



Interacting regional policies in containing a disease

Arun G. Chandrasekhar^{a,b,c,1}, Paul Goldsmith-Pinkham^{d,1}, Matthew O. Jackson^{a,e,1,2}, and Samuel Thau^{f,1}

^aDepartment of Economics, Stanford University, Stanford, CA 94305; ^bAbdul Latif Jameel Poverty Action Lab (J-PAL), Cambridge, MA 02142; ^cNational Bureau of Economic Research (NBER), Cambridge, MA 02138; ^dYale School of Management, Yale University, New Haven, CT 06511; ^eSanta Fe Institute, Santa Fe, NM 87501; and ^fApplied Mathematics, Harvard University, Cambridge, MA 02138

Contributed by Matthew O. Jackson, March 4, 2021 (sent for review October 15, 2020; reviewed by Ozan Candogan, Dean Eckles, and Evan Sadler)

Regional quarantine policies, in which a portion of a population surrounding infections is locked down, are an important tool to contain disease. However, jurisdictional governments—such as cities, counties, states, and countries—act with minimal coordination across borders. We show that a regional quarantine policy’s effectiveness depends on whether 1) the network of interactions satisfies a growth balance condition, 2) infections have a short delay in detection, and 3) the government has control over and knowledge of the necessary parts of the network (no leakage of behaviors). As these conditions generally fail to be satisfied, especially when interactions cross borders, we show that substantial improvements are possible if governments are outward looking and proactive: triggering quarantines in reaction to neighbors’ infection rates, in some cases even before infections are detected internally. We also show that even a few lax governments—those that wait for nontrivial internal infection rates before quarantining—impose substantial costs on the whole system. Our results illustrate the importance of understanding contagion across policy borders and offer a starting point in designing proactive policies for decentralized jurisdictions.

contagion | network | quarantine | lockdown | coordination

Global problems, from climate change to financial crises to disease control, are hard to address without policy coordination across borders. Carbon emissions in one region are everyone’s problem, as are financial collapses, as well as the spread of an infectious disease. Coordinating policies across jurisdictions in terms of both timing and scale is important whenever problems have spillovers. In this paper we shed light on this problem by examining how different types of decentralized policies fare compared to more centralized policies at containing the spread of an infectious disease.

In particular, pandemics, like COVID-19, are challenging to contain if governments fail to coordinate efforts. Without vaccines or herd immunity, governments have responded to infections by limiting constituents’ interactions in areas where an outbreak exceeds a threshold of infections. Such regional quarantine policies are used by towns, cities, counties, states, and countries and trace to the days of the black plague. Over the past 150 y, regional quarantines have been used to combat cholera, diphtheria, typhoid, flus, polio, Ebola, and COVID-19 (1–4), but rarely with coordination across borders.

Decentralized policies across jurisdictions have two major shortcomings. First, governments care primarily about their own citizens and do not account for how their infections impact other jurisdictions: The resulting lack of coordination can lead to worse overall outcomes than a global policy (5–7). Second, many governments are inward looking, paying attention only to internal situations, which leads them to underforecast their own infection rates.

We examine three types of quarantine policies to understand the impact of noncoordination: 1) those controlled by one actor with control of the whole society—“single-regime policies”; 2) those controlled by separate jurisdictions that are inward looking and react only to internal infection rates, or “reactive” for short; and 3) those controlled by separate jurisdictions that are outward

looking, tracking infections outside of their jurisdiction as well as within to forecast their infection rates when deciding when to quarantine, or “proactive” for short.

We use a general model of contagion through a network to study these policies. We first consider single-regime policies. A government can quarantine everyone at once under a “global quarantine,” but those are very costly (e.g., lost days of work, school, etc.). Less costly (in the short run), and hence more common, alternatives are “regional quarantines” in which only people within some distance of observed infections are quarantined. Regional quarantines, however, face two challenges. First, many diseases are difficult to detect, because either some individuals are asymptotically contagious (e.g., HIV, COVID-19) (8–10) or a government lacks resources to quickly identify infections (11, 12). Second, it may be infeasible to fully quarantine a part of the network, because of difficulties in identifying whom to quarantine (e.g., imperfect or inefficient contact tracing) or noncompliance by some people—by choice or necessity (13–18). Either way, tiny leakages can spread the disease.

We show that regional quarantines curb the spread of a disease if and only if 1) there is limited delay in observing infections, 2) there is sufficient knowledge and control of the network to prevent leakage of infection, and 3) the network has a certain “growth balance” structure. The failure of any of these conditions substantially limits regional quarantine effectiveness.

We then examine jurisdictional policies, which are regional quarantine policies conducted by multiple, uncoordinated regimes. The regions that need to be quarantined, however, often cross borders, leading to leakage that limits their effectiveness.

Significance

The lack of coordination between governments on systemic problems, from climate change to financial contagions, has large costs that are well exemplified by a pandemic. We identify which quarantine policies are effective in curbing an outbreak and use that to understand how governments’ prioritization of their own populace and failure to coordinate with neighbors cause quarantines to fail, leading to repeated waves of a disease. Instead, if governments share information and anticipate infections from other jurisdictions, then infections are reduced by an order of magnitude. However, just a few lax governments create jurisdictions that repeatedly incubate a disease and substantially worsen the outcome for everyone. Our techniques should be useful in studying networked interactions across other domains.

Author contributions: A.G.C., P.G.-P., M.O.J., and S.T. designed research, performed research, and wrote the paper.

Reviewers: O.C., University of Chicago; D.E., Massachusetts Institute of Technology; and E.S., Columbia University.

The authors declare no competing interest.

Published under the [PNAS license](#).

¹A.G.C., P.G.-P., M.O.J., and S.T. contributed equally to this work.

