Interacting Regional Policies in Containing a Disease

Arun G. Chandrasekhar*, Paul Goldsmith-Pinkham†,
Matthew O. Jackson‡,[1], and Samuel Thau§

*Department of Economics, Stanford University; J-PAL; NBER.

†Yale School of Management.

‡Department of Economics, Stanford University; Santa Fe Institute.

§Harvard University

[1]To whom correspondence should be addressed. All authors contributed equally to this manuscript.

September 15, 2020

Abstract: Regional quarantine policies, in which a portion of a population surrounding infections are 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 upon whether (i) the network of interactions satisfies a balanced-growth condition, (ii) infections have a short delay in detection, and (iii) 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 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.

Introduction

Global problems, from climate change to disease control, are hard to address without policy coordination across borders. In particular, pandemics, like COVID-19, are challenging to contain because 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 years, regional quarantines have been used to combat cholera, diphtheria, typhoid, flus, polio, ebola, and COVID-19 (*1, 2, 3, 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, 6, 7*). Second, some governments only pay attention to what goes on within their borders, which leads them to under-forecast their own infection rates.

We examine three types of quarantine policies to understand the impact of non-coordination: (i) those controlled by one actor with control of the whole society – "single regime policies," (ii) those controlled by separate jurisdictions that only react to internal infection rates – "myopic jurisdictional policies," and (iii) those controlled by separate jurisdictions that are proactive and track infections outside of their jurisdiction as well as within when deciding on when to quarantine – "proactive jurisdictional policies."

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). 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 individuals are either asymptomatically contagious (e.g., HIV, COVID-19) (*8, 9, 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 non-compliance by some people – by choice or necessity (*13, 14, 15, 16, 17, 18*). Either way, tiny leakages can spread the disease.

We show that regional quarantines curb the spread of a disease if and only if: (i) there is limited delay in observing infections, (ii) there is sufficient knowledge and control of the network to prevent leakage of infection, and (iii) the network has a certain "balanced-growth" structure. The failure of any of these conditions substantially limits quarantine effectiveness.

We then examine jurisdictional policies, which are regional quarantine policies conducted by multiple, uncoordinated regimes. The regions that need to be quarantined cross borders, leading to leakage that limits their effectiveness. As we show, myopic policies do much worse than proactive ones, as they do not forecast the impact of neighboring infection rates on

their own population. Moreover, a few lax jurisdictions, which wait for higher infection rates before quarantining, substantially worsen outcomes for all jurisdictions.

A Model

Consider a large network of nodes (individuals). Our theory is asymptotic, applying as the population grows (*19*). 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, though 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 healthcare access (*8, 10, 20, 11, 12, 21*). After that delay, the 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 (*22, 23, 24, 18, 25*), agent-based models (*26, 27, 28*), and others.

Results

Baseline: an ideal setting

We begin by analyzing a single jurisdiction with complete control.

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, the estimation of the infection origin is an additional challenge.

Whether a regional policy halts infection in this setting is fully characterized by what we call *growth-balance* (formally defined in (*19*)). This requires that the network have large enough expansion properties and that the expansion rate not drop too low in any part of the network.

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 one in ten cases are detected in a timely manner ($\alpha = 0.1$). First, consider a part of the network in which each infected person infects 3.5 others on average. If we monitor nodes within distance $k = 3$ of an infected node, a "typical" chain of infection would lead to roughly $3.5 + 3.5^2 + 3.5^3 = 58.625$ expected cases. The chance that this goes undetected is tiny: $0.9^{58.625} = 0.002$. Next, 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 chain of depth 3 leads to $1 + 1 + 1 = 3$ infections. The chance that this spread remains undetected is very high: $0.9^3 = 0.72$.

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 it is too low, then it is highly likely that any early infection will be detected before it goes beyond a distance of three away from the first infected node. If instead, the distribution of reproduction numbers gives a nontrivial chance that the disease starts out on a chain with lower reproduction numbers, like the 1, 1, 1, chain, then there is a high chance that it can travel several steps before being detected. This highlights the fact that the reproduction number R_0 alone is a crude concept, whose local distributions 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 (*29,30,31*), this allows it to be almost anywhere. Figure S1 in the supplementary materials 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 the supplementary material (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 network satisfies growth-balance (*19*).

Growth-balance is satisfied by many, but not all, sequences of random graph models, provided that the average degree d satisfies $d^k \to \infty$ (Corollary 1, (*19*)). Without growth balance, a regional policy fails non-trivially even under idealized conditions. Indeed, if contact networks are sparse and have some very low degree nodes (as they tend to empirically), then, unless quarantine regions are large ($k \to \infty$), growth-balance will fails and a regional quarantine will be ineffective.

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

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 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 network is growth-balanced and the delay in detection is shorter than the diameter of the network (Theorem 2, (*19*)). Given that real-world networks have short average distances between nodes (*32*), non-trivial delays in detection allow the disease to escape a regional quarantine.

Leakage

Next, we consider how leakage – inability to limit interactions (*13*) or mistakes in identifying portions of a network 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 (Theorem 3, (*19*)) that even a small amount of leakage leads to a nontrivial probability that a regional policy will fail. 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, with more nodes involved there is also more leakage that increase the chance that the infection escapes the quarantine.

Jurisdictions and Leakage

We can use the results from regional quarantines as a starting point to understand jurisdictional policies. For instance, leakage generally applies when interactions cross jurisdictions. Figure 1 pictures two jurisdictions that fail to nicely tessellate the network. Given leakage across jurisdictional borders, unless policies are fully coordinated across jurisdictions, our theoretical results indicate that they will fail to contain infections.

Figure 1: Inconsistency of Jurisdictions and Distances

(a) Jurisdictions with interactions that do not align **(b)** Figure 1a, but based on distance from infection

Figure 1: Nodes in two jurisdictions do not align with the distances from the initial infection. In Panel (a), the nodes are presented in a geographic sense, within their jurisdictions, and the interaction network does not comply with the jurisdictional boundaries. In Panel 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 Panel (b) could contain the infection; however, if it is only executed by the infected node's jurisdiction in panel (a) then it would fail for cross-jurisdictional connections.

Simulations

The theory provides insights into the various hurdles that quarantine policies face, but does not provide insight into how well different types policies will fare in slowing infection and at what costs.

To explore this, we simulate a contagion on a network of 140000 nodes that mimics real-world data (*33, 34, 35, 21*). These simulations illustrate our theoretical results and also show the improvements that proactive policies provide relative to myopic ones. The results are robust to choices of parameters (*19*).

The network is divided into 40 *locations*, each with a population of 3500. We generate the network using a geographic stochastic block model (*19*). 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 location (calibrated to data from India and the United States, including data collected during COVID-19 (*33, 34, 35, 21, 19*)). We set the basic reproduction rate $R_0 = 3.5$ to mimic COVID-19 (*36*), and $\theta = 5$ and $\alpha = 0.1$ (*20, 37, 38, 19*).

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

Before introducing jurisdictions, we first illustrate the effects of leakage as well as delays in detection. In Figure 2, the entire network is governed by a single policymaker using a $(k, x) = (3, 1)$-regional quarantine.

Figure 2a shows the outcomes for no delay in detection nor any leakage. Consistent with

Theorem 1 the policy is effective: on average 277 nodes per million are infected (0.028% of the population), with 803956 node-days of quarantine per million nodes. Figure 2b introduces a delay in detection. With a delay of $\tau = 3$, infections increase, with 2256 nodes per million eventually infected (0.23% of the population) and 2301414 node-days of quarantine per million nodes. 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. Figure 2c adds leakage to the setup of Figure 2b, making only 95% of the intended nodes quarantined. The number of cumulative infections per million nodes increases to 5138 (0.50% of the population). The leakage increases the number of quarantined node-days to 6478055 per million nodes.

