



Spatial heterogeneity can lead to substantial local variations in COVID-19 timing and severity

Loring J. Thomas^a, Peng Huang^a, Fan Yin^b, Xiaoshuang Iris Luo^c, Zack W. Almquist^d, John R. Hipp^c, and Carter T. Butts^{a,b,e,f,1}

^aDepartment of Sociology, University of California, Irvine, CA, 92697; ^bDepartment of Statistics, University of California, Irvine, CA, 92697; ^cDepartment of Criminology, Law, and Society, University of California, Irvine, CA, 92697; ^dDepartment of Sociology, Center for Studies in Demography and Ecology, Center for Statistics and Social Sciences, eScience, University of Washington, Seattle, WA, 98195; ^eDepartment of Computer Science, University of California, Irvine, CA, 92697; and ^fDepartment of Electrical Engineering and Computer Science, University of California, Irvine, CA, 92697

Edited by Douglas S. Massey, Princeton University, Princeton, NJ, and approved August 18, 2020 (received for review June 6, 2020)

Standard epidemiological models for COVID-19 employ variants of compartment (SIR or susceptible–infectious–recovered) models at local scales, implicitly assuming spatially uniform local mixing. Here, we examine the effect of employing more geographically detailed diffusion models based on known spatial features of interpersonal networks, most particularly the presence of a long-tailed but monotone decline in the probability of interaction with distance, on disease diffusion. Based on simulations of unrestricted COVID-19 diffusion in 19 US cities, we conclude that heterogeneity in population distribution can have large impacts on local pandemic timing and severity, even when aggregate behavior at larger scales mirrors a classic SIR-like pattern. Impacts observed include severe local outbreaks with long lag time relative to the aggregate infection curve, and the presence of numerous areas whose disease trajectories correlate poorly with those of neighboring areas. A simple catchment model for hospital demand illustrates potential implications for health care utilization, with substantial disparities in the timing and extremity of impacts even without distancing interventions. Likewise, analysis of social exposure to others who are morbid or deceased shows considerable variation in how the epidemic can appear to individuals on the ground, potentially affecting risk assessment and compliance with mitigation measures. These results demonstrate the potential for spatial network structure to generate highly nonuniform diffusion behavior even at the scale of cities, and suggest the importance of incorporating such structure when designing models to inform health care planning, predict community outcomes, or identify potential disparities.

COVID-19 | spatial heterogeneity | diffusion | health disparities | social networks

Since its emergence at the end of 2019, severe acute respiratory syndrome coronavirus 2 (SARS-CoV-2) has spread rapidly to all portions of the globe, infecting over 20 million people as of mid-August 2020 (1). The disease caused by this virus, denoted COVID-19, generally manifests as a respiratory illness that is spread primarily via airborne droplets. While most cases of COVID-19 are nonfatal, a significant fraction of those infected require extensive supportive care, and the mortality rate is substantially higher than more common infectious diseases such as seasonal influenza (2). Even for survivors, infection can lead to long-term damage to the lungs and other organs, leading to long convalescence times and enhanced risk of secondary complications (3, 4). By early March of 2020, COVID-19 outbreaks had appeared on almost every continent, including significant clusters within many cities (5). In the absence of an effective vaccine, public health measures to counteract the pandemic in developed nations have focused on social distancing measures that seek to slow diffusion sufficiently to avoid catastrophic failure of the health care delivery system. Both the planning and public acceptance of such measures have been highly dependent upon the use of epidemiological models to probe the potential impact of distancing interventions, and to anticipate when such measures may

be loosened with an acceptable level of public risk. As such, the assumptions and behavior of COVID-19 diffusion models are of significant concern.

Currently, dominant approaches to COVID-19 modeling (6–8) are based on compartment models (often called SIR models, after the conventional division of the population into susceptible, infected, and recovered groups in the most basic implementations) that implicitly treat individuals within a population as geographically well mixed. While some such models include differential contact by demographic groups (e.g., age), and may treat states, counties, or, occasionally, cities as distinct units (e.g., work by ref. 9), those models presently in wide use do not incorporate spatial heterogeneity at local scales (e.g., within cities). Past work, however, has shown evidence of substantial heterogeneity in social relationships at regional, urban, and suburban scales (10–12), with these variations in social network structure impacting outcomes as diverse as regional identification (13), disease spread (14), crime rates (15), neighborhood identification, and development (12, 16). If individuals are not socially “well-mixed” at local scales, then it is plausible that diffusion of SARS-CoV-2 via interpersonal contacts will likewise depart from the uniform mixing characteristic of the SIR models. Indeed, at least one computational study (17) using a fairly “generic” (non-COVID) diffusion process on realistic urban networks has shown considerable nonuniformity in diffusion times, suggesting that such effects could hypothetically be present. Variations across local regions on the pandemic timing, severity,

Significance

We examine the effects of an uneven population distribution on the spread of the COVID-19 disease spread, using a diffusion model based on interpersonal contact networks. Taking into account spatial heterogeneity, the spread of COVID-19 is much “burstier” than in standard epidemiological models, with substantial local disparities in timing and severity and long lags between local outbreaks. We show that spatial heterogeneity may produce dramatic differences in social exposures to those with the illness, and may stress health care delivery systems in ways that are not well captured by standard SIR-like models.

Author contributions: J.R.H. and C.T.B. designed research; L.J.T., P.H., F.Y., and Z.W.A. performed research; L.J.T., P.H., Z.W.A., and C.T.B. contributed new reagents/analytic tools; L.J.T., P.H., F.Y., X.I.L., and Z.W.A. analyzed data; and L.J.T., P.H., F.Y., X.I.L., Z.W.A., J.R.H., and C.T.B. wrote the paper.

The authors declare no competing interest.

This article is a PNAS Direct Submission.

This open access article is distributed under [Creative Commons Attribution License 4.0 \(CC BY\)](https://creativecommons.org/licenses/by/4.0/).

¹To whom correspondence may be addressed. Email: butts@uci.edu.

This article contains supporting information online at <https://www.pnas.org/lookup/suppl/doi:10.1073/pnas.2011656117/-DCSupplemental>.

and the hospital load could have huge impacts on the social outcomes of different population groups (e.g., racial/ethnic groups) in the pandemic, given the heterogeneity of their spatial distribution in urban and suburban areas (18, 19). However, it could also be hypothesized that such effects would be small perturbations to the broader infection curve captured by conventional compartment models, with little practical importance. The question of whether these effects are likely to be present for COVID-19, and, if so, their strength and size, has, to date, remained open.