²To whom correspondence may be addressed. Email: jacksonm@stanford.edu.

This article contains supporting information online at <https://www.pnas.org/lookup/suppl/doi:10.1073/pnas.2021520118/-DCSupplemental>.

Published April 27, 2021.

As we show, jurisdictional policies that are reactive do much worse than proactive ones, as they do not forecast the impact of neighboring jurisdictions' infection rates on their own population. Moreover, a few lax jurisdictions, which wait for higher infection rates before quarantining, worsen outcomes for all jurisdictions.

A Model

Consider a large network of n individuals or nodes. Our theory is asymptotic, stating properties that apply with a probability approaching one as the population grows ($n \rightarrow \infty$). Our theoretical results consider sequences of networks as n grows, while the simulations are on given networks with thousands of nodes.

An infectious disease begins with an infection of a node i_0 , the location of which is known, and expands via (directed) paths from i_0 . In each discrete time period, the infection spreads from each currently infected node to each of its susceptible contacts independently with probability p . A node is infectious for θ periods, after which it recovers and is no longer susceptible or contagious, although our results extend to the case in which a node can become susceptible again.

The disease may exhibit a delay of $\tau \leq \theta$ periods during which an infected and contagious person does not test positive. This can be a period of asymptomatic infectiousness, a delay in testing, or limits to healthcare access (8, 10–12, 19, 20). After that delay, each infected node's infection is detected with probability $\alpha < 1$ (for simplicity, in the first period after the delay). α incorporates testing accuracy, availability, and decisions to test.

This framework nests the susceptible–infected–recovered (SIR) model and its variations including exposure, multiple infectious stages, and death (18, 21–24); agent-based models (25–27); and others.

Results

Baseline: A Single Jurisdiction with Complete Control. We begin by analyzing a single jurisdiction with complete control or, equivalently, the entire network being one jurisdiction with one policymaker.

A (k, x) -regional policy is triggered once x or more infections are observed within distance k from the seed node i_0 , at which point it quarantines all nodes within distance $k + 1$ of the seed for θ periods. This captures a commonly used policy where regions that are exposed to the disease are shut down in response to detection. We begin by giving the policymaker the advantage of knowing which nodes are within distance $k + 1$ of the seed, which could reflect rapid and efficient contact tracing supplemented with rich network data. We later explore how errors in this knowledge change the results. We also give the policymaker knowledge of which node is the seed and study subsequent containment efforts. In practice, policymakers must estimate the origin of infection, which presents an additional challenge.

Whether a regional policy halts infection is fully characterized by whether the sequence of networks satisfies what we call growth balance. This requires that as the networks become large, all paths along which the disease might escape beyond the regional quarantine of distance $k + 1$ are such that at least some nodes have many neighbors. In particular, for the sequence of networks to be growth balanced with respect to k , there must exist $m(n) \rightarrow \infty$ such that in a network with n nodes, every path leading from i_0 to a node at distance $k + 2$ has at least one node with degree at least $m(n)$. This condition ensures that if an infection does spread along some path that can take it outside of the region, then, at some point along the way, it is very likely to infect many nodes and thus be detected before it reaches the edge of the region.

To better understand growth balance, consider an example of a disease that is beginning to spread with a reproduction number R_0 of 3.5 and such that 1 in 10 cases is detected in a timely man-

ner ($\alpha = 0.1$). First, consider a part of the network in which each infected person infects 3.5 others on average. If we monitor all nodes within distance $k = 3$ of an infected node, a “typical” path of infection would lead to roughly $3.5 + 3.5^2 + 3.5^3 = 58.625$ expected cases before it reaches the edge of the region. The chance that this goes undetected is tiny: $0.9^{58.625} = 0.002$. In contrast, suppose the infection starts in a part of the network where each infected person infects just one other, on average, so that the local reproduction number here is $R_0 = 1$ rather than 3.5. Now a path of length 3 leads to $1 + 1 + 1 = 3$ (expected) infections. The chance that such a spread remains undetected is much higher $0.9^3 = 0.72$.

Speaking loosely, many different networks can lead to the same average reproduction number, but have very different structures. If the distribution of reproduction numbers around the network has no pockets in which they are too low—i.e., if the growth structure of the disease around the network is well “balanced” and not too low—then it is highly likely that any early infection will be detected before it gets too far from the first infected node. If, instead, the distribution of reproduction numbers gives a nontrivial chance that the disease starts out on a path with all low reproduction numbers, like the 1, 1, 1, path, then there is a high chance that it can travel far from the seed before being detected. This highlights the fact that a reproduction number R_0 alone is a crude concept and that the specifics of the network structure matter considerably for whether a disease spreads or is containable. In particular, areas with low R_0 (but above one) can lead to more containment failures and lead to broader infections. Given the short distances in many networks (28–30), a lack of growth balance allows a disease to spread far before detection. *SI Appendix, Fig. S1* pictures a network that has a high average reproduction number, but is not growth balanced and allows the infection to travel far from the initially infected node without detection.

In *SI Appendix, Theorem 1* we prove that, with no delays in detection and no leakage, a (k, x) -regional policy halts infection among all nodes beyond distance $k + 1$ from i_0 with probability approaching 1 (as the population grows) if and only if the sequence of networks satisfies growth balance with respect to k . In fact, we prove a stronger version in which the quarantine distance $k(n)$, average degree $d(n)$, transmission probability $p(n)$, delay $\theta(n)$, and detection probability $\alpha(n)$ are all allowed to vary with n .

Growth balance is satisfied by some, but not all, sequences of prominent random graph models, provided that the average degrees $d(n)$ satisfy $d(n)^{k(n)} \rightarrow \infty$ (*SI Appendix*). However, the additional heterogeneity in human contact networks makes the property unlikely to hold in real networks even if average connectivity is high. Indeed, if contact networks have some low degree nodes (as they tend to empirically), then, unless quarantine regions are large ($k(n)$ grows sufficiently as n grows), growth balance fails and a regional quarantine will be ineffective at halting a spread.