Figure 2: The Impact of Detection Delay and Leakage

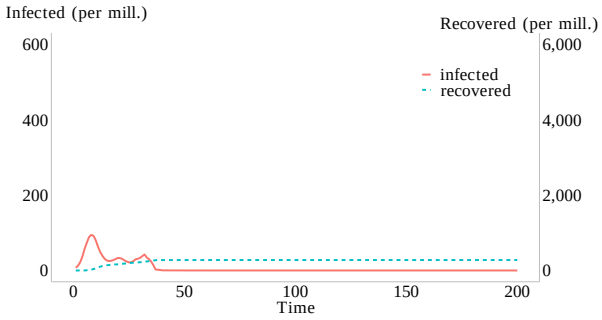

(a) $(k, x) = (3, 1)$-quarantine with no delay in detection and no leakage

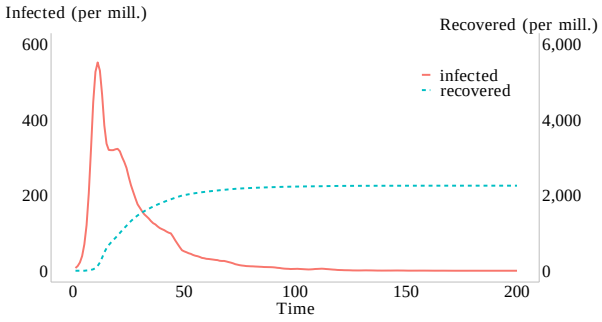

(b) $(k, x) = (3, 1)$-quarantine policy with a detection delay of 3 periods

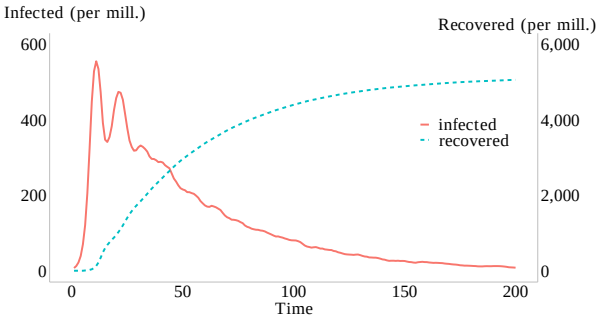

(c) $(k, x) = (3, 1)$-quarantine policy with a detection delay of 3 periods and leakage

Figure 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 Panel 2a, there is no detection delay and no leakage. In Panel 2b, we introduce a detection delay of $\tau = 3$. This represents the 3 day pre-symptomatic window during which an infected node can transmit, as well as an expected delay in seeking healthcare and testing upon symptom onset (*20, 19*). Panel 2c adds leakage of $\epsilon = 0.05$ to the setup of Panel 2b. For each figure, we simulate 10000 times on the same network with random initial infections, and present the average number of infections and recovered people over time, scaled per million.

Jurisdictional Policies

We now introduce jurisdictions to the same network as before, and each location becomes its own jurisdiction.

We compare two types of jurisdictional policies. In myopic policies each jurisdiction quarantines based entirely on internal infections. In proactive policies, jurisdictions track infections in other jurisdictions and predict their own – possibly undetected – infections and base their quarantines off of predicted infections (calculation details in (*19*)). In both cases, if a jurisdiction enters quarantine, all links within and to the jurisdiction are removed.

Figure 3 illustrates the improvement a proactive policy offers relative to myopic internal jurisdictional policies. In Figure 3a, jurisdictions use myopic policies, while in Figure 3b jurisdictions use proactive policies. In the myopic case, there are 118447 infections per million nodes (11.85% of the population), with 65634600 person-day quarantines per million nodes. Multiple waves are common: 67.4% of jurisdictions have multiple quarantines. Proactive quarantining dramatically improves outcomes (Figure 3b): only 6300 nodes per million are infected (0.630% of the population), with 37816130 person-day quarantines per million nodes. Multiple shutdowns are less frequent: 56.9% of jurisdictions quarantine more than once.

Lax Jurisdictions

Finally, we also add a few "lax" jurisdictions to the setting. These are jurisdictions that are myopic and have a high threshold of internal infections before quarantining. We examine how these few lax jurisdictions worsen the outcomes for all jurisdictions.

In Figures 3c and 3d, four lax jurisdictions react only to infections within their own borders and wait until they have detected five infections before quarantining (Simulation Details, (*19*)). Figure 3c shows the outcomes when the remaining 36 jurisdictions using myopic internal strategies, while in Figure 3d the remaining 36 jurisdictions using proactive strategies. Comparing Figures 3a to 3c, infections are much worse under the myopic internal policies. 209389 nodes per million are infected (20.9% of the population), and 72.8% of regions shut down multiple times. Of the infections in Figure 3c, 84.2% happen in low threshold-jurisdictions. Comparing Figures 3b to 3d shows that things deteriorate less for the proactive jurisdictional policies. The 27312 total infections per million nodes (2.73% of the population) is well below either set of myopic policies: 67.4% of jurisdictions have multiple quarantines; and 73.4% of the infections in Figure 3d happen in the proactive regions.

Figure 4a displays the dynamics of quarantines for each of the policy configurations from Figure 3, and Figure 4b displays the number of person-day infections versus the number of person-day quarantines.

Figure 3: The Effectiveness of Myopic vs Proactive Quarantines

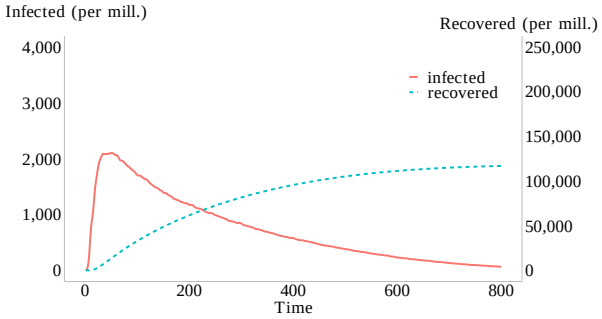

(a) Each jurisdiction myopically quarantines once it observes any infections internally, ignores other jurisdictions

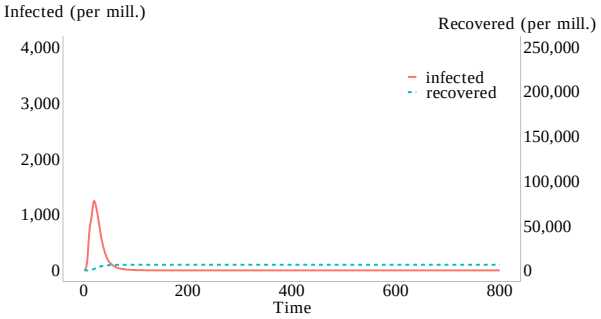

(b) Each jurisdiction proactively quarantines by estimating internal infections based on observation of other jurisdictions

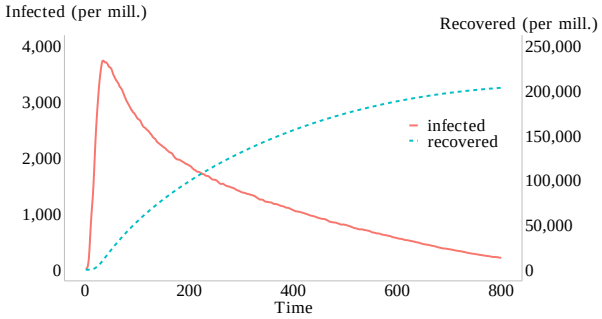

(c) 36 jurisdictions myopically quarantine once observing any internal infections, 4 lax jurisdictions only quarantine once they reach 5 internal infections

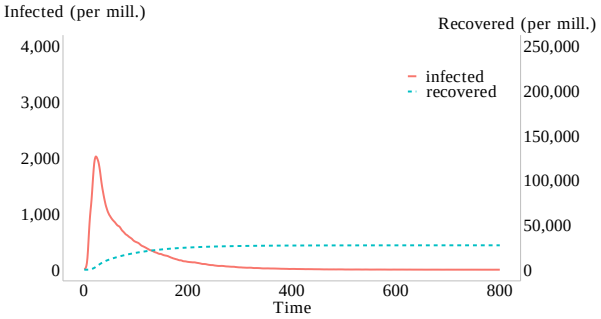

(d) 36 jurisdictions proactively quarantine by estimating internal infections based on observation of other jurisdictions, 4 lax jurisdictions only quarantine once they reach 5 internal infections

Figure 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 Panel 3a, all jurisdictions use a myopic internal policy. In Panel 3b, all jurisdictions use a proactive policy. In Panel 3c, we implement the same policies as Panel 3a, but have four lax jurisdictions that use $x = 5$ (0.14% of the jurisdiction population) instead of $x = 1$ (*19*). Panel 3d has 36 jurisdictions with proactive policies and four with lax policies. For each figure, we simulate 10000 times on the same network with random initial infections, and present the average number of infections and recovered people over time, scaled per million.

