



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}

^aSloan School of Management, Massachusetts Institute of Technology, Cambridge, MA 02142; ^bInitiative on the Digital Economy, Massachusetts Institute of Technology, Cambridge, MA 02142; ^cArgyros School of Business and Economics, Chapman University, Orange, CA 92866; ^dDepartment of Industrial Engineering, University of Pittsburgh, Pittsburgh, PA 15261; ^eMcCombs School of Business, The University of Texas at Austin, Austin, TX 78712; ^fComputer Science, Harvard University, Cambridge, MA 02138; ^gMedia Lab, Massachusetts Institute of Technology, Cambridge, MA 02139; ^hSchool of Information, University of Michigan, Ann Arbor, MI 48109; and ⁱSchool 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.

Published under the PNAS license.

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

²To 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*.

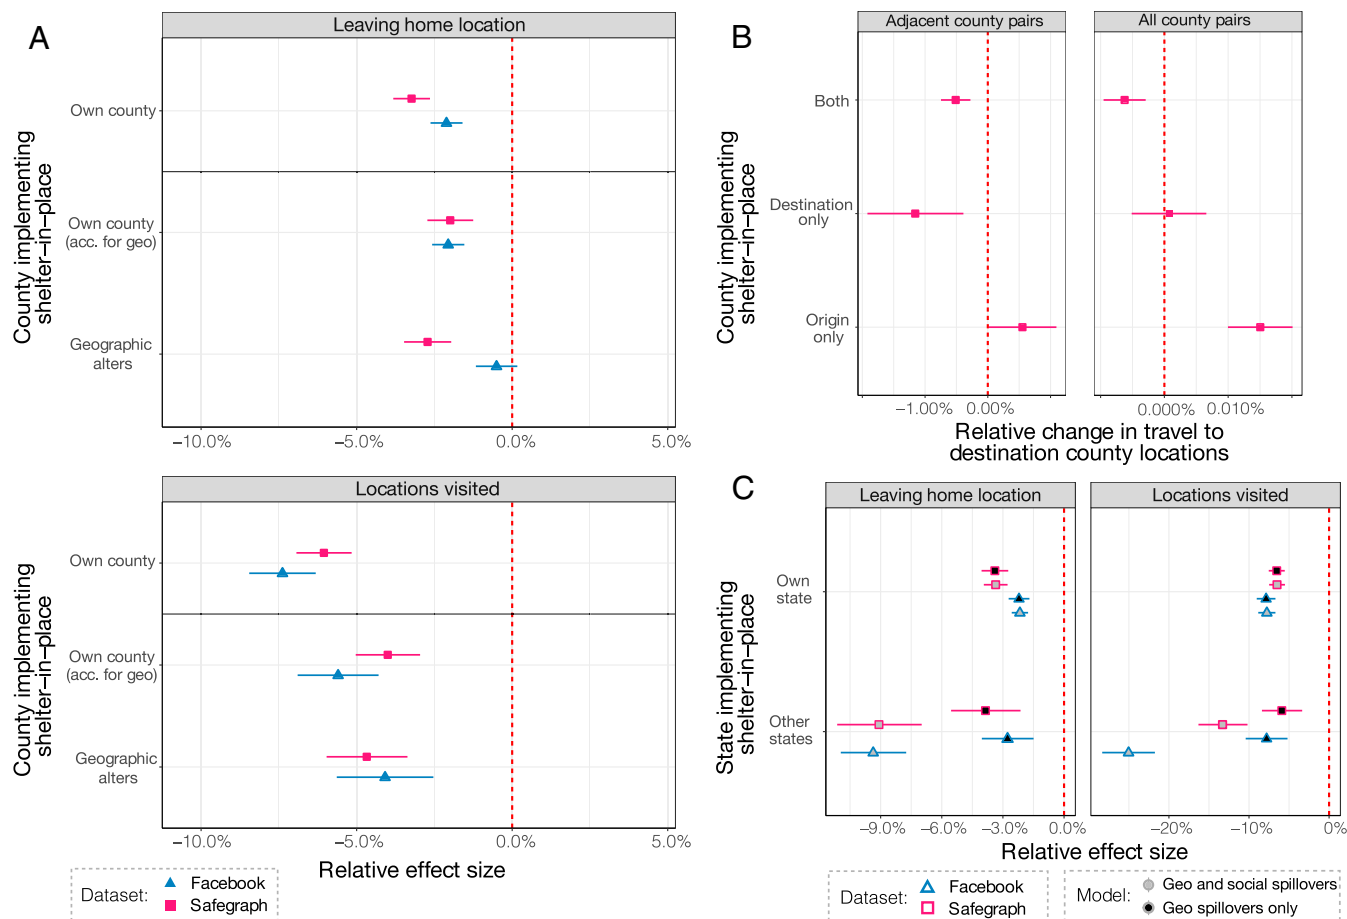


Fig. 2. (A) Comparison of the results of our DiD model that ignores spillovers and estimates the effect of a policy on its “own county” and our DiD model that includes geographic spillovers and separates the effects of a policy on its own county (“own county [acc. for geo]”) from the effects of the policies 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, sec-*

tion 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. 3A). 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. 3C 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