Next, we show that the effectiveness of a regional policy breaks down, even if a network is growth balanced, once there is sufficient delay in detection or leakage (due to imperfect information, enforcement, or jurisdictional boundaries).

Delays in Detection and Wider Quarantines. To understand how delays in detection affect a regional policy, consider two extremes. If the delay is short relative to the infectious period, the policymaker can still anticipate the disease and adjust by simply enlarging the area of the quarantine to include a buffer. An easy extension of the above theorem is that a regional policy with a buffer works if and only if the sequence of networks is growth balanced and the delay in detection plus $k + 1$ is shorter than the diameter of the network (*SI Appendix, Theorem 2*). Given that real-world networks have short average distances between

nodes (31), this condition can be even harder to satisfy and non-trivial delays in detection allow the disease to escape a regional quarantine.

Leakage. Next, we consider how leakage—the inability to limit interactions (13) or mistakes in identifying the portion of the network and nodes to quarantine (17, 18)—diminishes the effectiveness of regional policies. Although minimizing leakage increases the chance that a regional quarantine will be successful, we show that even a small amount of leakage leads to a nontrivial probability that the disease will escape the quarantine. In particular, we show that if even a small share $\varepsilon > 0$ of nodes within distance k of the seed i_0 ignore the quarantine and are connected to nodes outside of the quarantine, then the policy will fail to halt the spread with a probability bounded away from 0 as n grows (*SI Appendix, Theorem 3, part 2*). The result also highlights a tension in containment strategies: The infection is easier to detect when there are many nodes and interactions within the potential quarantine radius; but there is also more leakage and a higher chance that the infection escapes the quarantine.

Jurisdictions and Leakage. We use the theory results as a starting point to understand jurisdictional policies. It is important to note that the results on leakage (*SI Appendix, Theorem 3*) apply when interactions cross jurisdictions. To show this, Fig. 1 displays two jurisdictions that fail to nicely tessellate the network: Geographic location and network distance from the seed are not perfectly aligned. Therefore, a quarantine in one region will necessarily have leakage, missing nodes that interact across jurisdictions. Given leakage across borders, unless policies are coordinated across jurisdictions, our theoretical results indicate that they will fail to contain infections, which is then the starting motivation of the simulations.

Simulations. The theory shows that for most natural settings, anything short of a global quarantine is unlikely to contain the disease. Thus, it becomes important to understand how well different policies do at curbing the number of infections over time. In particular, we next study—via simulations—how stylized versions of the containment policies that are used in practice fare in terms of minimizing infections over time and at what costs in terms of person days of quarantine (a crude measure that accounts only for the total and not temporal or geographic incidences, which might also be important considerations).

To explore this, we simulate a contagion on a network of 140,000 nodes that mimics real-world data (20, 32–34). These

simulations illustrate our theoretical results and also show the improvements that proactive policies provide relative to reactive ones. *SI Appendix, section 2E* presents the robustness of the results to some variations of parameters.

The network is divided into 40 locations, each with a population of 3,500. We generate the network using a geographic stochastic block model. The probability of interacting declines with distance. The average degree is 20.49 and nodes have 79.08% of their interactions within their own locations and 20.92% outside of their locations [calibrated to data from India and the United States, including data collected during COVID-19 (*SI Appendix, section 2A*) (20, 32–34)]. We fix this network and use it for all simulations.

We conduct 10,000 simulations of each policy and then take the average over the simulations, with each simulation using an infection seed selected uniformly at random. The simulations progress in four stages: First, any node that has been infected for exactly τ periods is detected with probability α ; next, policymakers use the information they have about detected infections to decide whether to enact a quarantine (if one is not already in place in their jurisdiction); third, the disease progresses and currently infected people can infect their neighbors and people who have been infected for θ periods recover; finally, quarantines can end and new quarantines are implemented. We set the rate that a node infects its neighbors to get a basic reproduction rate of $R_0 = 3.5$ (to mimic COVID-19) (35), and we set $\theta = 5$, $\tau = 3$ (when used), and $\alpha = 0.1$ (*SI Appendix, section 2B*) (19, 36, 37).

The simulated network is fairly symmetric in degree and therefore approximates satisfying growth balance. Thus, the attention in our simulations is focused on leakage across jurisdictions and detection delay.

Before introducing jurisdictions, we first illustrate the effects of leakage as well as delays in detection on a regional policy. In Fig. 2, the entire network is governed by a single policymaker using a $(k, x) = (3, 1)$ -regional quarantine. The policymaker eventually knows the location of the infection seed, so that it can properly center the quarantine, but does not detect the initial infection. This is meant to emulate the difficulties of finding an initial infection in real time, but we give the policymaker the advantage of being able to trace back to the epicenter and center the quarantine once the policymaker decides to enact a quarantine. As an addition to the policy, if the initial (k, x) quarantine fails to contain the disease, the policymaker treats detected infected people outside of quarantine as new seeds and quarantines all nodes within distance $k + 1$ of them. In *SI Appendix*, we include simulations that relax the assumption that the policymaker knows the location of the original seed i_0 , along