In this paper, we examine the potential impact of local spatial heterogeneity on COVID-19, modeling the diffusion of SARS-CoV-2 in populations whose contacts are based on spatially plausible network structures. We focus here on the urban context, examining 19 different cities in the United States. We simulate the population of each city in detail (i.e., at the individual level), simulating hypothetical outbreaks on the contact network in each city in the absence of measures such as social distancing. Despite allowing the population to be well mixed in all other respects (i.e., not imposing mixing constraints based on demographic or other characteristics), we find that spatial heterogeneity alone is sufficient to induce substantial departures from spatially homogeneous SIR behavior. Among the phenomena observed are “long lag” outbreaks that appear in previously unharmed communities after the aggregate infection wave has largely subsided, frequently low correlations between infection timing in spatially adjacent communities, and distinct subpatterns of outbreaks found in some urban areas that are uncorrelated with the broader infection pattern. Gaps between infection peaks at the intraurban level can be large, for example, on the order of weeks or months in extreme cases, even for communities that are within kilometers of each other. Such heterogeneity is potentially consequential for the management of health care delivery services: As we show, using a simple “catchment” model of hospital demand, local variations in infection timing can easily overload hospitals in some areas, generating “hospital deserts” (20), while leaving others relatively empty (absent active reallocation of patients). Likewise, we show that individuals’ social exposures to others who are morbid or deceased vary greatly over the course of the pandemic, potentially leading to differences in risk assessment and bereavement burden for persons residing in different locations. Differences in outbreak timing and severity may exacerbate health disparities (since, e.g., surge capacity varies by community) and may even affect perception of and support for prophylactic behaviors among the population at large, with those in so-far untouched communities falsely assuming that the pandemic threat is either past or was exaggerated to begin with, or attributing natural variation in disease timing to the impact of health interventions.

We note at the outset that the models used here are intended to probe the hypothetical impact of spatial heterogeneity on COVID-19 diffusion within particular scenarios, rather than to produce high-accuracy predictions or forecasts. For the latter applications, it is desirable to incorporate many additional features that are here simplified to facilitate insight into the phenomenon of central interest. In particular, we do not incorporate either demographic effects or social distancing (21, 22), allowing us to consider a setting that is as well mixed as possible (and hence as close as possible to an idealized SIR model), with the exception of spatial heterogeneity. As we show, even this basic scenario is sufficient to produce large deviations from the SIR model. Despite the simplicity of our models, we do note that the approach employed here could be integrated with other factors and calibrated to produce models intended for forecasting or similar applications.

Materials and Methods

Spatial Network Data. Networks are generated using population distributions from the most recent US Census in 2010. Network construction

followed the same methodology as Butts et al. (23). Hospital information was obtained from the Homeland Infrastructure Foundation-Level Data (HIFLD) database (24). HIFLD is an initiative that collects geospatial information on critical infrastructure across multiple levels of government. We employ the national-level hospital facility database, which contains locations of hospitals for the 50 US states; Washington, DC; and US territories of Puerto Rico, Guam, American Samoa, Northern Mariana Islands, Palau, and Virgin Islands; underlying data are collated from various state departments or federal sources (e.g., Oak Ridge National Laboratory). We employ all hospitals within our 19 target cities, excluding facilities closed since 2019. Latitude/longitude coordinates and capacity information were employed to create a spatial database that includes information on the number of beds in each hospital. The capacity information includes the number of beds that each hospital has available, and can be used to assess strain that a surge in hospitalizations could create.

The dates of the first confirmed case and all of the death cases for King County, where Seattle is located, were obtained from *The New York Times*, based on reports from state and local health agencies (25). The death rate was calculated based on population size of each county from the 2018 American Community Survey, and employed to calibrate the infection rate (the only free parameter in the models used here); details are provided in *SI Appendix*.

We ran 10 replicates of the COVID-19 diffusion process in each of our 19 cities, seeding with 25 randomly selected infections in each replicate and following the course of the diffusion until no infectious individuals remained. Simulations were performed using a combination of custom scripts for the R statistical computing system (26) and the statnet library (27–29). Analyses were performed using R.

Methods. COVID-19 is typically transmitted via direct contact with infected individuals, with the greatest risk occurring when an uninfected person is within approximately six feet of an infected person for an extended period. Such interactions can be modeled as events within a social network, where individuals are tied to those with whom they have a high hazard of intensive interaction. In prior work, this approach has been successfully employed for modeling infectious diseases ranging from HIV (30) and influenza (31) to Zika (32). To model networks of potential contacts at scale, we employ spatial network models (33), which are both computationally tractable and able to capture the effects of geography and population heterogeneity on network structure (23). Such models have been successfully used to capture social phenomena ranging from neighborhood-level variation in crime rates (15) and regional identification (13) to the flow of information among homeless persons (34).

The spatial network models used here allow for complex social dependence through a kernel function, referred to as the *social interaction function* (SIF). The SIF formally defines the relationship between two individuals based on spatial proximity. For example it has been shown that many social interaction patterns obey the Zipf law (35), where individuals are more likely to interact with others close by rather than far away [a pattern that holds even for online interactions (10)]. Here, we use this approach to model a network that represents the combination of frequent interactions due to ongoing social ties and contacts resulting from frequent incidental encounters (e.g., interactions with neighbors and community members).

We follow the protocol of refs. 15 and 23 to simulate social network data that combine the actual distribution of residents in a city with a prespecified SIF. We employ the model of ref. 15 with decennial Census data to produce large-scale social networks for 19 cities and counties in the United States—providing a representation of major urban areas in the United States (*SI Appendix*). Given these simulated networks, we then implement an individual-level SIR-like framework to examine COVID-19 diffusion. At each moment in time, each individual can be in a susceptible, infected but not infectious, infectious, deceased, or recovered state. The disease diffuses through the contact network, with currently infectious individuals infecting susceptible neighbors as a continuous time Poisson process with a rate estimated from mortality data (*SI Appendix*); recovered or deceased individuals are not considered infectious for modeling purposes. Upon infection, an individual’s transitions between subsequent states (and into mortality or recovery) are governed by waiting time distributions based on epidemiological data as described in *SI Appendix*. To begin each simulated trajectory, we randomly infect 25 individuals, with all others being considered susceptible. Simulation proceeds until no infectious individuals remain.

From the simulated trajectory data, we produce several metrics to assess spatial heterogeneity in disease outcomes. First, we present infection curves

for illustrative cities, showing the detailed progress of the infection and its difference from what an SIR model would posit. We also present choropleth maps showing spatial variation in peak infection times, as well as the correlations between the infection trajectory within local areal units and the aggregate infection trajectory for the city as a whole. While an SIR model would predict an absence of systematic variation in the infection curves or the peak infection day for different areal units in the same city, geographically realistic models show considerable disparities in infection progress from one neighborhood to another. To quantify the degree of heterogeneity more broadly, we examine spatial variation in outcomes for each of our city networks. We show that large variations in peak infection days across tracts are typical (often spanning weeks or even months), and that overall correlations of within-tract infection trajectories with the aggregate urban trajectory are generally modest (a substantial departure from what would be expected from an SIR model).

In addition to these relatively abstract metrics, we also examine a simple measure of the potential load on the health care system in each city. Given the locations of each hospital in each city, we attribute infections to each hospital using a Voronoi tessellation (i.e., under the simple model that individuals are most likely to be taken to the nearest hospital if they become seriously ill). Examination of the potential hospital demand over time shows

substantial differences in load, with some hospitals severely impacted while others have few cases. Finally, we consider the *social exposure* of individuals to COVID-19, by computing the fraction of individuals with a personal contact who is respectively morbid or deceased. Our model shows considerable differences in these metrics over time, revealing that the pandemic can appear very different to those “on the ground”—evaluating its progress by its impact on their own personal contacts—than what would be suggested by aggregate statistics.