Global quarantines (closing the entire network at once) and single-jurisdiction regional quarantines (with leakage) do the best on both dimensions. Once jurisdictions are introduced, proactive jurisdictions quarantine earlier and have fewer recurrences than myopic jurisdictions. Lax jurisdictions cause an overall higher number of quarantines, over a longer time. The proactive jurisdictional policy trades off more quarantine days for substantially fewer infection days compared to the myopic internal policy, but proactive policies do significantly better than myopic policies on both dimensions when mixed with jurisdictions using lax policies.

Figure 4: The Impact and Costs of Quarantine Policies with and without Lax Jurisdictions

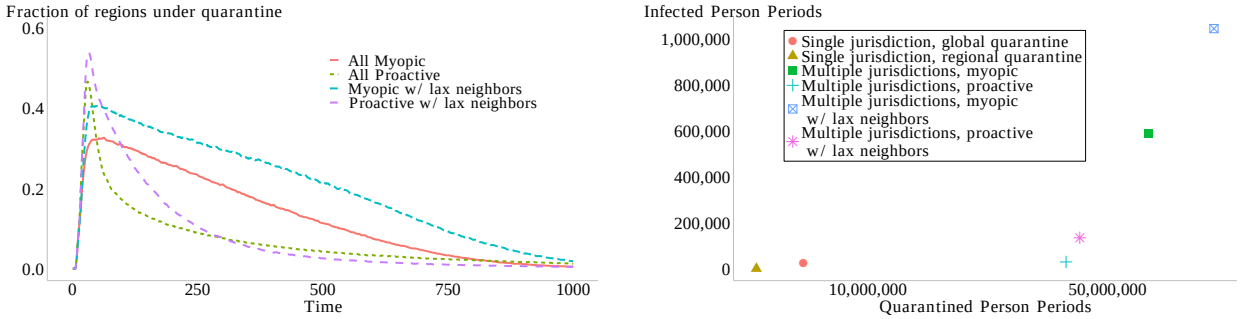

(a) Dynamics of quarantines in each of the policy configurations

(b) Person-day infections vs. person-day quarantines (per million)

Figure 4: Figure 4a displays the dynamics of quarantines for each of the policy configurations. Figure 4b 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 myopic strategy (which does worse than the global because of leakage). With 40 jurisdictions, both proactive policies outperform the internal, myopic policies. By far the worst, on both dimensions, is the internal, myopic policy with some lax jurisdictions. These results come from the same solutions that produce figures 2 and 3.

Discussion

We have shown that regional quarantine policies are likely to fail unless leakage and delays in detection are limited. Multiple jurisdictions using independent policies are even less effective, as leakage occurs across jurisdictional borders. We have also shown that there are substantial improvements from proactive policies, and that a few lax jurisdictions greatly worsen the outcomes for all jurisdictions.

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 (*39*). 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), all help attenuate contagion, but unless they maintain the reproduction number below one, the problems identified here remain. Even tiny fractions of interactions across boundaries are enough to lead to spreading in large populations. With modern inter- and intranational trade being a sizable portion of all economies, such interaction is difficult to avoid. Nonetheless, our analysis 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 come close to satisfying the conditions of our first theorem.

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. Our growth-balance condition also points out that not all parts of a network are equal in their potential for undetected transmission. In places where the reproduction number is lower, so is the probability of observing outbreaks, enabling undetected leakage of infections.

References and Notes

1. Hardy, A. Cholera, quarantine and the English preventive system, 1850–1895. *Medical History* **37**, 250–269 (1993).

2. Gensini, G. F., Yacoub, M. H. & Conti, A. A. The concept of quarantine in history: from plague to SARS. *Journal of Infection* **49**, 257–261 (2004).

3. Tognotti, E. Lessons from the history of quarantine, from plague to influenza A. *Emerging Infectious Diseases* **19**, 254 (2013).

4. Drazen, J. M. *et al.* Ebola and quarantine (2014).

5. Jackson, M. O. & Lopez-Pintado, D. Diffusion and contagion in networks with heterogeneous agents and homophily. *Network Science* **1:1**, 49–67 (2013).

6. Holtz, D. *et al.* Interdependence and the cost of uncoordinated responses to covid-19. *Proceedings of the National Academy of Sciences* **117**, 19837–19843 (2020). URL `https://www.pnas.org/content/117/33/19837`. `https://www.pnas.org/content/117/33/19837.full.pdf`.

7. Cheng, C., Li, J. & Zhang, C. Variations in governmental responses to and the diffusion of covid-19: The role of political decentralization. *Available at SSRN 3665067* (2020).

8. Bridges, C., Kuehnert, M. & Hall, C. Transmission of influenza: implications for control in health care settings. *Clinical Infectious Diseases* **37**, 1094–1101 (2003).

9. Ten Bosch, Q. A. *et al.* Contributions from the silent majority dominate dengue virus transmission. *PLoS pathogens* **14**, e1006965 (2018).

10. Bai, Y. *et al.* Presumed asymptomatic carrier transmission of COVID-19. *Journal of the American Medical Association* **323**, 1406–1407 (2020).

11. Wu, Z. & McGoogan, J. M. 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. *Journal of the American Medical Association* **323**, 1239–1242 (2020).

12. Lazer, D. *et al.* Failing the test: Waiting times for COVID diagnostic tests across the U.S. *The State of the Nation: A 50-State COVID-19 Survey Report* (2020).

13. Parmet, W. E. & Sinha, M. S. COVID-19—the law and limits of quarantine. *New England Journal of Medicine* **382**, e28 (2020).

14. Wilder-Smith, A., Chiew, C. J. & Lee, V. J. Can we contain the COVID-19 outbreak with the same measures as for SARS? *The Lancet Infectious Diseases* (2020).

15. Ghani, A. C., Swinton, J. & Garnett, G. P. The role of sexual partnership networks in the epidemiology of gonorrhea. *Sexually Transmitted Diseases* **24**, 45–56 (1997).

16. Jolly, A. & Wylie, J. Gonorrhoea and chlamydia core groups and sexual networks in Manitoba. *Sexually Transmitted Infections* **78**, i145 (2002).

17. Halloran, M. E., Longini, I. M., Nizam, A. & Yang, Y. Containing bioterrorist smallpox. *Science* **298**, 1428–1432 (2002).

18. Keeling, M. J. & Eames, K. T. Networks and epidemic models. *Journal of the Royal Society Interface* **2**, 295–307 (2005).

19. Materials and methods are available as supplementary materials at the science website.

20. Lauer, S. A. *et al.* The incubation period of coronavirus disease 2019 (COVID-19) from publicly reported confirmed cases: estimation and application. *Annals of Internal Medicine* **172**, 577–582 (2020).

21. Banerjee, A. *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. Tech. Rep., National Bureau of Economic Research (2020).

22. Kermack, W. O. & McKendrick, A. G. A contribution to the mathematical theory of epidemics. *Proceedings of the Royal Society of London. Series A, Containing papers of a mathematical and physical character* **115**, 700–721 (1927).