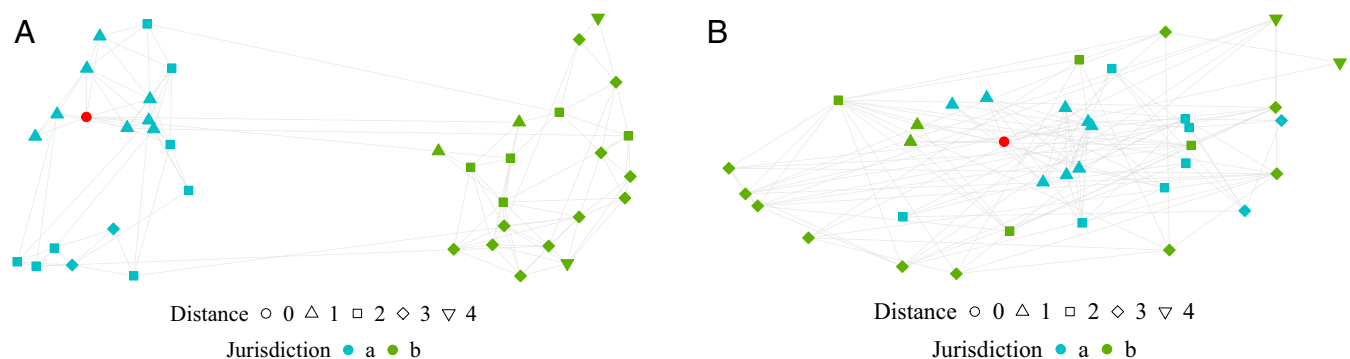


Fig. 1. Nodes in two jurisdictions do not align with the distances from the initial infection. In *A*, the nodes are presented in a geographic sense, within their jurisdictions, and the interaction network does not comply with the jurisdictional boundaries. In *B*, we show the network as a function of directed distance from the initial infection. A coordinated quarantine of distance 2 over the network in *B* could contain the infection; however, if it is executed only by the infected node's jurisdiction in *A*, then it would fail for cross-jurisdictional connections.

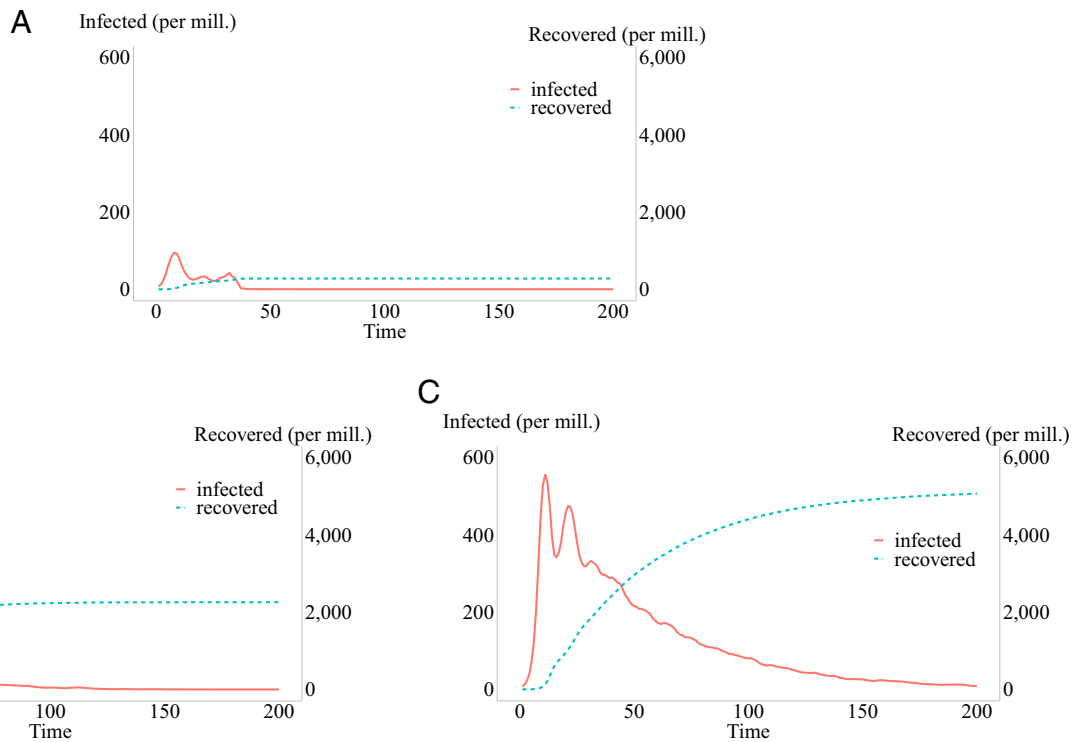


Fig. 2. We picture daily infections and cumulative recoveries under three scenarios. The entire network is governed by a single policymaker using a $(k, x) = (3, 1)$ -regional quarantine. In *A*, there is no detection delay and no leakage. In *B*, we introduce a detection delay of $\tau = 3$. This represents the 3-d presymptomatic window during which an infected node can transmit, as well as an expected delay in seeking healthcare and testing upon symptom onset (*SI Appendix, section 2B*) (19). *C* adds leakage to the setup of *B*, by having a randomly selected fraction $\epsilon = 0.05$ never quarantine. For each plot, we simulate 10,000 times on the same network with random initial infections and present the average number of infections and recovered people over time, scaled per million.

with variants of the transmission and detection parameters—varying θ , τ , α , and R_0 (*SI Appendix, sections 2D and 2E*). In our simulations, we primarily focus on $x = 1$ and $k = 3$ (other parameters appear in *SI Appendix, section 2E*). We use $x = 1$ because, as we discuss in *SI Appendix*, a policymaker attempting to minimize the number of infected person periods and quarantined person periods should attempt to minimize the number of infections. We choose $k = 3$ due to the limited size of our simulated network—a larger boundary would cover too much of the overall network.