Results

Smooth Aggregate Infection Trajectories Can Mask Local Outbreak Dynamics. When taken over even moderately sized regions, aggregate infection curves can appear relatively smooth. Although this suggests homogeneous mixing (as assumed, e.g., by standard SIR models), appearances can be deceiving. Fig. 1 shows typical realizations of infection curves for two cities (Seattle, WA, and Washington, DC), showing both the aggregate trajectory (red) and trajectories within individual Census tracts (black). While the infection curves in both cases are relatively




Fig. 1. (Top Left) Infection curves for Seattle, WA. The red line is the curve for the whole city, while the black lines are the infection curves for each tract in the city. While the red curve is relatively smooth, this smoothness hides a significant amount of heterogeneity in the timing of the infection curves for each census tract. (Top Right) Infection curves for Washington, DC. As with Seattle, the city-level curve conceals considerable spatial variability in the infection's progress. (Bottom Left) Histogram showing the mean pairwise correlation of infection curves for each tract within each city, across our entire sample. The infection curve in any given tract is likely to have a correlation of only around 0.2 with any other tract in the city. This histogram includes a single data point for each tract in the sample. (Bottom Right) Histogram of variance accounted for by the principal component of the standardized tract-level curve set. None of the principal components account for more than 60% of the variance, with most accounting for around only 35% of the total variance. The data points included here include a single amount of variance explained for each city.

smooth, and suggestive of a fairly simple process involving a sharp early onset followed by an initially sharp but mildly slowing decline in infections, within-tract trajectories tell a different story. Instead of one common curve, we see that tracts vary wildly in onset time and curve width, with some tracts showing peaks weeks or months after the initial aggregate spike has passed.

The cases of Fig. 1 are emblematic of a more systematic phenomenon: The progress of the infection within any given areal unit often has relatively little relationship to its progress in the city as a whole. Fig. 1, *Bottom* assesses this phenomenon over our entire sample, using two different consensus metrics. First, we simply compute the correlation between the infection curves in each pair of tracts (assessed at daily resolution), taking the mean for each tract of its correlation with all other tracts within the city; if the progress of the infection were uniform across the city, the mean correlations would be large and positive. Second, we provide a more direct assessment of the extent to which the set of infection curves can be summarized by a common pattern by taking the variance on the first principal component of the correlation matrix generated from the tract-level correlations discussed immediately above. As before, where different parts of the city experience similar patterns of growth and decline in infections, we expect the dimension of greatest shared variance to account for the overwhelming majority of variation in infection rates. Contrary to these expectations, however, Fig. 1 shows that there is little coherence in tract-level infection patterns. Mean correlations of local infection curves across tracts typically range from ~ 0 to 0.5, with a mean of approximately 0.2, indicating very little correspondence between infection timing in one tract and that of another. The principal component analysis tells a similar story: Overall, we see that the first component accounts for relatively little of the total variance in trajectories, with, on average, only around 35% of variation in infection curves lying on the first principal component (and no observed case of the first component accounting for more than 60% of the variance). Interestingly, this variation is not explained by time required for the diffusion process to reach each tract (*SI Appendix, Fig. S4*), in contrast to the hypothesized importance of similar delays in a cross-national context (9). This confirms that local infection curves are consistently distinct, with behavior that is only weakly related to infections in the city as a whole. This is a substantially different scenario than what is commonly assumed in traditional SIR models.




Fig. 2. (Left) Choropleth showing the peak day of infection in each tract in the city of Seattle. The map shows significant variability in peak times, with nearby regions sometimes having sharply different patterns. In outlying parts of Seattle, the infection does not peak until almost a year past the first infections, while, in the more eastern and central parts of the city, the infection peaks much earlier. (Right) Times to peak infection for Washington, DC tracts. The southern and eastern part of the city has infections that are more delayed than in the central and northern parts of the city. Both of these maps show that there is a high degree of spatial heterogeneity present in the infection curves.




Fig. 3. Marginal distributions of days to peak infection by tract, across 10 simulated trajectories. Although locales vary both in terms of overall median peak time and range of tract-level variation, large differences in peak time are nearly ubiquitous. (Trajectory specific distributions are shown in Fig. S2.)

Peak Infection Days Can Vary Substantially, Even among Nearby Regions. These differences in local infection curves are a consequence of the unevenness of the “social fabric” that spans the city: While the disease can spread rapidly within regions of high local connectivity, it can easily become stalled upon reaching the boundaries of these regions. Further transmission requires that a successful infection event occur via a bridging tie, an event with a potentially long waiting time. Such delays create potential opportunities for public health interventions (trace/isolate/treat strategies), but they can also create a false sense of security for those on the opposite side of the bridge (who may incorrectly assume that their area was passed over by the infection). Indeed, examining the time to peak infection across the cities of Seattle and Washington, DC (Fig. 2) shows that, while peak times are visibly autocorrelated, tracts with different peak times frequently border each other. Residents on opposite sides of the divide may be exposed to very different local infection curves, making risk assessment difficult.

The cases of Seattle and Washington, DC are not anomalous. Looking across multiple trajectories over our entire sample, Fig. 3 shows consistently high variation in per-tract peak infection times for nearly all study communities. (This variation is also seen within individual trajectories, as shown in *SI Appendix, Fig. S2*.) Although peak times in some cities are concentrated within an interval of several days to a week, it is more common for peak times to vary by several months or longer. Extreme delays in peak time arise from “slow burn” trajectories associated with very long infection chains, in which sequential transmission to only one or two alters at a time can sustain infection for greatly extended periods. Such gaps can arise naturally in inhomogeneous networks, but are far from what would be expected under uniform local mixing.

Heterogeneous Impact Timing May Affect Hospital Load. Variation in the timing of COVID-19 impacts across the urban landscape has potential ramifications for health care delivery, creating unequally distributed loads that overburden some providers while leaving others with excess resources. To obtain a sense of how spatial heterogeneity in the infection curve could potentially

impact hospitals, we employ a simple “catchment” model in which seriously ill patients are taken to the nearest hospital, subsequently recovering and/or dying as assumed throughout our modeling framework. Based on prior estimates (36), we assume that 14% of all infections are severe enough to require hospitalization (robustness to alternative rate estimates is shown in *SI Appendix*). While hospitals draw from (and hence average across) areas that are larger than tracts, the heterogeneity shown in Fig. 1 suggests the potential for substantial differences in hospital load over time. Indeed, our models suggest that such differences will occur. Fig. 4 shows the number of patients arriving at each hospital in Seattle and Washington, DC during a typical simulation trajectory. While some hospitals do have demand curves that mirror the city’s overall infection curve, others show very different patterns of demand. In particular, some hospitals experience relatively little demand in the early months of the pandemic, only to be hit hard when infections in the city as a whole are winding down.

Just as hospital load varies, hospital capacities vary as well. As a simple measure of strain on hospital resources, we consider the difference between the number of COVID-19 hospitalizations and the total capacity of the hospital (in beds), truncating at zero when demand outstrips supply. (For ease of interpretation as a measure of strain, we take the difference such that higher values indicate fewer available beds.) Using data on hospital locations and capacities, we show, in Fig. 4, strain on all hospitals

in Seattle and Washington, DC during a typical infection trajectory. While some hospitals are hardest hit early on (as would be expected from the aggregate infection curve), others do not peak for several months. Likewise, hospitals proximate to areas of the city with very different infection trajectories experience natural “curve flattening,” with a more distributed load, while those that happen to draw from positively correlated areas experience very sharp increases and declines in demand. These conditions in some cases combine to keep hospitals well under capacity for the duration of the pandemic, while others are overloaded for long stretches of time. These marked differences in strain for hospitals within the same city highlight the potentially complex consequences of heterogeneous diffusion for health care providers.