23. Bailey, N. T. The mathematical theory of epidemics. Tech. Rep. (1957).

24. Anderson, R. M. & May, R. M. *Infectious diseases of humans: dynamics and control* (Oxford university press, 1992).

25. McAdams, D. Economic epidemiology of infection. *Annual Review of Economics* (2020).

26. Newman, M. E. & Watts, D. J. Scaling and percolation in the small-world network model. *Physical Review E* **60**, 7332 (1999).

27. Newman, M. E. Spread of epidemic disease on networks. *Physical Review E* **66**, 016128 (2002).

28. Flaxman, S. *et al.* Estimating the effects of non-pharmaceutical interventions on COVID-19 in Europe. *Nature* 1–5 (2020).

29. Watts, D. J. & Strogatz, S. H. Collective dynamics of 'small-world' networks. *Nature* **393**, 440–442 (1998).

30. Amaral, L. A. N., Scala, A., Barthelemy, M. & Stanley, H. E. Classes of small-world networks. *Proceedings of the National Academy of Sciences* **97**, 11149–11152 (2000).

31. Chung, F. & Lu, L. The average distances in random graphs with given expected degrees. *Proceedings of the National Academy of Sciences* **99**, 15879–15882 (2002).

32. Watts, D. J. *Small worlds: the dynamics of networks between order and randomness* (Princeton university press, 2004).

33. McCormick, T. H., Salganik, M. J. & Zheng, T. How many people do you know?: Efficiently estimating personal network size. *Journal of the American Statistical Association* **105**, 59–70 (2010).

34. Banerjee, A. V., Chandrasekhar, A. G., Duflo, E. & Jackson, M. O. Changes in social network structure in response to exposure to formal credit markets. *Available at SSRN 3245656* (2018).

35. Beaman, L., BenYishay, A., Magruder, J. & Mobarak, A. M. Can network theory-based targeting increase technology adoption? Tech. Rep., National Bureau of Economic Research (2018).

36. Hao, X. *et al.* Reconstruction of the full transmission dynamics of COVID-19 in Wuhan. *Nature* 1–7 (2020).

37. Hortaçsu, A., Liu, J. & Schwieg, T. Estimating the fraction of unreported infections in epidemics with a known epicenter: an application to COVID-19. Tech. Rep., National Bureau of Economic Research (2020).

38. Li, R. *et al.* Substantial undocumented infection facilitates the rapid dissemination of novel coronavirus (SARS-CoV-2). *Science* **368**, 489–493 (2020).

39. Elliott, M. & Golub, B. A network approach to public goods. *Journal of Political Economy* **127**, 730–776 (2019).

40. Ugander, J., Karrer, B., Backstrom, L. & Kleinberg, J. Graph cluster randomization: Network exposure to multiple universes. In *Proceedings of the 19th ACM SIGKDD international conference on Knowledge discovery and data mining*, 329–337 (2013).

41. Hoff, P. D., Raftery, A. E. & Handcock, M. S. Latent space approaches to social network analysis. *Journal of the American Statistical Association* **97:460**, 1090–1098 (2002).

42. Penrose, M. *Random geometric graphs*, vol. 5 (Oxford university press, 2003).

43. Leskovec, J., Lang, K. J., Dasgupta, A. & Mahoney, M. W. Statistical properties of community structure in large social and information networks. In *Proceedings of the 17th International Conference on World Wide Web*, 695–704 (2008).

44. Banerjee, A. V., Chandrasekhar, A. G., Duflo, E. & Jackson, M. O. Diffusion of microfinance. *Science* **341**, DOI: 10.1126/science.1236498, July 26 2013 (2013).

45. Chandrasekhar, A. G. & Jackson, M. O. A network formation model based on subgraphs. *SSRN paper no 2660381.* (2016).

46. Breza, E., Chandrasekhar, A. G., McCormick, T. H. & Pan, M. Using aggregated relational data to feasibly identify network structure without network data. *American Economic Review* **Forthcoming** (2019).

47. Erdös, P. & Rényi, A. On random graphs. *Publ. Math. Debrecen* **6**, 156 (1959).

48. Furukawa, N. W., Brooks, J. T. & Soble, J. Evidence supporting transmission of severe acute respiratory syndrome coronavirus 2 while presymptomatic or asymptomatic. *Emerging Infectious Diseases* **26** (2020).

49. Cuomo, A. Continuing temporary suspension and modification of laws relating to the disaster emergency. `https://www.governor.ny.gov/news/no-2028-continuing-temporary-suspension-and-modification-laws-relating-disaster-emergency` (2020).

50. The COVID Tracking Project. Totals by state (2020). URL `https://covidtracking.com/data/`.

51. DeSantis, R. Phase 1: Safe. Smart. Step-by-step. Plan for Florida's recovery. `https://www.flgov.com/wp-content/uploads/2020/04/EO-20-112.pdf` (2020).

52. Emergency Information from Swedish Authorities. Restrictions and prohibitions. `https://www.krisinformation.se/en/hazards-and-risks/disasters-and-incidents/2020/official-information-on-the-new-coronavirus/restriktioner-och-forbud` (2020).

Acknowledgments

We thank Marcella Alsan, Abhijit Banerjee, Gabriel Carroll, Bharat Chandar, Dean Eckles, Ben Golub, Dave Holtz, Ali Jadbabaie, Ed Kaplan, Jon Kleinberg, Devavrat Shah, and Johan Ugander for helpful discussions. We thank 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 FASRC Cannon cluster supported by the FAS Division of Science Research Computing Group at Harvard University.

Funding

We gratefully acknowledge financial support from the NSF under grants SES-1629446, SES-2018554, and RAPID # 2029880.

Competing Interests

No authors have competing interests to report.

Data and Materials Availability

Code and data can be made available at the editor or referee's request.

List of Supplementary Materials

1. Materials and Methods: Modeling an epidemic and quarantine policy
2. Materials and Methods: Simulation details
3. Tables S1-S5
4. Fig S1

Supplementary Material

Interacting Regional Policies in Containing a Disease
by Chandrasekhar, Goldsmith-Pinkham, Jackson, Thau

A Modeling An Epidemic and Quarantine Policy

The Model

People and Interactions

There are $n > 1$ nodes (individuals) in an unweighted, and possibly directed, network.

We study the course of a disease through the network. Time is discrete, with periods indexed by $t \in \mathbb{N}$. An initial infected node, indexed by $i_0 \in V$, is the only node infected at time 0. We call this node the *seed*.

We track the network via neighborhoods that expand outwards via (directed) paths from i_0. Let N_k be all the nodes who are at (directed) distance k from node i_0. We use n_k to denote the cardinality of N_k.

For any node in $j \in N_{k'}$, for $k' < k$, let n^j be the number of its direct descendants and n_k^j be the number of its (possibly indirect) descendants in N_k that are reached by never passing beyond distance k from i_0.

All unweighted network models are admitted here. Additionally, all results extend directly to any weighted model in which weights are bounded above and below (e.g., probabilities of interaction). Note also, that the network can be directed or undirected.

The infection process proceeds as follows. In every time period $t \in \{1, 2, \ldots\}$, an infected node i transmits the disease to each of i's neighbors independently with probability p. A newly infected node is infectious for $\theta \geq 1$ periods after which the node recovers and is never again infectious. The model can easily be extended to accommodate renewed susceptibility.

There may be a *delay* in the ability to detect the disease. The number of periods of delay is given by τ with $0 \leq \tau \leq \theta$. Delay is a general term that can capture many things. For example, it can correspond to (a) asymptomatic infectiousness, (b) a delay in accessing health care given the onset of an infectious period, (c) any delay in the administration of testing, and so on.

In the first period of an infected node's infectious period – after delay – there is a probability α that the policymaker detects it as being infected. So, potential detection happens exactly once during the first period in which the node can be detected. Detection is independently and identically distributed. Our results are easily extended to have a random period for detection after the delay.