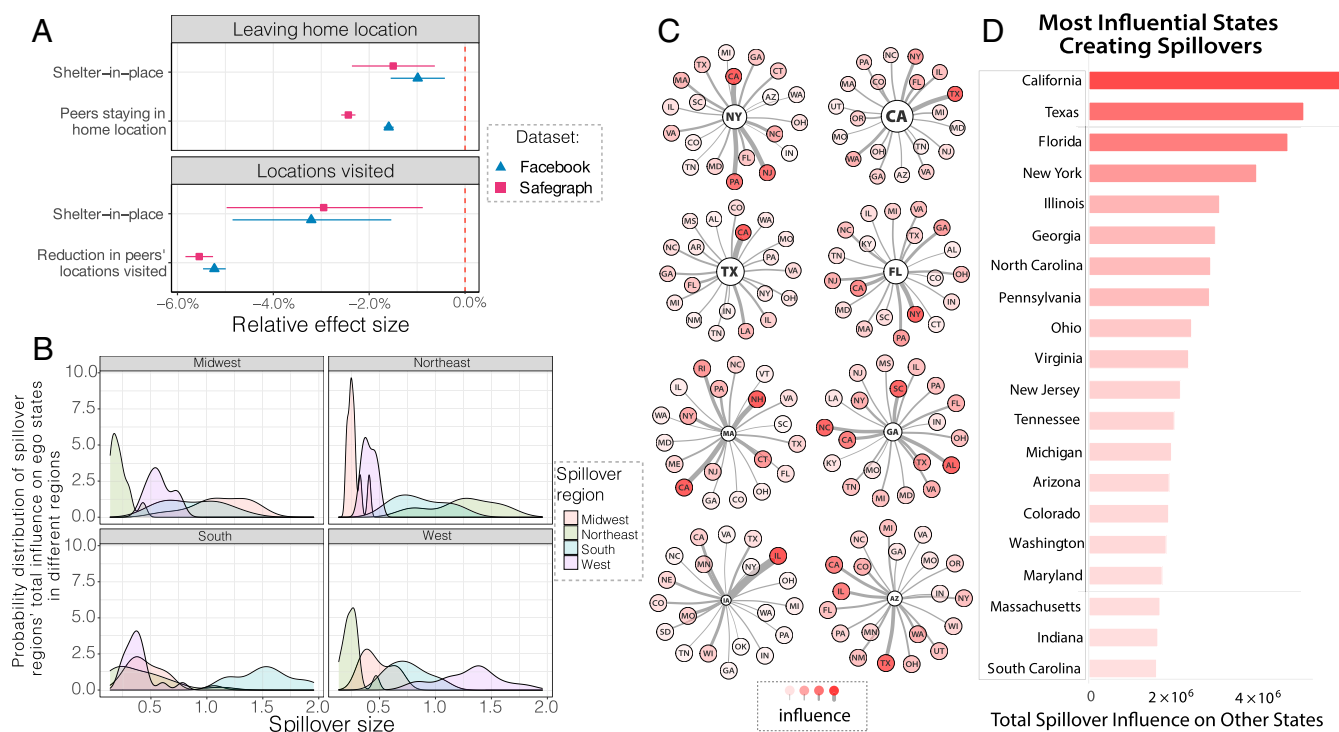


Fig. 3. (A) The causal effect of endogenous mobility levels in Facebook alter counties on mobility levels in focal counties estimated using an IV framework. The magnitudes of endogenous peer effects are scaled to the direct effect of shelter-in-place orders in their own state. (B) The region-level probability 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. 3C). 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 state's 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. 4A). 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