Looking across cities, we see the same high-variability patterns as observed in Seattle and Washington. In particular, we note that local variation in disease timing leads to a heavy-tailed distribution for the duration at which hospitals will be at capacity. Fig. 4 shows the marginal distribution of hospital overload periods (defined as total number of days at capacity during the pandemic), over the entire sample. While the most common outcome is for hospitals to be stressed for a brief period (not always to the breaking point), a significant fraction of hospitals end up being overloaded for months—or even, in a small fraction of cases, nearly the whole duration of the pandemic.

It should be reiterated that the hospital load model used here is extremely simplified, and that we are employing a




Fig. 4. (Top Left) Numbers of infections attributed to each hospital in the city of Seattle, with each curve representing a different hospital. Hospital peak demand times vary markedly, with some getting the majority of their hospitalizations before day 100, and others peaking almost a year into the pandemic. (Top Middle) Hospitalizations in Washington, DC. As in Seattle, each hospital has a unique demand trajectory, with some hospitals not getting their peak of infections until more than a year after the infection begins. (Bottom Left) Hospital strain in Seattle, WA. Values closer to zero indicate that hospitals are more strained and have fewer open beds, while lower values suggest more resources are available; color varies from blue (low average strain) to red (high average strain). Much like the number of infections, there is a high degree of heterogeneity present here, with hospitals freeing up resources at different points across the first year of the pandemic. (Bottom Middle) Hospital strain for Washington, DC. Most hospitals get overwhelmed in the first 25 days of the pandemic, but then are able to recover at different times, usually within the second hundred days of the pandemic; some, however, are hit hard by a second wave, and others remain overwhelmed for several months. (Right) Marginal distribution of number of days without available beds, for all hospitals in our sample. While most hospitals will have only brief periods of overload, some will be at or over capacity for the entire pandemic, potentially several years.

no-mitigation scenario. However, these results quite graphically demonstrate that the importance of curve-flattening interventions does not abate once geographical factors are taken into account. On the other hand, these results suggest that differences in hospital load may be substantially more profound than would be anticipated from uniform mixing models, creating logistical challenges and possibly exacerbating existing differences in resource levels across hospitals. At the same time, such heterogeneity implies that resource sharing and patient transfer arrangements could prove more effective as load management strategies than would be suggested by spatially homogeneous models, as hospitals are predicted to vary considerably in the timing of patient demand.

Social Exposures to Morbidity and Mortality Vary by Location. In addition to health care strain, the *subjective experience* of the pandemic will potentially differ for individuals residing in different locations. In particular, social exposures to outcomes such as morbidity or mortality may shape individuals' understandings of the risks posed by COVID-19, and their willingness to undertake protective actions to combat infection. Such exposures may furthermore act as stressors, with potential implications for physical and/or mental health. As a simple measure of social exposure, we consider the question of whether a focal individual (ego) either has experienced a negative outcome themselves or has at least one personal contact (alter) who has experienced the outcome in question. (Given the highly salient nature of COVID-19 morbidity and mortality, we focus on the transition to first exposure rather than, e.g., the total number of such exposures, as the first exposure is likely to have the greatest impact on ego's assessment of the potential severity of the disease.)

To examine how social exposure varies by location, we compute the fraction of individuals in each tract who are socially exposed to morbidity or mortality. Fig. 5 shows these proportions for Baltimore, MD, over the course of the pandemic. As with other outcomes examined here, we see considerable variation in timing, with many tracts seeing a rapid increase in exposure to infections, while others go for weeks or months with relatively few persons having a personal contact with the disease. Another notable axis of variation is sharpness. In many tracts, the fraction of individuals with at least one morbid contact transitions from near zero to near one within a matter of days, creating an extremely sharp social transition between the "preexposure world" (in which almost no one present knows someone with the illness) to a "postexposure world" in which almost everyone knows someone with the illness). By contrast, other tracts show a much more gradual increase (sometimes punctuated by jumps), as more and more individuals come to

know someone with the disease. In a few tracts that are never hit hard by the pandemic, few people ever have an infected alter; residents of these areas obviously have a very different experience than those of high-prevalence tracts. These distinctions are even more stark for mortality, which takes longer to manifest and which does so much more unevenly. Tracts vary greatly in the fraction of individuals who ultimately lose a personal contact to the disease, and in the rapidity with which that fraction is reached. In many cases, it may take a year or more for this quantity to be realized; until that point, many residents may be skeptical to the notion that the pandemic poses a great risk to them personally.

By way of assessing the milieu within each tract, it is useful to consider the "cross-over" point at which at least half of the residents who will be socially exposed in a given tract have been socially exposed to either COVID-19 morbidity or mortality. Fig. 6 maps these values for Baltimore, MD. It is immediately apparent that social exposures are more strongly spatially autocorrelated than other outcomes considered here, due to the presence of long-range ties within individuals' personal networks. Even so, however, we see strong spatial differentiation, with residents in the urban core being exposed to both morbidity and mortality much more quickly than those on the periphery. This suggests that the social experience of the pandemic will be quite different for those in city centers compared to those in more outlying areas, with the latter taking far longer to be exposed to serious consequences of COVID-19. This may manifest in differences in willingness to adopt protective actions, with those in the urban core being more highly motivated to take action (and perhaps resistant to rhetoric downplaying the severity of the disease) than those on the outskirts of the city.

Discussion

Our simulation results all underscore the potential effects of local spatial heterogeneity on disease spread. The spatial heterogeneity driving these results occurs on a very small scale (i.e., Census blocks), operating well below the level of the city as a whole. As the infection spreads, relatively small differences in local network connectivity and the prevalence of bridging ties driven by uneven population distribution can lead to substantial differences in infection timing and severity, leading different areas in each city to have vastly different experiences of the pandemic. Resources will be utilized differently in different areas, as some areas will experience the bulk of their infections far later than others, and the subjective experience of a given individual regarding the pandemic threat may differ substantially from someone in another area. These behaviors are in striking contrast to what is assumed by models based on the assumption of




Fig. 5. (Left) Trajectories showing the fraction of people in each tract in Baltimore who have an infected person in their personal network across time. We see a large degree of spatial heterogeneity, as some tracts are more insulated from others in terms of social exposure. However, by the end of the pandemic, most people across all tracts have been exposed to someone who has had the disease. (Right) The fraction of persons in each tract who have an alter who died from COVID-19 in their personal network. On average, only around 20 to 30% of people in any given tract know someone who died, by the end of the pandemic, although this varies widely across tracts.




Fig. 6. (Left) Choropleth showing the time for half of those in each tract to be socially exposed to COVID-19 morbidity in Baltimore, MD. The central parts of the city are exposed far sooner than the northwestern part of the city. (Right) Choropleth showing the time for half of those in each tract to be socially exposed to COVID-19 mortality. Central Baltimore is exposed to deaths in personal networks far sooner than the more outlying areas of the city.

spatially homogeneous mixing, which posit uniform progress of the infection within local areas.