Finally, the policymaker may face some error in their knowledge of the network. This can come from their inability to enforce exactly the interactions they wish to allow or limit, this can come from random variation in data collected to estimate interaction networks, or

this can come from misspecification. If there is error, we will track a share ϵ of nodes that are within a k-neighborhood of the seed but are estimated by the policymaker to be outside the k-neighborhood.

Regional Quarantine Policy

Let *regional policy of distance k and threshold x* be such that once there are at least x infections (other than the seed) detected within distance k from the initial seed, then all nodes within distance $k+1$ of i_0 are quarantined for at least θ periods. A quarantine implies all connections between nodes are severed to avoid any further transmission and the infection waits out its duration θ and dies out.

Implicit in this definition is that a quarantine is not instantaneous, but that infected people could have infected their friends before being shut down, which is why the nodes at distance $k+1$ are quarantined. All the results below extend if we assume that it is instantaneous, but with quarantines moved back one step and path lengths in definitions correspondingly adjusted.

We have assumed the policymaker knows the "seed," for simplicity - and which may take some time in reality. This provides an advantage to the policymaker, but we see substantial. containment failures despite this advantage.

Growth Balance

In order to conduct asymptotic analysis, a useful device to study the probabilities of events in question in large networks, we study a sequence of networks $G(n)$ with $n \to \infty$ and an associated sequence of parameters $(\alpha, p, \tau, \theta, k) = (\alpha(n), p(n), \tau(n), \theta(n), k(n))$. In what follows when we drop the index n, and it is implied unless otherwise stated.

Consider a network and a distance k from the initially infected node i_0. A *path of potential infection to $k+2$* is a sequence of nodes $i_0, i_1, \ldots i_\ell$ with $i_\ell \in N_{k+1}$, i_{j+1} being a direct descendant of i_j for each $j \in \{0, \ldots, \ell-1\}$, and for which i_ℓ has a descendant in N_{k+2}.

Consider a sequence of networks and $k(n)$s. We say that there are *bounded paths of potential infection to $k(n)+2$* if there exists some finite M and for each n there is a path of potential infection to $k(n)+2$, $i_0, i_1, \ldots i_\ell$ of length less than M, with $n^j < M$ for every $j \in \{0, \ldots, \ell-2\}$.

We say that a sequence of networks is *growth-balanced* relative to some $k(n)$ if there are no bounded paths of potential infection to $k(n)+2$.

Growth balance is essentially a condition that requires a minimum bound of expansion along all paths from some initial infection. The intuition behind the condition is clear: in order to be sure to detect an infection, within distance k of the seed, it has to be that many of the nodes within distance k have been exposed to the disease by the time it reaches distance k. What is ruled out is a relatively short path that gets directly to that distance without

having many nodes be exposed along that path.[1]

Figure S1 presents an illustration of a network that is not growth-balanced.

Results

A Benchmark: No Delay in Detection; Perfect Information and Enforcement

We begin with a benchmark case in which there is no delay in detection ($\tau = 0$) and the policymaker can completely enforce a quarantine at some distance $k + 1$.[2]

We allow the size of the quarantine region k to depend on n in any way, as the theorem still applies. We work with an arbitrary but fixed threshold x, in order to allow infections to be detected. What is important is that x not grow too rapidly, as otherwise there is no chance of observing that many infections within some distance of the seed.[3]

THEOREM **1.** *Consider any sequence of networks and associated $k(n) < K(n) - 1$ where $K(n)$ is the maximum k for which $n_k > 0$;*[4] *such that each node in $N_{k(n)+1}$ has at least one descendent at distance $k(n) + 2$, and let x be any fixed positive integer. Let the sequence of associated diseases have $\alpha(n)$ and $p(n)$ bounded away from 0 and 1,*[5] *no delay in detection, and any $\theta(n) \geq 1$. A regional quarantining policy of distance $k(n)$ and threshold x halts all infections past distance $k(n)+1$ with a probability tending to 1* if and only if *the sequence is growth-balanced with respect to $k(n)$.*

Note that the growth-balance condition implies that the number of nodes within distance $k(n)$ from i_0 must be growing without bound. Theorem 1 thus implies that in order for a regional policy to work, the region must be growing without bound, and also must satisfy a particular balance condition.

Proof of Theorem 1. To prove the first part, note that if the infection never reaches distance k then the result holds directly since it can then not go beyond $k + 1$. We show that if the sequence of networks is growth-balanced relative to k, then conditional upon an infection reaching level k with the possibility of reaching $k + 2$ within two periods, the probability that it infects more than x nodes within distance k before any nodes beyond k tends to 1. Suppose that infection reaches some node at distance k that can reach a node in N_{k+1}. Consider the corresponding sequence of paths of infected nodes $i_0, i_1, \ldots i_\ell$ with $i_\ell \in N_{k+1}$, i_{j+1} being a direct descendant of i_j for each $j \in \{0, \ldots, \ell - 1\}$, and note that by assumption

[1] This is very different from conditions that concern long paths within short distances, such as (*40*), as ours is ruling out *short* paths with low expansion.

[2] Note that this requires knowledge of the neighborhood structure around the seed node, but no other knowledge of the network by a policy maker.

[3] The theorem extends to allow x to grow with n, provided the growth is sufficiently slow, and then that growth-balance condition becomes more complicated, as the M in that definition adjusts with the rate of growth of x.

[4] Otherwise, it is actually a global policy.

[5] The cases of p or α equal to 1 are degenerate.

i_ℓ has a descendant in N_{k+2}. By the growth-balance condition, for any M, there is a large enough n for which either the length of the path is longer than M or else there is at least one i_j with $j \leq \ell - 2$ along the path that has more than M descendants. In the latter case, the probability that i_j has more than x descendants who become infected and are detected is at least $1 - F_{M,m}(x)$ where $F_{M,m}$ is the binomial distribution with M draws each with probability m, where $p\alpha > m$ for some fixed m. Given that x and m are fixed, this tends to probability 1 as M grows. In the former case, the sequence exceeds length M, all of which are infected and so given that α is bounded below, the probability that at least x of them are detected goes to 1 as M grows. In both cases, as n grows, the minimal M across such paths of potential infection to $k+1$ grows without bound, and so the probability that there are at least x infections that are detected by the time that $i_{\ell-1}$ is reached tends to 1 as n grows.

To prove the converse, suppose that the network is not growth-balanced. Consider a sequence of bounded paths of potential infection to $k+2$, with associated sequences of nodes $i_0, i_1, \ldots i_\ell$ of length less than M with $i_\ell \in N_{k+1}$, i_{j+1} being a direct descendant of i_j for each $j \in \{0, \ldots, \ell - 1\}$, with $n^j < M$ for every $j \in \{0, \ldots, \ell - 2\}$, and for which i_ℓ has a descendant in N_{k+2}. The probability that each of the nodes $i_1, \ldots i_{\ell-2}$ becomes infected and no other nodes are infected within distance $k-1$, and that all infected nodes are undetected is at least $(p(1-\alpha)(1-p)^M)^M$. This is fixed and so bounded away from 0. This implies that probability that the infection gets to nodes at distance k, and $i_{\ell-1}$ in particular, without any detections is bounded below. Thus, there is a probability bounded below of reaching i_ℓ before any detections, and then by the time the quarantine is enacted, there is at least a p times this probability that it escapes past N_{k+1}, which is thus also bounded away from 0. □

We note that Theorem 1 admits essentially all sequences of (unweighted) networks. Thus, for every type of network, one can determine whether a regional policy of some k, x will succeed or fail. The only thing that one needs to check is growth-balance. If it is satisfied, a regional policy works, and otherwise it will fail with nontrivial probability.

The following corollary details the implications of the theorem for some prominent random network models.