Fig. 2*A* shows the outcomes for no delay in detection nor any leakage. Consistent with *SI Appendix, Theorem 1* the policy is effective: On average 277 people per million are infected (0.028% of the population), with 803,956 person days of quarantine per million people. Fig. 2*B* introduces a delay in detection. With a delay of $\tau = 3$, infections increase, with 2,256 people per million eventually infected (0.23% of the population) and 2,301,414 person days of quarantine per million people. Adding a buffer to correspond to the detection delay effectively makes the regional policy global, as the buffered region contains 99.98% of the population on average. Fig. 2*C* adds leakage to the setup of Fig. 2*B*, by having 5% of people never quarantine. The number of cumulative infections per million people increases to 5,138 (0.50% of the population). The leakage increases the number of quarantined person days to 6,478,055 per million nodes.

Jurisdictional Policies. We now introduce jurisdictions to the same network as before, and each of the 40 locations in the network becomes its own jurisdiction.

We compare two types of jurisdictional policies. In reactive policies, each jurisdiction decides on when it quarantines based entirely on internal infections. If a jurisdiction has a threshold

of x cases, and observes at least x cases within the jurisdiction, the jurisdiction goes into quarantine. In proactive policies, jurisdictions track infections in other jurisdictions and predict their own—possibly undetected—infections and base their quarantines off of predicted infections. Intuitively, jurisdictions are constantly estimating infection rates (including undetected cases) in each jurisdiction based on the history of observed infections, using knowledge of the interaction rates within and across borders and the infection and latency properties of the disease. More specifically, at each time step t , the number of estimated infections at time t is the sum of the estimated infections at $t - 1$, plus the expected number of new estimated infections minus the number of expected recoveries. The expected number of new infections is calculated using the connection rates to nonquarantined jurisdictions (including own jurisdiction) and the estimated infections in those jurisdictions at $t - 1$. If at any point it is clear that the actual number of detected infections in some jurisdiction is above the estimated rate, then the estimation is updated. All jurisdictions begin by estimating that there are no infections, until at least one infection is observed. Details of this calculation are in *SI Appendix, section 2D*. Proactive jurisdictions quarantine if they infer (or observe) at least x cases.

We set $x = 1$ for both the reactive and proactive simulations unless otherwise specified. For both reactive and proactive jurisdictions, when a jurisdiction enters quarantine, all connections to and within the jurisdiction are severed for θ periods. As before, we set $R_0 = 3.5$, $\theta = 5$, $\tau = 3$, and $\alpha = 0.1$. The policymaker does not detect i_0 for both the reactive and proactive policies, but does know its location when setting the quarantine.

Fig. 3 illustrates the improvement that proactive jurisdictional policies offer relative to reactive jurisdictional policies.

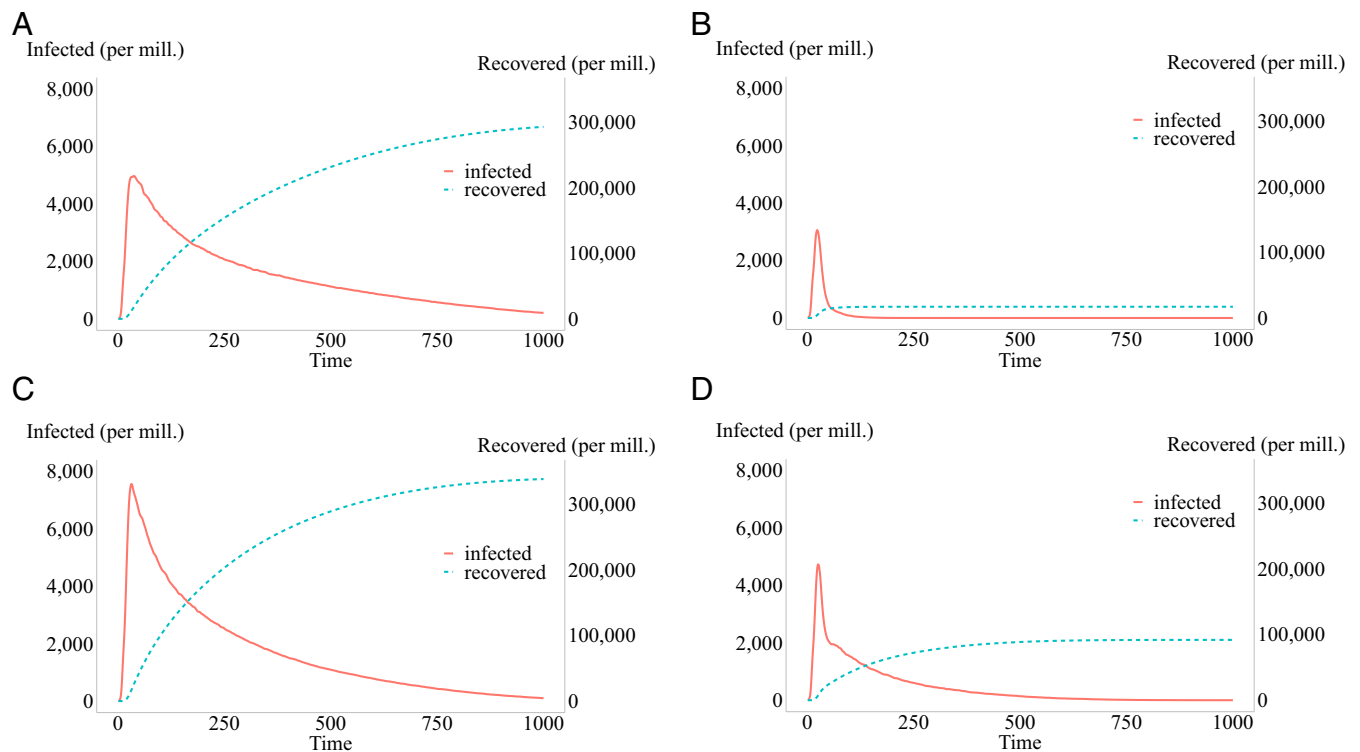


Fig. 3. We picture daily infections and cumulative recoveries under four quarantine policies with 40 jurisdictions. When a jurisdiction quarantines, it locks down the entire jurisdiction. In *A*, all jurisdictions use a reactive policy. In *B*, all jurisdictions use a proactive policy. In *C*, we implement the same policies as in *A*, but have four lax jurisdictions that use $x = 5$ (0.14% of the jurisdiction population) instead of $x = 1$. *D* has 36 jurisdictions with proactive policies and 4 with lax policies. For each plot, we simulate 10,000 times on the same network with random initial infections and present the average number of infections and recovered people over time, scaled per million.