As noted at the outset, our model is based on a no-mitigation scenario, and is not intended to capture the impact of social distancing. While distancing measures by definition limit transmission rates—and will hence slow diffusion—contacts occurring through spatially correlated networks like those modeled here are still likely to show patterns of heterogeneity like those described. One notable observation from our simulations is the long outbreak delay that some census tracts experience, even in the absence of social distancing. This would suggest that relaxation of mitigation measures leading to a resumption of “normal” diffusion may initially appear to have few negative effects, only to lead to deadly outbreaks weeks or months later. Public health messaging may need to stress that apparent lulls in disease progress are not necessarily indicators that the threat has subsided, and that areas “passed over” by past outbreaks could be impacted at any time.

Finally, we stress that conventional diffusion models using locally homogeneous mixing have been of considerable value in both pandemic planning and scenario evaluation. Our findings should not be taken as an argument against the use of such models. However, the observation that incorporating geographical heterogeneity in contact rates leads to radically different local behavior would seem to suggest that there is value in including such effects in models intended to capture outcomes at the city or county level. Since these are the scales on which decisions regarding infrastructure management, health care logistics, and other policies are often made, improved geographical realism could potentially have a substantial impact on our ability to reduce lives lost to the COVID-19 pandemic.

Data Availability. Social networks and analysis code data have been deposited in Harvard Dataverse (<https://doi.org/10.7910/DVN/B9XKSR>) (37).

ACKNOWLEDGMENTS. This material is based on research supported by NSF Awards IIS-1939237 and SES-1826589 to C.T.B., and by a University of California, Irvine Council on Research, Computing and Libraries grant.

- World Health Organization, “Coronavirus disease 2019 (COVID-19): Situation report” (Rep. 205, World Health Organization, 2020).
- G. Onder, G. Rezza, S. Brusaferro, Case-fatality rate and characteristics of patients dying in relation to COVID-19 in Italy. *JAMA* **323**, 1775–1776 (2020).
- F. Jiang et al., Review of the clinical characteristics of coronavirus disease 2019 (COVID-19). *J. Gen. Intern. Med.* **35**, 1545–1549 (2020).
- Y. J. Geng et al., Pathophysiological characteristics and therapeutic approaches for pulmonary injury and cardiovascular complications of coronavirus disease 2019. *Cardiovasc. Pathol.* **47**, 107228 (2020).
- World Health Organization, “Coronavirus disease 2019 (COVID-19): Situation report” (Rep. 43, World Health Organization, 2020).
- M. L. Jackson et al., Effects of weather-related social distancing on city-scale transmission of respiratory viruses. *medRxiv*:10.1101/2020.03.02.20027599 (3 March 2020).
- Y. Zhang, B. Jiang, J. Yuan, Y. Tao, The impact of social distancing and epicenter lockdown on the COVID-19 epidemic in mainland China: A data-driven SEIQR model study. *medRxiv*:10.1101/2020.03.04.20031187 (6 March 2020).
- B. S. Pujari, S. M. Shekatkar, Multi-city modeling of epidemics using spatial networks: Application to 2019-nCoV (COVID-19) coronavirus in India. *medRxiv*:10.1101/2020.03.13.20035386 (17 March 2020).
- D. Brockmann, D. Helbing, The hidden geometry of complex, network-driven contagion phenomena. *Science* **342**, 1337–1342 (2013).
- E. S. Spiro, Z. W. Almqvist, C. T. Butts, The persistence of division: Geography, institutions, and online friendship ties. *Socius* **2**, 2378023116634340 (2016).
- E. J. Smith et al., The relationship of age to personal network size, relational multiplexity, and proximity to alters in the Western United States. *J. Gerontol. B Psychol. Sci. Soc. Sci.* **70**, 91–99 (2015).
- Q. Wang, N. E. Phillips, M. L. Small, R. J. Sampson, Urban mobility and neighborhood isolation in America’s 50 largest cities. *Proc. Natl. Acad. Sci. U.S.A.* **115**, 7735–7740 (2018).
- Z. W. Almqvist, C. T. Butts, Predicting regional self-identification from spatial network models. *Geogr. Anal.* **47**, 50–72 (2015).
- S. Riley, Large-scale spatial-transmission models of infectious disease. *Science* **316**, 1298–1301 (2007).
- J. R. Hipp, C. T. Butts, R. Acton, N. N. Nagle, A. Boessen, Extrapolative simulation of neighborhood networks based on population spatial distribution: Do they predict crime? *Soc. Network.* **35**, 614–625 (2013).
- R. J. Sampson, P. Sharkey, Neighborhood selection and the social reproduction of concentrated racial inequality. *Demography* **45**, 1–29 (2008).
- Z. W. Almqvist, C. T. Butts, Point process models for household distributions within small areal units. *Demogr. Res.* **26**, 593–632 (2012).
- D. S. Massey, N. A. Denton, *American Apartheid: Segregation and the Making of the Underclass* (Harvard University Press, 1993).
- D. S. Massey, J. Tannen, Suburbanization and segregation in the United States: 1970–2010. *Ethn. Racial Stud.* **41**, 1594–1611 (2018).
- M. D. Verhagen, D. M. Brazel, J. B. Dowd, I. Kashnitsky, M. Mills, Mapping hospital demand: Demographics, spatial variation, and the risk of “hospital deserts” during COVID-19 in England and Wales. *OSF*:10.31219/osf.io/g8s96 (21 March 2020).
- J. B. Dowd et al., Demographic science aids in understanding the spread and fatality rates of COVID-19. *Proc. Natl. Acad. Sci. U.S.A.* **117**, 9696–9698 (2020).
- P. Block et al., Social network-based distancing strategies to flatten the COVID 19 curve in a post-lockdown world. *Nat. Human Behav.* **4**, 588–596 (2020).

23. C. T. Butts, R. M. Acton, J. R. Hipp, N. N. Nagle, Geographical variability and network structure. *Soc. Network.* **34**, 82–100 (2012).
24. Department of Homeland Security, Homeland infrastructure foundation-level data (HIFLD). <https://hifld-geoplatform.opendata.arcgis.com/datasets/hospitals>. Accessed 3 April 2020.
25. New York Times, Coronavirus (COVID-19) data in the United States. <https://github.com/nytimes/covid-19-data>. Accessed 5 May 2020.
26. R Core Team, *R: A Language and Environment for Statistical Computing*, (R Foundation for Statistical Computing, Vienna, Austria, 2020).
27. C. T. Butts, network: A package for managing relational data in R. *J. Stat. Software* **24.2**, 1–36 (2008).
28. C. T. Butts, Social network analysis with sna. *J. Stat. Software* **24.6**, 1–51 (2008).
29. M. S. Handcock, D. R. Hunter, C. T. Butts, S. M. Goodreau, M. Morris, statnet: Software tools for the representation, visualization, analysis and simulation of network data. *J. Stat. Software* **24**, 1–11 (2008).
30. M. Morris, *Network Epidemiology: A Handbook for Survey Design and Data Collection* (Oxford University Press, 2004).
31. C. Viboud *et al.*, Synchrony, waves, and spatial hierarchies in the spread of influenza. *Science* **312**, 447–451 (2006).
32. L. Li *et al.*, Analysis of transmission dynamics for Zika virus on networks. *Appl. Math. Comput.* **347**, 566–577 (2019).
33. C. T. Butts, R. M. Acton, “Spatial modeling of social networks” in *The Sage Handbook of GIS and Society Research*, T. L. Nyerges, Ed. (SAGE, Thousand Oaks, CA, 2011), pp. 222–250.
34. Z. W. Almquist, Large-scale spatial network models: An application to modeling information diffusion through the homeless population of San Francisco. *Environ. Plann. B* **47**, 523–540 (2020).
35. G. K. Zipf, *Human Behavior and the Principle of Least Effort: An Introduction to Human Ecology* (Ravenio, 2016).
36. E. K. Stokes *et al.*, Coronavirus disease 2019 case surveillance—United States, January 22–May 30, 2020. *MMWR (Morb. Mortal. Wkly. Rep.)* **69**, 759–765 (2020).
37. L. J. Thomas *et al.*, Replication data for: Spatial heterogeneity can lead to substantial local variations in COVID-19 timing and severity. Harvard Dataverse. <https://doi.org/10.7910/DVN/B9XKSR>. Deposited 26 August 2020.