Corollary **1.**

1. *For a sequence of block models (with Erdos-Renyi as a special case),*[6] *a regional policy with a bounded k has a probability going to 1 of halting the disease on the randomly realized network if and only if the seed node's expected out degree d is such that $d^k \to \infty$.*

2. *For a regular expander graph with outdegree d, a regional policy works if and only if the expansion rate $d^k \to \infty$.*

[6]Consider a sequence of block models such that the ratio of expected out degree of a node in one neighborhood compared to another in some other block cannot grow without bound.

3. *For a regular lattice of degree d, a regional policy works if and only if* $d^k \to \infty$.

4. *For a rewired lattice with a fraction links that are randomly rewired, a regional policy with a bounded k has a probability going to 1 of halting the disease on the randomly realized network if and only if* $d^k \to \infty$.

5. *For a sequence of random networks with a scale-free degree distribution, a regional policy works (with probability 1) if and only if* $k \to \infty$.

Thus, whether a regional policy works in almost any network model requires that either the degree of almost all nodes grows without bound, or else the size of the quarantine grows without bound. For a scale free distribution, there is always a nontrivial probability on small degrees, and hence in order for a regional policy to work, the size of the neighborhood must grow without bound.

In practice, even very sparse networks will have a large d^k (e.g., if people have hundreds of contacts, 100^3 is already a million and even with a very low α many infections will be detected within a few steps of the initial node).[7] What the growth-balance condition rules out is that some nontrivial part of the network have neighborhoods with many fewer contacts - so there cannot be people who have just a few contacts, since that will allow for a nontrivial probability of undetected escape (e.g., $2^3 = 8$ and so with only 8 infections, it is possible that none are detected and the disease escapes beyond 3 steps). As many real-world network structures have substantial heterogeneity, with some people having very low numbers of interactions, such an escape becomes possible even under idealized assumptions of no delay in detection and no leakage (*41, 42, 43, 44, 45*).

Delay in Detection

The detection delay, τ, is distributed over the support $\{1, \ldots, \tau^{\max}\}$. This includes degenerate distributions with $\tau^{\max}$ being the maximal value of the support with positive mass. The policymaker may or may not know $\tau^{\max}$ and we study both cases. The latter is important as in practice we estimate delay periods so there is bound to be uncertainty. When τ is known, we can simply say $\tau = \tau^{\max}$.

Let a *regional policy with trigger k and threshold x and buffer h* be such that once there are at least x infections detected within distance $k + h$ from the initial seed, then all nodes within distance $k + h + 1$ of i_0 are quarantined/locked down for at least θ periods.

There are two differences between this definition of regional policy from the one considered before. First, it is triggered by infections within distance $k + h$ (not within distance k), and it also has a buffer in how far the quarantine extends beyond the k-th neighborhood.

We extend the definition of growth balance to account for buffers.

[7] This is still extremely sparse, as having 100 contacts out of millions or billions of potential other nodes is a small fraction.

Consider a network and a distance k from the initially infected node i_0 and an $h \geq 1$. A *path of potential infection to* $k+h+2$ is a sequence of nodes $i_0, i_1, \ldots i_\ell$ with $i_\ell \in N_{k+h+1}$, i_{j+1} being a direct descendant of i_j for each $j \in \{0, \ldots, \ell-1\}$.

Consider a sequence of networks, n, and associated $k(n), h(n)$. We say that there are *bounded paths of potential infection* to $k(n)+h(n)+2$ if there exists some finite M and for each n there is a path of potential infection to $k+h+2$, $i_0, i_1, \ldots i_\ell$ of length less than M, with $n^j < M$ for every $j \in \{0, \ldots, \ell-h-2\}$. We say that a sequence of networks is *growth-balanced* relative to some $k(n)$ and buffers $h(n)$ if there are no bounded paths of potential infection to $k(n)+h(n)+2$.

THEOREM **2.** *Consider any sequence of networks and $k(n) < K(n) - h - 1$ where $K(n)$ is the maximum k for which $n_k > 0$, such that each node in $N_{k'}$ for $k' > k$ has at least one descendent at distance $k'+1$, and let x be any fixed positive integer. Let the sequence of associated diseases have $\alpha(n)$ and $p(n)$ bounded away from 0 and 1, $\theta(n) \geq 1$, and have a detection delay distributed over some set $\{1, \ldots, \tau^{\max}\}$ with $\tau^{\max} > 1$ (with probability on $\tau^{\max}$ bounded away from 0).*[8] *A regional policy with trigger $k(n)$, threshold x, and buffer $\tau^{\max}$ halts all infections past distance $k(n)+\tau^{\max}+1$ with a probability tending to 1* if and only if *the sequence is growth-balanced with respect to $k(n)$.*

The Proof of Theorem 2 is a straightforward extension of the previous proof and so it is omitted.

This result shows several things. First, if the detection delay is small relative to the diameter of the graph, one can use a regional quarantine policy – adjusted for the detection delay – along the lines of that from Theorem 1 and ensure no further spread. This is true even if the period is stochastic as long as the upper bound is known to be small.

Second, and in contrast, if the detection delay is large compared to the diameter of the graph, then a regional policy is insufficient. By the time infections are observed, it is too late to quarantine a subset of the graph. This condition will tend to bind in the case of real world networks, as they exhibit small world properties and have small diameters (*30, 31*). As a result, even short detection delays may correspond to rapidly moving wavefronts that spread undetected.

Leakage in the Quarantine

Next we turn to the case of in which there is some leakage in the quarantine, which may come for a variety of reasons. The policymaker may have measurement error in knowing the network structure of the network and who should be quarantined. Second, and distinctly, the policymaker may be unable to control some nodes or interactions. Third, the network may leak across jurisdictions and some nodes within distance k of i_0 may be outside of the policymaker's jurisdiction.

[8] A special case is in which $\tau^{\max}$ is known.

To keep the analysis uncluttered, we assume no detection delay, but the arguments extend directly to the delay case with the appropriate buffer.

THEOREM **3.** *Consider any sequence of networks. Let the sequence of associated diseases have α and p bounded away from 0 and 1, and be such that $\theta \geq 1$, with no detection delay. Consider any $k(n) < K - 1$ where K is the maximum k for which $n_k > 0$, suppose that each node in $N_{k(n)}$ has at least one descendent at distance $k(n)+1$, and let x be any positive integer.*

Suppose that a random share of ε_n of nodes within distance k of i_0 are not included in a regional quarantine policy and connected to nodes of distance greater than $k+1$ – because of a lack of jurisdiction, misclassification by a policymaker, or lack of complete control over people's behaviors. Then:

1. *If $\varepsilon_n = o((\sum_{k' \leq k} n_{k'})^{-1})$ and the network is growth-balanced, then a regional policy of distance k and threshold x halts all infections past distance $k+1$ with a probability tending to 1.*

2. *If $\varepsilon_n \geq \min[1/x, \eta]$ for all n for some $\eta > 0$ or the network is not growth-balanced, then a regional policy of distance $k(n)$ and threshold x fails to halt all infections past distance $k(n)+1$ with a probability bounded away from 0.*

Proof of Theorem 3. Part 1 follows from the fact that if $\varepsilon_n = o((\sum_{k' \leq k} n_{k'})^{-1})$ then the probability of having all nodes in N_k correctly identified as being in N_k tends to 1, and then Theorem 1 can be applied.

For Part 2, suppose that some x infections are detected. The probability that at least one of them is misclassified is at least $1 - (1 - \varepsilon_n)^x$. Given that $\varepsilon_n \geq \min[1/x, \eta]$ for any $\eta > 0$, it follows that $(1 - \varepsilon_n)^x$ is bounded away from 1. There is a probability bounded away from 0 that at least one of the infected nodes is misclassified, and not subject to the quarantine, and connected to a node outside of distance $k+1$. □

The theorem implies that the effectiveness of a regional policy is sensitive to any small fixed ε amount of leakage.