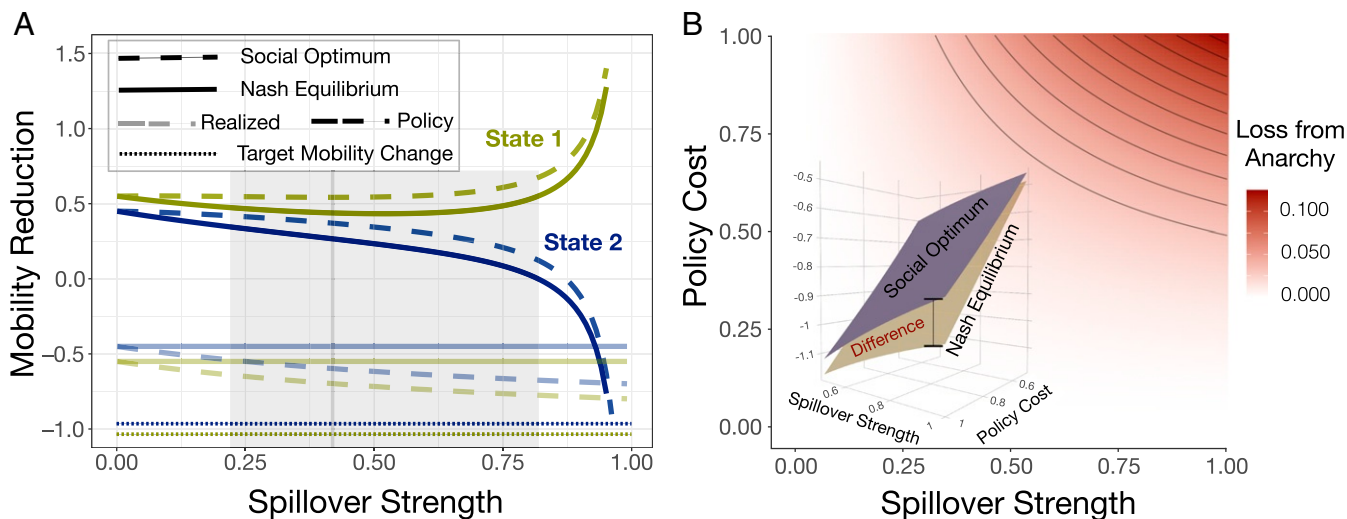


Fig. 4. Mobility reduction targets, optimal policy choices, equilibrium mobility reductions, and utility under anarchy (Nash Equilibrium) and coordination (Social Optimum). (A) Mobility outcomes achieved (faint) under varying levels of spillover strength (x axis) for a pair of states with similar but 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. 2C). (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:

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

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_{-it}^{geo} is the geographic adjacency weighted average of peer county policies, D_{-it}^{social} 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:

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

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

where Y_{it} , D_{it} , D_{-it}^{geo} , D_{-it}^{social} , $f(W_{it})$, 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 repre-

sents 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/mfzha0/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 Safe-graph 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.

- J. S. Jia *et al.*, Population flow drives spatio-temporal distribution of COVID-19 in China. *Nature* **582**, 389–394 (2020).
- 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).
- J. Adda, Economic activity and the spread of viral diseases: Evidence from high frequency data. *Q. J. Econ.* **131**, 891–941 (2016).
- 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).
- R. Florida, T. Gulden, C. Mellander, The rise of the mega-region. *Camb. J. Reg. Econ. Soc.* **1**, 459–476 (2008).
- 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.
- C. O. Buckee *et al.*, Aggregated mobility data could help fight COVID-19. *Science* **368**, 145.2–146 (2020).
- 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.
- 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.
- 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).
- N. Oliver *et al.*, Mobile phone data for informing public health actions across the COVID-19 pandemic life cycle. *Sci. Adv.* **6**, eab0764 (2020).
- A. L. Olsen, F. Hjorth, Willingness to distance in the COVID-19 pandemic. <https://osf.io/xpwwg2/>. Accessed 7 May 2020.
- L. Chiou, C. Tucker, "Social distancing, internet access and inequality" (Tech. Rep. 26982, National Bureau of Economic Research, 2020).
- S. Aral, C. Nicolaides, Exercise contagion in a global social network. *Nat. Commun.* **8**, 1–8 (2017).
- S. Aral, D. Walker, Identifying influential and susceptible members of social networks. *Science* **337**, 337–341 (2012).
- R. M. Bond *et al.*, A 61-million-person experiment in social influence and political mobilization. *Nature* **489**, 295–298 (2012).
- S. G. Bronars, J. R. Lott, Criminal deterrence, geographic spillovers, and the right to carry concealed handguns. *Am. Econ. Rev.* **88**, 475–479 (1998).
- M. K. Chen, R. Rohla, The effect of partisanship and political advertising on close family ties. *Science* **360**, 1020–1024 (2018).
- 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.
- M. Bailey, R. Cao, T. Kuchler, J. Stroebel, A. Wong, Social connectedness: Measurement, determinants, and effects. *J. Econ. Perspect.* **32**, 259–80 (2018).
- 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).
- B. D. Killeen *et al.*, A county-level dataset for informing the United States' response to COVID-19. arXiv:2004.00756 (1 April 2020).
- A. Brzezinski, G. Deiana, V. Kecht, D. Van Dijke, "The COVID-19 pandemic: Government vs. Community action across the United States" (Tech. Rep. 2020-06, Institute for New Economic Thinking, 2020).
- A. Goodman-Bacon, "Difference-in-differences with variation in treatment timing" (Report 25018, National Bureau of Economic Research, 2018).
- 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.
- V. Chernozhukov *et al.*, Double/debiased machine learning for treatment and structural parameters. *Econom. J.* **21**, C1–C68 (2018).
- T. G. Conley, GMM estimation with cross sectional dependence. *J. Econom.* **92**, 1–45 (1999).
- L. Coviello *et al.*, Detecting emotional contagion in massive social networks. *PLoS One* **9**, e90315 (2014).
- 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.
- T. J. Bartik, *Who Benefits from State and Local Economic Development Policies?* (W.E. Upjohn Institute for Employment Research, 1991).
- P. Goldsmith-Pinkham, I. Sorkin, H. Swift, "Bartik Instruments: What, when, why, and how" (Report 24408, National Bureau of Economic Research, 2018).
- 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).