1

2 **Supplementary Information for**

3 **Spatial Heterogeneity Can Lead to Substantial Local Variations in COVID-19 Timing and**

4 **Severity**

5 Loring J. Thomas, Peng Huang, Fan Yin, Xiaoshuang Iris Luo, Zack W. Almquist, John R. Hipp, Carter T. Butts

6 **Carter T. Butts.**

7 **E-mail: buttsc@uci.edu**

8 **This PDF file includes:**

- 9 Supplementary text
- 10 Figs. S1 to S5
- 11 Tables S1 to S2
- 12 SI References

13 Supporting Information Text

14 Introduction

15 In this supplement, we go into more depth on Spatial Interaction Functions, Spatial Bernoulli Models, the setup and
16 parameterization of our simulations, and the parameter estimates that were used for this paper. We also provide additional
17 analyses regarding the role of local first-passage times as determinants of individuals' waiting times to infection, and the impact
18 of alternative estimates of hospitalization rates on numbers of days at which hospitals would be expected to be at capacity.

19 Spatial Interaction Function

20 A Spatial Interaction Function (SIF) describes the marginal probability of a tie between any two nodes, given the distance
21 between them. We denote the SIF by $\mathcal{F}(\mathcal{D}_{ij}, \theta)$, with $\mathcal{D}_{ij}$ being the distance between vertices i and j , and θ being the
22 parameters for the function. Prior literature has found that spatial interaction functions for social networks like those of
23 interest here tend to be of the power law or attenuated power law form (1). Following this, we employ SIFs of the form
24 $\mathcal{F}(\mathcal{D}_{ij}, \theta) = \frac{p_b}{(1 + \alpha \mathcal{D}_{ij})^\gamma}$. Here, p_b represents the base tie probability, which can be thought of as the probability of a tie between
25 two individuals residing at the same location. α is a scaling parameter that determines the phenomenological unit of distance
26 for the decay in tie probability, and γ is the parameter that determines the weight of the tail (higher values imply fewer
27 long-range ties, *ceteris paribus*).

28 We employ two SIFs in this paper, using models for social interactions and face-to-face interactions employed in prior
29 studies (2, 3). The social interaction SIF declines with a γ of 2.788, while the face-to-face SIF declines with γ of 6.437. The
30 parameters for the social interaction SIF are $p_b = 0.533$, $\alpha = 0.032$, $\gamma = 2.788$, and the parameters for the face-to-face SIF are
31 $p_b = 0.859$, $\alpha = 0.035$, $\gamma = 6.437$ (3).

32 Spatial Bernoulli Models

33 The Bernoulli Network Models are a class of random graph models in which each edge occurs as a Bernoulli trial, possibly with
34 a distinct probability of occurrence. In a spatial Bernoulli graph, tie probabilities are determined by a Spatial Interaction
35 Function, applied to the pairwise distances between individuals within some space (here, geographically determined using
36 Census data). Spatial Bernoulli models are highly scalable due to the conditional independence of edges, but allow for extremely
37 complex structure due to the heterogeneity in edge probabilities induced by the SIF; likewise, they naturally produce properties
38 such as local cohesion and degree heterogeneity observed in many types of social networks (2). Formally, we can specify a
39 Spatial Bernoulli Model by $\Pr(Y_{ij} = 1) = \mathcal{F}(\mathcal{D}_{ij}, \theta)$, where Y_{ij} is a dichotomous indicator for the presence of the i, j edge, and
40 $\mathcal{F}(\mathcal{D}_{ij}, \theta)$ is a Spatial Interaction Function taking as inputs the i, j distance $\mathcal{D}_{ij}$ and parameters θ .

41 Network Simulations

42 To simulate diffusion of COVID-19, we require a contact network. Here, we employ the above-described spatial Bernoulli
43 graphs, with node locations for each of our 19 study locations drawn based on block-level Census data (including clustering
44 within households, an important factor in disease diffusion). We follow the protocols described in (2, 4) to generate node
45 positions, specifically using the quasirandom (Halton) placement algorithm. Node placement begins with the households in
46 each census block, using Census 2010 data with regions defined per (3). The quasirandom placement algorithm uses a Halton
47 sequence to place households in space within the areal unit in which they reside. If any two households are placed within a
48 critical radius of each other, then the algorithm “stacks” the households on top of each other by introducing artificial elevation
49 (simulating e.g. a multistory apartment building). Once all households are placed, individuals within households are placed at
50 jittered locations about the household centroid. (Individuals not otherwise attached to households are treated as households of
51 size 1.)

52 Given an assignment of individuals to spatial locations, we simulate spatial Bernoulli graphs using the models specified above.
53 We generate two networks for each city, one with the social interaction SIF, and the other with the face-to-face interaction
54 SIF. To form a network of potential high-risk contacts, we then merge these networks (which share the same node set) by
55 taking their union, leading to a network in which two individuals are tied if they either have an ongoing social relationship or
56 would be likely to have extensive face-to-face interactions for other reasons (e.g., interacting with neighbors). This process is
57 performed for each city in our sample.

58 **List of Cities.** Table S1 lists the cities that we use for our simulations. These data are drawn from (3), with population data
59 updated to reflect the most recent (2010) decennial Census.

60 Disease Simulations

61 We conduct a series of simulations to examine the spread of COVID-19 across city-sized networks. These simulations use a
62 simple continuous-time network diffusion process, the general description of which are described in the main text. The input for
63 the diffusion simulation is a network and a vector of initial disease states (*susceptible*, *latent* (infected but not yet infectious),
64 *infectious*, *recovered*, and *deceased*), and the output is detailed history of the diffusion process up to the point at which a steady
65 state is obtained (i.e., no infectious individuals remain). Infection occurs via the network, with currently infectious individuals

Table S1. List of study communities.