In Fig. 3*A*, jurisdictions use reactive policies, while in Fig. 3*B* jurisdictions use proactive policies. In the reactive case, there are 298,911 infections per million people (28.89% of the population), with 131,303,638 person days of quarantine per million nodes. Proactive quarantining dramatically improves outcomes (Fig. 3*B*): Only 17,105 people per million are infected (1.71% of the population), with 51,328,755 person days of quarantine per million people.

Lax Jurisdictions. Finally, we also add four “lax” jurisdictions to the setting. These are jurisdictions that are reactive and have a high threshold of internal infections before quarantining, using a threshold of $x = 5$. We examine how these few lax jurisdictions worsen the outcomes for all jurisdictions. Fig. 3*C* shows the outcomes when the remaining 36 jurisdictions use reactive policies, while in Fig. 3*D* the remaining 36 jurisdictions use proactive policies. Comparing Fig. 3*A–C*, infections are worse under the reactive policies. There are 340,587 people per million infected (34.1% of the population), compared to 298,911 (29.9%) without the lax jurisdictions. Comparing this change to Fig. 3*B* and *D* shows that things deteriorate relatively more for the proactive jurisdictional policies. The 91,887 total infections per million people (9.18% of the population) are a larger increase from 17,105 (1.17%) without lax jurisdictions. Nonetheless, even with lax jurisdictions, the proactive policies fare better than the reactive policies (even if those do not have lax jurisdictions).

Fig. 4*A* displays the dynamics of quarantines for each of the policy configurations from Fig. 3, and Fig. 4*B* displays the number of person days of infection versus the number of person days of quarantine. Single jurisdiction policies (global quarantines and (k, x) regional quarantines [with $\varepsilon = 0.05$ leakage]) do the best on both dimensions. Once multiple jurisdictions are intro-

duced, proactive policies perform better than reactive ones for both infected and quarantined person days. Even with lax jurisdictions, the proactive policy is better than the reactive policy without lax jurisdictions. For both proactive and reactive policies, introducing lax jurisdictions increases infected person days. However, this effect is not uniform: With proactive jurisdictions, lax jurisdictions cause a larger increase in infections (both absolutely and proportionally). With respect to quarantined person days, lax jurisdictions have different effects, depending on other jurisdictions’ policies. With proactive jurisdictions, quarantined person days increase, while with reactive jurisdictions, they slightly decrease. With reactive policies, infections spread so rapidly from lax jurisdictions that by coincidence, large numbers of jurisdictions quarantine at once—and eventually an almost global quarantine occurs, halting the disease more quickly than in the scenarios without lax jurisdictions, but with higher infections. This relative ordering of quarantined-person days for reactive policies depends on parameters, as demonstrated in *SI Appendix*: Increasing α causes the number of quarantined person days to increase once lax jurisdictions are added to reactive policies, rather than decrease (*SI Appendix, section 2E*).

Discussion

We have shown that regional quarantine policies are likely to fail to halt the spread of a virus in most empirical settings, unless there is extremely rapid and efficient detection of the disease and governments can halt all contact within the quarantined region. This failure is due to the failure of what we call growth balance, which ensures that there are no infection paths leading from infected individuals to others outside the quarantined region that are likely to be undetected. Multiple governments using independent policies are even less effective, as leakage occurs across

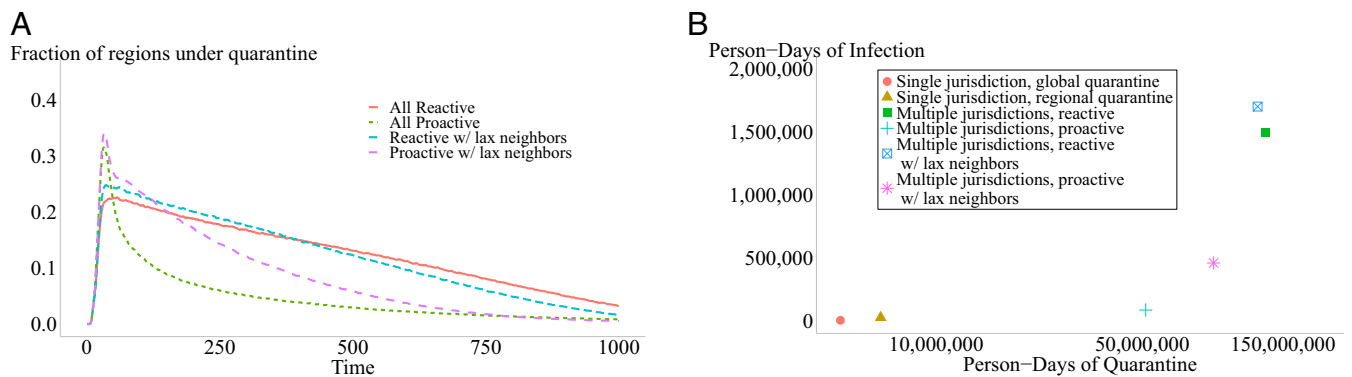


Fig. 4. *A* displays the dynamics of quarantines for each of the policy configurations. *B* plots the number of person-day infections (per million) against the number of person-day quarantines (per million) for six key policy scenarios. The global policy does the best on both dimensions, and the second best is the single-jurisdiction reactive strategy (which does worse than the global one because of leakage). With 40 jurisdictions, both proactive policies outperform the internal, reactive policies. By far the worst, on both dimensions, is the internal, reactive policy with some lax jurisdictions. These results come from the same simulations that produce Figs. 2 and 3.

their borders. We have also shown that if governments are more attentive to their neighbors, there can be substantial improvements to their infection rates. However, the existence of some lax jurisdictions imposes a significant cost on everyone else.