B Simulation Details

To illustrate the processes described in the main text, we run several simulations. First, we construct a large network with many jurisdictions. We directly study the content of the theorems with several versions of (k, x) quarantines with an SIR infection process on a network. We use the same process and network to show the issues with regional containment, studying regional and adaptive policies.

Network Model

We model real world network structure as follows.

1. There are L locations distributed uniformly at random on the unit sphere.

   Each location has a population of m nodes with a total of $n = mL$ nodes in the network.

2. The linking rates across locations are given as in a spatial model (*41, 46*). The probability of nodes $i \in \ell$ and $j \in \ell'$ for locations $\ell \neq \ell'$ linking depends only on the locations of the two nodes and declines in distance:

   $$q_{\ell,\ell'} = \exp(a + b \cdot \text{dist}(\ell, \ell'))$$

   where $\text{dist}(\ell, \ell')$ is the distance between the two locations on the sphere and $a, b < 0$.

   Every interaction between every pair of nodes is drawn independently from the observed spatial distribution, with distances being along the surface of the unit sphere.

3. The linking patterns within a location are given as in a mixture of random geometric (RGG) (*42*) and Erdos-Renyi (ER) random graphs (*47*). Specifically, as spheres are locally Euclidean, we model nodes in a location (e.g., in a city) as residing in a square in the tangent space to the location. The probability that two nodes within a location link declines in their distance in this square.

   We set d_{RGG} as the desired degree from the RGG. Nodes are uniformly distributed on the unit square $[0, 1]^2$, and links are formed between nodes within radius r_ℓ (*42*). To obtain the desired degree we set:

   $$r_\ell = \sqrt{\frac{d_{RGG}}{m_\ell \pi}}.$$

   The remaining links within location are drawn identically and independently with probability

   $$\pi = \frac{d_\ell - d_{RGG}}{m_\ell}$$

Where d_ℓ is the desired average degree for all nodes within location ℓ.

4. Next, we uniformly add links to create a small world effect, with identical and independently distributed probability $s = \frac{1}{cn}$, where c is an arbitrary constant and n is the total number of nodes in the network (*29*).

5. Finally, we designate a single location as a "hub," to emulate the idea that certain metro areas may have more connections to *all* other regions. To do so, we select a hub uniformly at random and add links independently and identically distributed with probability h from the hub location to every other location.

We first take $L = 40$ and $m = 3500$ for all locations. We set $a = -4$ and $b = -15$. Next, we set $d_\ell = 15.5$, and $d_{RGG} = 13.5$ for all locations. Next, we set $c = 2$. Finally, we set $h = 2.85 \times 10^{-6}$. This process results in a graph that emulates real world networks in the United States and India (*33, 34, 35, 21*). This includes data from India during the COVID-19 lockdowns about interactions within six feet, meaning that it is conservative (*21*).

We fix a graph to use in all versions of the simulations. The network we generate is sparse, clustered, and has small average distances, as demonstrated by information detailed in Table S1.

Finally, we recalculate the connection probability matrix to accurately reflect rates of connection across regions, which we call q.

Disease Process

We set parameters as follows: the duration of infection is $\theta = 5$, detection delay (when incorporated) is $\tau = 3$, and set thresholds x for quarantine based on the simulation type.

We set transmission probability p as

$$p = 1 - \left(1 - \frac{R_0}{\bar{d}}\right)^{\frac{1}{\theta}}$$

where $\bar{d}$ is the mean degree. We take $R_0 = 3.5$, based on estimates of COVID-19 (*36*).

Following estimates from the literature (5-15%), we set $\alpha = 0.1$ (*37, 38*). In the simulations, each node is either detected or not during the first period in which it can be detected, and no information comes after that.

As outlined in the main text, we begin by using $\theta = 5$ and $\tau = 3$ (*20, 37, 38, 48*).

Simulation Progression

Each time period in the simulation progresses in four parts, which happen sequentially. The simulations run as follows:

1. The policy maker sees the detected infections from the previous period, and calculates if a quarantine is necessary in the next period.

2. The disease progresses for a period. This includes new infections and recoveries.

3. Infected nodes that have just finished their detection delay of τ periods are independently detected with probability α.

4. New quarantines are enacted based on decisions made in part one of the process. Quarantines that have taken place for θ periods end.

A node that becomes infected in period t with a detection delay of τ and total disease length θ, is tested in period $t + \tau$, results are processed in $t + \tau + 1$, and they will be quarantined (if necessary) starting at the end of $t + \tau + 1$ (under the fourth item above). This means that they have $\tau + 1$ time periods during which they can infect other nodes. For instance, if $\tau = 0$ this allows a node that becomes infected but (that was not already under quarantine for other reasons) one opportunity to infect others. This process reflects that neither detection nor quarantining of individuals (or jurisdictions) happens instantaneously. In addition, we stipulate that the seed node, i_0 is not counted in the quarantining testing and calculations. This is meant to reflect that it may be unclear whether the disease is spreading or not. Nodes that are detected are marked as such until recovery.

Containment Policies

A random node i_0 is selected and the epidemic begins there. We study the epidemic curve, the number total node-days of infection, and the number of node-days of quarantine for a variety of containment strategies.

(k, x) Policies

We examine a number of scenarios using the (k, x) policy model outlined in Theorems 1-3.

In the case that the quarantine fails, but there are infections outside of the quarantine radius, the policy maker deals with them individually. The policy maker treats each detected case outside of the initial quarantine as a new seed, and immediately quarantines all nodes with the same radius as the initial quarantine.

Begin by using a simple objective function to find the optimal threshold for triggering the initial quarantine. We minimize a linear combination of the number of infected person periods and quarantined person periods. For all linear combinations where some weight is given to both terms, the optimal threshold is $x = 1$. The logic is as follows: if the initial quarantine is successful, the number of quarantined person periods will be fixed and also the minimum number of quarantined person periods. Therefore, the problem reduces to minimizing the number of infections, which is done by setting $x = 1$.

We study three versions of a (k, x) policy. First, we simulate $(k, x) = (3, 1)$ with no detection delay. Then, we incorporate a detection delay of $\tau = 3$, still using a policy of $(k, x) = (3, 1)$ with no buffer. Lastly, we study a $(3, 1)$ policy with no buffer and enforcement failures. In this case, a fraction $\epsilon = 0.05$ of nodes do not ever quarantine.

Global Quarantine Policy

A global quarantine policy imagines the state as an actor which quarantines every node for θ periods when more than $x = 1$ infections are detected globally. We study this in the case with a detection delay, to compare to the (k, x), regional, jurisdiction based, and proactive policies.

Myopic-Internal and Proactive Quarantine Policies

For both the myopic-internal and proactive policies, we take each location as a single jurisdiction.

Myopic Internal Quarantine Policies. Jurisdictions respond only to detections within their own borders, setting x independently of one another. In addition, states act independently: jurisdictions do not take detected infections outside of their borders into account. We set $x = 1$ for all jurisdictions, the most conservative possible threshold unless otherwise specified.

Proactive Quarantine Policies. We examine a more sophisticated approach to deciding when to quarantine. With this policy, each jurisdiction decides to quarantine based on not only defections within their borders, but within neighboring jurisdictions as well. In each period, each jurisdiction ℓ calculates their expected detected infections w_ℓ as follows:

$$w_{\ell,t} = \max\{w_{\ell,t-1} + y_{\ell,t} - r_{\ell,t}, z_{\ell,t}\}$$

We use $y_{\ell,t}$ to denote the number of expected new infections in region ℓ at time t, and use $r_{\ell,t}$ to denote the number of expected recoveries in ℓ at t. Each state calculates $y_{\ell,t}$ as:

$$y_{\ell,t} = p \sum_{\ell' \text{ s.t. } \ell' \text{ not quarantined at t-1}} m_{\ell'} q_{\ell,\ell'} w_{\ell',t-1}$$

The summation includes the term for spread from ℓ to still within ℓ. If ℓ is quarantined at time t, then $y_{\ell,t} = 0$. Expected recovery at each period $r_{\ell,t}$ is calculated as:

$$r_{\ell,t} = w_{\ell,t-\theta} - w_{\ell,t-\theta-1} + r_{\ell,t-\theta}.$$

Finally, we set $w_{\ell,t} < 0.01$ to be zero, to avoid implementation issues with floating point calculations. Setting a lower value to truncate at would improve the performance of the proactive jurisdiction policies, as they would be more sensitive to detected cases in other jurisdictions.