	City/County
1	Buffalo
2	Baltimore
3	Cincinnati
4	Cleveland
5	Denver
6	Indianapolis
7	Miami
8	Milwaukee
9	Nashville
10	Pittsburgh
11	Rochester
12	Sacramento
13	Salinas
14	San Diego City
15	Seattle
16	St. Petersburg
17	Tampa
18	Tucson
19	Washington DC

66 infecting susceptible alters as Poisson events with a fixed rate. The transitions between latent and infectious, and infectious
67 and either recovery or mortality are governed by gamma distributions estimated from epidemiological data. Table S2 shows
68 the estimated shape and scale parameters for the gamma distributions employed here. The parameters for waiting time to
69 infectiousness are directly available in the Appendix of (5), while those for the recovery and death are estimated by matching
70 the mean and standard deviation of durations reported in the literature (6). Selection into death versus recovery was made via
71 a Bernoulli trial drawn at time of infection (thereby determining which waiting time distribution was used), with the estimated
72 mortality probability being 0.0138 using the case fatality rate adjusted for under-ascertainment reported in (6). Under these
73 parameters, the median time to infectiousness is 5.2 days, with 95% of cases falling between 2 and 10.1 days; once infectious,
74 respective median times to recovery and death are 25.1 days (95% range 9.6 to 52.2) and 16.9 days (95% range 5.5 to 38.3).
75 We note in particular that these distributions incorporate the clinically observed skewness in recovery times, with many cases
76 resolving in less than two weeks but a non-trivial fraction persisting for six weeks or longer.

Table S2. Shape and Scale parameters for Gamma distributions for durations (unit: day).

	Death	Recovery	Infectious
Shape	4.566	5.834	5.807
Rate	0.251	0.219	1.055
Scale (i.e., 1/Rate)	3.984	4.566	0.948

77 Infection Rate Parameter Estimation

78 To determine the infection rate (the only free parameter for the models used in our simulations), we simulate the diffusion of
79 virus in Seattle and fit it to the over-time death rate of the King County, WA before the first shelter-in-place order went into
80 effect on March 23, 2020. We limit our data to this time period because our simulation employs a no-mitigation scenario. A
81 grid search strategy was employed to determine the expected days to transmission (which is the inverse of infection rate), and
82 the number of days between the existence of the first infected cases and the first confirmed cases (aka the time lag, a nuisance
83 parameter that is relevant only for estimation of the infection rate). The time lag is treated as an integer and the expected
84 days to transmission as a continuous variable. For each lag/rate pair, we randomly take 5 draws from the expected infection
85 waiting time distribution, add them to the lag time (i.e. the introduction of the true patient zero for the initial outbreak),
86 and simulate 50 realizations of the diffusion process (redrawing the network each time). The diffusion rate parameter was
87 selected based on minimizing the mean squared error between the simulated death rate and the observed death rate over the
88 selected period. The first round of grid-search divided the expected days of search into 100 intervals, from (0,1) to (99,100),
89 with days of lag ranging from 1 to 100 days. The second round of grid-search, based on the performance of the first round,
90 divided the expected days of search into 240 intervals, from (40.00,40.25) to (99.75,100.00), with days of lag ranging from 1 to
91 60 days. The grid-search suggests that the expected days to transmission is 82.875 (82.75,83.00) days (Fig S1); that is, in a
92 hypothetical scenario in which a single infective ego remained indefinitely in the infective state, and a single alter remained
93 otherwise susceptible, the average waiting time for ego to infect alter would be approximately 80 days. While this may at
94 first blush appear to be a long delay, it should be borne in mind that this embodies the reality that no individual is likely to
95 infect any *given* alter within a short period (since, indeed, ego and alter may not happen to interact within a narrow window).
96 With many alters, however, the chance of passing on the disease is quite high. Likewise, we note that the thought experiment

above should not be taken to imply that actors remain infectious for such an extended period of time; per the above-cited epidemiological data, individuals typically remain infectious for roughly 25.1 days (95% range 9.6-52.2). When both delay times are considered, the net probability of infecting any given alter prior to recovering is approximately 27%. We further calculated the corresponding basic reproductive number (R_0), which is the product of the probability of infection (27%) and the mean degree of the networks (10). The corresponding basic reproductive number in the diffusion simulation model is 2.7, compatible with estimates of R_0 at the pre-mitigation stage in other literature (7–9).

Simulation Replicates

To supplement the results on the variation in the peak infection time given in the main text, we ran a series of simulation replicates. Figure 3 in the main text shows the data from the figure below aggregated across all replicates. In the supplemental Figure S2, we break out the peak infection days in each city, by replicate. These data show that the significant variation in Figure 3 is not due to the number of replicates that were run, but instead due to the intrinsic variation that is present (i.e., spatial heterogeneity).

Timing and Shape of Infection Curves

While our simulations show substantial heterogeneity in infection patterns across tracts, it may be hypothesized that this pattern is driven by a characteristic pattern of infection *within tracts*, with the heterogeneity being driven by the differences in *arrival time* of the infection to different tracts. (Such a pattern has been hypothesized by e.g. (10) at the national level to explain heterogeneity at global scales.) To examine this possibility, we ran several additional analyses. First, to assess whether diffusion patterns within tracts follow a *universal curve*, we compare standardized within-tract prevalence curves across the entire sample of cities. To standardize the infection curves, we apply the following transformations. First, all selected infection prevalence curves (i.e., active cases as a function of time) were translated to the origin, such that the first infection occurs at time 0. Next, we standardized the maximum prevalence by dividing the number of infections active on any day by the maximum prevalence for the tract. Finally, we standardized the time scale of the prevalence curve so that the mean prevalence time (i.e., the centroid of the temporal distribution) occurs at unit time. The resulting standardized curves reflect the pure shape of the infection trajectory; if infection follows a universal pattern at local scales (up to an affine transformation), then the standardized curves should be approximately identical.

To assess the standardized curves, we chose a stratified sample of tracts from across the sample, selecting 5 tracts at random from each city. To avoid artifacts from tracts with insufficient infections, we exclude tracts with fewer than 20 infections during the time course of the simulation. Figure S3 shows the resulting distribution of standardized prevalence curves. In line with the universality hypothesis, there is a dominant functional form approximately followed by a large fraction of tracts. However, we also observe a large fraction of tracts that deviate from this central pattern. Modes of deviation are idiosyncratic, and include multi-modality, substantial differences in skewness, and differences in “roughness” over time. Quantitatively, relatively large variations in standardized time to maximum prevalence are observed, with maxima for the majority of the sample falling between approximately 0.5 and 1.5 time-to-mean-prevalence units. These observations suggest that, while there is a fairly *common* tract-level infection pattern, this pattern coexists with a wide range of other types of trajectories and is better thought of as a central tendency than a universal phenomenon.

Even if there is not a universal tract-level diffusion curve, it could still be the case that the majority of heterogeneity in infection times could be explained by the time taken for the infection to reach each tract: in particular, if the city-level diffusion process reflects a union of very small outbreaks with the waiting time for the infection to “jump” from one area to another being large compared to the time needed for the local outbreaks to reach all susceptibles within the local area, then the primary determinant of individual infection time would be the time taken for the infection to reach the individual’s local region. We examine this hypothesis in figure S4. The left-most boxplot of figure S4 shows the distribution of times to infection for all individuals in our sample, timed from the onset of the first infection. The right-most boxplot shows the corresponding waiting times for the same individuals, *net of the first appearance of the infection in their tracts*. If infection time were determined primarily by the time it takes for the infection to reach an individual’s local area, we would see a substantially compressed distribution relative to the total waiting time; on the contrary, the two are nearly the same, falsifying the conjecture that diffusion to the local area is the driver of waiting time heterogeneity.