We view the (k, x) quarantine model as a stylized approximation of some policymakers' policies. These policymakers estimate the disease's location and try to shut down the surrounding area. This can range from city blocks to subdistricts to larger regions, depending on the country. For example, in multiple states across India, such as Haryana, Karnataka, Andhra Pradesh, Gujarat, and West Bengal (with a combined population of around 285 million as of the 2011 Indian Census), district officials are entrusted to define containment zones. This is decentralized and many of their efforts are "microtargeted." More globally, each district plays the role of a jurisdiction in the multijurisdiction case. Another example is a sports league where teams or collections of players are contact traced and quarantined—parsimoniously modeled as (k, x) from i_0 .

Jurisdictional policies tend to be aimed at the welfare of their internal populations, yet the external effects are large. Our results underscore the importance of timely information sharing and coordination in both the design and execution of policies across jurisdictional boundaries (38). The results also underscore the global importance of aiding poor jurisdictions. Indeed, there is mounting evidence that a lack of coordination across boundaries has been damaging in the case of COVID-19 (6).

The use of masks (decreasing p), social distancing (decreasing d), and increasing testing (increasing α) and vaccinations (decreasing p) all help attenuate contagion, but unless they maintain the reproduction number below one, the problems identified here remain. Even tiny fractions of interactions across borders are enough to lead to spreading in large populations. With modern inter- and intranational trade and travel being a sizable portion of all economies, such interaction is difficult to avoid. Nonetheless, our analysis also offers insights into managing infections at smaller scales, e.g., within schools, sports, and businesses. By creating a network of interactions that is highly modular, keeping cross-modular interactions to a minimum and making sure that they are highly traceable, together with aggressive testing (especially of cross-module actors), one can eliminate leakage and effectively bound the set of interactions. This would

divide the network into small components of diameter less than k , so that growth balance is satisfied by default.

We note that the effect of lax jurisdictions can be mitigated if other jurisdictions eliminate contact with that jurisdiction, e.g., with travel restrictions. Then, for the nonlax jurisdictions, the situation returns to one without any lax jurisdictions, which is better as we have shown. This will work only if all nonlax jurisdictions participate in a travel ban, as otherwise infection will continue to resurge in jurisdictions that continue contact with the lax jurisdictions, which then pass infection along to others.

Our results also suggest caution in using statistical models to identify regions to quarantine. Although contagion models are helpful for informing policy about the magnitude of an epidemic and broad dynamics, the models can give false comfort in our ability to engage in highly targeted policies, whose results can be influenced by small deviations from idealized assumptions (e.g., leakage). Our growth balance condition also points out that not all parts of a network are equal in their potential for undetected transmission. Growth balance offers insight into when containment will frequently fail. While full containment is sometimes, but not always, the full policy goal, it helps us understand what features of the network aid or hinder containment efforts even under ideal conditions. Policymakers must be conscious of limited paths of interactions that can introduce the disease into a larger population, precisely because detection is very difficult along such paths. Therefore, such paths are important to monitor. In places where the reproduction number is lower, the probability of observing outbreaks is also lower, enabling a leakage of undetected infections.

Data Availability. There are no data underlying this work.

ACKNOWLEDGMENTS. We thank Marcella Alsan, Abhijit Banerjee, Gabriel Carroll, Bharat Chandar, Dean Eckles, Ben Golub, Hans Haller, Dave Holtz, Ali Jadbabaie, Ed Kaplan, Jon Kleinberg, Jacob Leshno, Sudipta Sarangi, Devavrat Shah, and Johan Ugander for helpful discussions. We thank Abdul Latif Jameel Poverty Action Lab (J-PAL SA), Tithee Mukhopadhyay, Shreya Chaturvedi, Vasu Chaudhary, Shobitha Cherian, Arnesh Chowdhury, Anoop Singh Rawat, and Meghna Yadav for research assistance. The computations in this paper were run on the Faculty of Arts and Sciences Research Computing (FASRC) Cannon cluster supported by the Faculty of Arts and Sciences (FAS) Division of Science Research Computing Group at Harvard University. We gratefully acknowledge financial support from the NSF under Grants SES-1629446, SES-2018554, and RAPID 2029880. A.G.C. thanks the Alfred P. Sloan foundation for support.

1. A. Hardy, Cholera, quarantine and the English preventive system, 1850–1895. *Med. Hist.* **37**, 250–269 (1993).
2. G. F. Gensini, M. H. Yacoub, A. A. Conti, The concept of quarantine in history: From plague to SARS. *J. Infect.* **49**, 257–261 (2004).

3. E. Tognotti, Lessons from the history of quarantine, from plague to influenza. *A. Emerg. Infect. Dis.* **19**, 254–259 (2013).
4. J. M. Drazen *et al.*, Ebola and quarantine. *N. Engl. J. Med.* **371**, 2029–2030 (2014).