Uniform and Lax Policies We run two simulation variants for both the proactive and internally-based policy: one in which all states are as conservative as possible, setting $x = 1$ and a second in which four regions set a higher threshold of $x = 5$; in the proactive case, these lax regions also act myopically, following the internal jurisdiction-based policy.

We choose $x = 5$ to simulate lax thresholds. In the United States, New York state issued a stay at home order when 0.07% of the state population was infected, which scaled to our populations of 3500 that is equivalent to a threshold of 2.73 (*49, 50*). When scaled to match our population of 3500, Florida began re-opening with a threshold of 6.15, and some countries never locked down (*51, 50, 52*). In our stylized model, quarantines are more aggressive as they cut contact completely.

Results and Sensitivity Analysis

We include the results of the simulations detailed in the main text in the tables below. In addition, we run simulations with two sets of varied parameters: first, we take $\alpha = 0.05$, second we take $\theta = 8$ and $\tau = 5$. Within the United States, estimates for the detection rate range from 5% to 15%, and in countries with less developed testing infrastructure, the detection rate is undoubtedly lower (*37*). Because disease parameters are estimated, we use a different estimated of the disease lifespan of COVID-19 (*48*). For all simulations, we fix $R_0 = 3.5$. As shown in Tables S2-S5, our results are robust to these different sets of parameters.

Supplementary Tables

Table S1: Graph Statistics

Property	Value
Average Degree	20.49
Average Local Clustering Coefficient	0.208
Diameter	9
Average Path Length	5.33

Graph statistics for the graph used in all simulations. Similar to real world networks, it is sparse, clustered and has short average distances between nodes.

Table S2: Regional Policy Simulation Results

θ	τ	α	ϵ	Percent Infected	Infection Person Days	Quarantined Person Days	Escape Rate
5	0	0.1	0	0.0276	1384.05	803955.61	0.0953
5	3	0.1	0	0.226	11282.19	2301413.60	0.458
5	3	0.1	0.05	0.514	25688.08	6478054.64	0.551
5	0	0.05	0	0.0684	3421.10	11231131.73	0.225
5	3	0.05	0	2.81	140667.17	20297075.03	0.623
5	3	0.05	0.05	7.80	390155.83	66067046.93	0.706
8	0	0.1	0	0.0277	2213.92	1243574.65	0.0904
8	5	0.1	0	0.285	22834.58	4187189.53	0.506
8	5	0.1	0.05	0.559	44709.41	10653981.92	0.582

Results for the parameters used in the main text are the average over 10000 simulations. Results for the parameters only used in this section are the average over 2500 simulations. For all simulations, we set $k = 3$ and $x = 1$. Infection person days and quarantined person days are scaled to be per million individuals. The escape rate is defined as the frequency with which the disease escapes the initial quarantine.

Table S3: Global Policy Simulation Results

θ	τ	α	Percent Infected	Infection Person Days	Quarantined Person Days
5	3	0.1	0.0456	2278.84	4725000.00
5	3	0.05	0.0804	4019.67	4698000.00
8	5	0.1	0.0489	3914.26	7507200.00

Results for the parameters used in the main text are the average over 10000 simulations. Results for the parameters only used in this section are the average over 2500 simulations. Infection person days and quarantined person days are scaled to be per million individuals. There are fewer quarantined person days on average with $\alpha = 0.05$, rather than $\alpha = 0.1$ as there is a greater chance of the disease going completely undetected before dying out.

Table S4: Internal and Proactive Policy Simulation Results

Policy	θ	τ	α	Percent Infected	Infection Person Days	Quarantined Person Days	Fraction Requarantined
Internal	5	3	0.1	11.85	592234.08	65634600.00	0.674
Proactive	5	3	0.1	0.630	31501.43	37816130.00	0.569
Internal	5	3	0.05	30.61	1530914.46	133060800.00	0.804
Proactive	5	3	0.05	1.93	96626.34	56930840.00	0.719
Internal	8	5	0.1	10.13	810885.33	87270880.00	0.684
Proactive	8	5	0.1	0.792	63328.75	19992240.00	0.568

Results for the parameters used in the main text are the average over 10000 simulations. Results for the parameters only used in this section are the average over 2500 simulations. For all simulations, every jurisdiction sets $x = 1$. Infection person days and quarantined person days are scaled to be per million individuals.

Table S5: Internal and Proactive Policies with Lax Jurisdictions Simulation Results

Policy	θ	τ	α	Percent Infected	Infection Person Days	Quarantined Person Days	Fraction Requarantined	Low Threshold Case Fraction
Internal	5	3	0.1	20.94	1046945.43	102106475.00	0.728	0.842
Proactive	5	3	0.1	2.73	136557.93	41470367.50	0.674	0.734
Internal	5	3	0.05	34.02	1700815.97	117682100.00	0.822	0.852
Proactive	5	3	0.05	11.64	582078.59	101899270.00	0.806	0.720
Internal	8	5	0.1	19.65	1572688.84	155807360.00	0.744	0.840
Proactive	8	5	0.1	3.23	258570.58	49832880.00	0.677	0.763

Results for the parameters used in the main text are the average over 10000 simulations. Results for the parameters only used in this section are the average over 2500 simulations. For all simulations, 36 jurisdictions set $x = 1$ and the remained set $x = 5$. In the proactive case, jurisdictions with $x = 5$ follow myopic, internal policies. Infection person days and quarantined person days are scaled to be per million individuals.

Supplementary Figures

Figure S1: Growth Balance

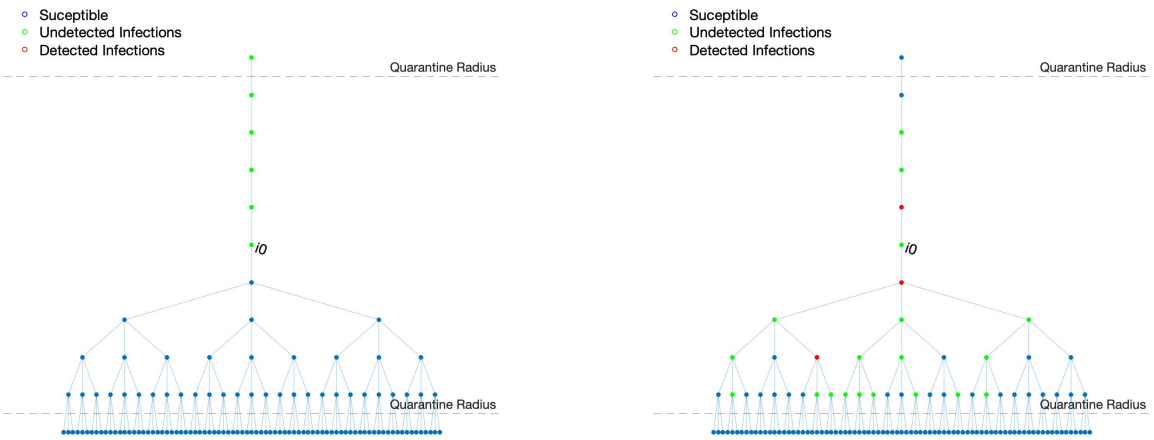

(a) Regional Policy Fails

(b) Regional Policy Succeeds

Figure S1: Panel (a) demonstrates the possible failure of growth-balance. The infection escapes up the line undetected beyond the quarantine radius. If the infection happens to spread downwards, as in Panel (b), it is much more likely to be detected. However, that only happens with some moderate probability in this network, and so growth balance fails.