R Journal* **5**, 104 (2013).

53. T. Chen, *et al.*, *xgboost: Extreme Gradient Boosting* (2019). R package version 0.90.0.1.

54. N. Simon, J. Friedman, T. Hastie, R. Tibshirani, *Journal of statistical software* **39**, 1 (2011).

55. M. Morgan, *R package* (2014).

56. M. Lesnoff, R. Lancelot, M. R. Lancelot, M. Suggests, Package 'aod' (2010).

57. F. E. Harrell, C. Dupont, *et al.*, *R package version* **3** (2008).

58. R. Analytics, *R package version* **1** (2014).

59. M. Hlavac, *R package version* **5** (2015).

60. H. Wickham, *ggplot2: elegant graphics for data analysis* (Springer, 2016).

61. J. D. Hunter, *Computing in science & engineering* **9**, 90 (2007).

62. C. O. Wilke, M. C. O. Wilke, *Streamlined Plot Theme and Plot Annotations for 'ggplot2* (2019).

63. B. G. Peterson, *et al.*, *R Team Cooperation* (2018).