The reason for this lack of compression can be seen in the middle boxplot of figure S4, which shows the distribution of waiting times for the first case in each tract. We can see that, while individual infection times vary markedly, first passage times to tracts are both short and highly compressed: once introduced, the infection diffuses rapidly to nearly all tracts. Further, we can see that nearly all tracts are reached long before the majority of individuals are infected – thus, diffusion to local areas is not the primary limiting factor determining individual waiting times. Instead, the relative permeability of those areas to diffusion (which may depend upon factors such as local population density and the presence of barriers to interaction) appears to play a much greater role in governing the distribution of infection times.

Robustness of Spatial Heterogeneity on Hospital Load

Currently, considerable uncertainty exists regarding the fraction of SARS-CoV-2 infections leading to hospitalization; as such, it is useful to verify that the overall patterns of spatial heterogeneity seen here are robust to the hospitalization rate. On the high end, the World Health Organization and the CDC have estimated the rate to be 20% and 20-30%, respectively, in the absence

of mitigation (11, 12). Arguably, these estimates are inflated by under-reporting of asymptomatic or mildly symptomatic infections, particularly given the poor state of testing in the early stage of the pandemic. At the opposite extreme, one recent serology study contended that less than 10% of infections were reported in USA (13), implying a potential hospitalization rate in the neighborhood of 2%. This study, too, faces problems with selection, as population prevalence was based on a convenience sample from patients seeking health care during the pandemic (and who are hence disproportionately likely to be infected); the true extent of under-reporting is thus likely to be smaller than this estimate (and the hospitalization rate correspondingly higher). Taking these two estimates as upper and lower bounds on the likely rate of COVID-19 hospitalizations, we replicated our analysis on hospital load using 20% and 2% (i.e., 10% of 20%) as respective probabilities of hospitalization per infection.

Fig. S5 shows the respective marginal distributions of hospital overload periods with 20% and 2% hospitalization rates. While the scale of full-capacity days varies with the selection of the hospitalization rate, the shape of these distributions closely resembles that of Fig. 7 in the main text: despite the majority of hospitals running at capacity for a relatively short period of time, a sizable fraction of hospitals experience overloads for very long periods. The persistence of this pattern over a full order of magnitude variation in hospitalization rates demonstrates the potential of unmitigated COVID-19 infections to severely strain local resources even under fairly optimistic scenarios, and suggests that substantial inequalities in healthcare service demand are robust to detailed hospitalization rates. However, we also note that the current level of uncertainty in hospitalization rates serves as a significant obstacle to quantitative prediction of the spatial distribution of hospital load for planning or response purposes. Particularly given the large gap between anticipated load on “typical” units and those expected to be hit hardest, more refined rate estimates would seem to have the potential to inform important resource allocation decisions.

Code and Data Availability

Code and data needed to replicate the simulation and analysis for this paper can be found at:

Loring J. Thomas; Peng Huang; Fan Yin; Xiaoshuang Iris Luo; Zack W. Almquist; John R. Hipp; Carter T. Butts, 2020, “Replication Data for: Spatial Heterogeneity Can Lead to Substantial Local Variations in COVID-19 Timing and Severity,” <https://doi.org/10.7910/DVN/B9XKSR>, Harvard Dataverse.

References

- Butts CT, Acton RM (2011) Spatial modeling of social networks. *The Sage Handbook of GIS and Society Research*. Thousand Oaks, CA: SAGE Publications pp. 222–250.
- Butts CT, Acton RM, Hipp JR, Nagle NN (2012) Geographical variability and network structure. *Social Networks* 34(1):82–100.
- Hipp JR, Butts CT, Acton R, Nagle NN, Boessen A (2013) Extrapolative simulation of neighborhood networks based on population spatial distribution: Do they predict crime? *Social Networks* 35(4):614–625.
- Almquist ZW, Butts CT (2012) Point process models for household distributions within small areal units. *Demographic Research* 26:593–632.
- Lauer SA, et al. (2020) The incubation period of coronavirus disease 2019 (COVID-19) from publicly reported confirmed cases: estimation and application. *Annals of Internal Medicine*.
- Verity R, et al. (2020) Estimates of the severity of coronavirus disease 2019: a model-based analysis. *The Lancet Infectious Diseases*.
- Li R, et al. (2020) Substantial undocumented infection facilitates the rapid dissemination of novel coronavirus (SARS-CoV-2). *Science* 368(6490):489–493.
- Li Q, et al. (2020) Early transmission dynamics in Wuhan, China, of novel coronavirus-infected pneumonia. *New England Journal of Medicine* 382(13):1199–1207.
- Wu JT, Leung K, Leung GM (2020) Nowcasting and forecasting the potential domestic and international spread of the 2019-nCoV outbreak originating in Wuhan, China: a modelling study. *The Lancet* 395(10225):689–697.
- Brockmann D, Helbing D (2013) The hidden geometry of complex, network-driven contagion phenomena. *Science* 342(6164):1337–1342.
- CDC COVID and Response Team (2020) Severe outcomes among patients with coronavirus disease 2019 (COVID-19)—United States, February 12–March 16, 2020. *MMWR Morb Mortal Wkly Rep* 69(12):343–346.
- World Health Organization (2020) *Media Statement: Knowing the risks for COVID-19*.
- Havers FP, et al. (2020) Seroprevalence of antibodies to SARS-CoV-2 in 10 sites in the United States, March 23–May 12, 2020. *JAMA Internal Medicine*.




Fig. S1. Grid-search for infection rate. The mean squared error (in logarithm form) for simulations with combinations of days of lag and expected days to transmission, for the first round (A) and the second round (C). The cross-sectional analysis of the minimum log error for each days of lag and expected days to transmission (B) suggests the interval of both variables for the second round of search (areas in gray). Curves of death rate based on the best-fit parameter: 44 days of lag, 82.875 expected days to transmission (D).




Fig. S2. Boxplot showing the peak infection days across 10 replicates for each city in the sample. There is a large degree of heterogeneity within each city, showing that the day that the infection peaks for any given tract is not uniform at all. Within any given city, there is a consistently high amount of variance in the peak infection day. In other words, the variance that we show here is a property of the spread of the disease, rather than the number of simulation replicates.

Normalized Tract Prevalence Curves across the Sample




Fig. S3. Standardized infection curves at the tract level. Many tracts approximately follow the stereotypical “bell-shaped” pattern, but a large number of tracts deviate by being irregular, multimodal, and/or long-tailed. (Note: all trajectories scaled to unit maximum; apparent truncation actually reflects between-curve variation in time to maximum prevalence.)




Fig. S4. Respective distributions of individual infection times (left), tract arrival times (middle), and tract specific infection times (right) for the tract sample. Individual infection times relative to arrival within tracts show little difference from infection times relative to the start of the larger outbreak, as tract arrival times are substantially shorter than the time needed to diffuse within tracts.




Fig. S5. (Left) The distribution of days that hospitals are at capacity, with a 20% hospitalization rate. (Right) The distribution of days that hospitals are at capacity with a 2% hospitalization rate. While higher hospitalization rates lead to higher levels of strain overall, extreme inequality in load persists in both scenarios.