5. M. O. Jackson, D. Lopez-Pintado, Diffusion and contagion in networks with heterogeneous agents and homophily. *Netw. Sci.* **1**, 49–67 (2013).
6. D. Holtz *et al.*, Interdependence and the cost of uncoordinated responses to COVID-19. *Proc. Natl. Acad. Sci. U.S.A.* **117**, 19837–19843 (2020).
7. C. Cheng, J. Li, C. Zhang, Z. Chuanchuan, Variations in governmental responses to and the diffusion of COVID-19: The role of political decentralization. *SSRN*, 3665067 (2020).
8. C. Bridges, M. Kuehnert, C. Hall, Transmission of influenza: Implications for control in health care settings. *Clin. Infect. Dis.* **37**, 1094–1101 (2003).
9. Q. A. T. Bosch *et al.*, Contributions from the silent majority dominate dengue virus transmission. *PLoS Pathog.* **14**, e1006965 (2018).
10. Y. Bai *et al.*, Presumed asymptomatic carrier transmission of COVID-19. *J. Am. Med. Assoc.* **323**, 1406–1407 (2020).
11. Z. Wu, J. M. McGoogan, Characteristics of and important lessons from the coronavirus disease 2019 (COVID-19) outbreak in China: Summary of a report of 72 314 cases from the Chinese Center for Disease Control and Prevention. *J. Am. Med. Assoc.* **323**, 1239–1242 (2020).
12. D. Lazer *et al.*, “Failing the test: Waiting times for COVID diagnostic tests across the U.S.” in *The State of the Nation: A 50-State COVID-19 Survey Report*, D. Lazer, Ed. (The COVID-19 Consortium for Understanding the Public’s Policy Preferences Across States, 2020).
13. W. E. Parment, M. S. Sinha, COVID-19—the law and limits of quarantine. *N. Engl. J. Med.* **382**, e28 (2020).
14. A. Wilder-Smith, C. J. Chiew, V. J. Lee, Can we contain the COVID-19 outbreak with the same measures as for SARS? *Lancet Infect. Dis.* **20**, e102–e107 (2020).
15. A. C. Ghani, J. Swinton, G. P. Garnett, The role of sexual partnership networks in the epidemiology of gonorrhoea. *Sex. Transm. Dis.* **24**, 45–56 (1997).
16. A. Jolly, J. Wylie, Gonorrhoea and chlamydia core groups and sexual networks in Manitoba. *Sex. Transm. Infect.* **78**, i145–i151 (2002).
17. M. E. Halloran, I. M. Longini, A. Nizam, Y. Yang, Containing bioterrorist smallpox. *Science* **298**, 1428–1432 (2002).
18. M. J. Keeling, K. T. Eames, Networks and epidemic models. *J. R. Soc. Interface* **2**, 295–307 (2005).
19. S. A. Lauer *et al.*, The incubation period of coronavirus disease 2019 (COVID-19) from publicly reported confirmed cases: Estimation and application. *Ann. Intern. Med.* **172**, 577–582 (2020).
20. A. Banerjee *et al.*, Messages on COVID-19 prevention in India increased symptoms reporting and adherence to preventive behaviors among 25 million recipients with similar effects on non-recipient members of their communities. *SSRN*, 3649860 (2020).
21. W. O. Kermack, A. G. McKendrick, A contribution to the mathematical theory of epidemics. *Proc. R. Soc. Lond. Ser. A Contain. Pap. Math. Phys. Character* **115**, 700–721 (1927).
22. N. T. Bailey, *The Mathematical Theory of Epidemics* (Charles Griffin & Company, London, 1957).
23. R. M. Anderson, R. M. May, *Infectious Diseases of Humans: Dynamics and Control* (Oxford University Press, 1992).
24. D. McAdams, Economic epidemiology of infection. *Covid Economics*, **48** (2020).
25. M. E. Newman, D. J. Watts, Scaling and percolation in the small-world network model. *Phys. Rev. E* **60**, 7332 (1999).
26. M. E. Newman, Spread of epidemic disease on networks. *Phys. Rev. E* **66**, 016128 (2002).
27. S. Flaxman *et al.*, Estimating the effects of non-pharmaceutical interventions on COVID-19 in Europe. *Nature* **584**, 257–261 (2020).
28. D. J. Watts, S. H. Strogatz, Collective dynamics of ‘small-world’ networks. *Nature* **393**, 440–442 (1998).
29. L. A. N. Amaral, A. Scala, M. Barthelemy, H. E. Stanley, Classes of small-world networks. *Proc. Natl. Acad. Sci. U.S.A.* **97**, 11149–11152 (2000).
30. F. Chung, L. Lu, The average distances in random graphs with given expected degrees. *Proc. Natl. Acad. Sci. U.S.A.* **99**, 15879–15882 (2002).
31. D. J. Watts, *Small Worlds: the Dynamics of Networks between Order and Randomness* (Princeton University Press, 2004).
32. T. H. McCormick, M. J. Salganik, T. Zheng, How many people do you know?: Efficiently estimating personal network size. *J. Am. Stat. Assoc.* **105**, 59–70 (2010).
33. A. V. Banerjee *et al.*, Changes in social network structure in response to exposure to formal credit markets. *SSRN*, 3245656 (2018).
34. L. A. Beaman, A. BenYishay, J. Magruder, A. M. Mobarak, Can network theory-based targeting increase technology adoption? *SSRN*, 3225815 (2018).
35. X. Hao *et al.*, Reconstruction of the full transmission dynamics of COVID-19 in Wuhan. *Nature* **584**, 420–424 (2020).
36. A. Hortaçsu, J. Liu, T. Schwieg, Estimating the fraction of unreported infections in epidemics with a known epicenter: An application to COVID-19. *J. Econom.* **220**, 106–129 (2021).
37. R. Li *et al.*, Substantial undocumented infection facilitates the rapid dissemination of novel coronavirus (SARS-CoV-2). *Science* **368**, 489–493 (2020).
38. M. Elliott, B. Golub, A network approach to public goods. *J. Polit. Econ.* **127**, 730–776 (2019).