February 2021

Perspective
EXPERT INSIGHTS ON A TIMELY POLICY ISSUE

TIMOTHY R. GULDEN, GAVIN S. HARTNETT, RAFFAELE VARDAVAS, DAVID KRAVITZ

Protecting the Most Vulnerable by Vaccinating the Most Active

As the United States begins to administer vaccines for coronavirus disease 2019 (COVID-19), it is essential that good decisions are made about how to deploy them. How the United States manages the first rounds of vaccination will either save tens of thousands of lives and enable a full reopening of the economy in a matter of months or cost tens of thousands of lives and require continued precautions for another year or more. Our new research, which models the spread of disease through a realistic contact network, indicates that the United States might be able to provide more protection for vulnerable people by vaccinating those with many contacts than by vaccinating vulnerable people directly.

Recent models of COVID-19 vaccination have explored whom to prioritize for vaccination depending on how fast the vaccine will be rolled out and made available, but these models have tended not to take the person-to-person contact structure into account (Bubar et al., 2020).[1] These models, like most applied models in epidemiology, use simplifying assumptions about how the web of human contact is organized, with sophisticated models (such as those cited above) accounting for different contact rates within and between age groups but retaining the assumption that people interact at random. The models use carefully calibrated estimates of how much the average person in each age group has contact with people in the same age group and in each of the other age groups—but treat all of the people in each age group as having essentially the same (or nearly the same) amount of contact as one another. As a hypothetical example, if data indicate that people between the ages of 20 and 29 have an average of 15 contacts per day with others in their age group but only two contacts per day with those in the 75 and older age group, most models would assume that everyone between the ages of 20 and 29 has approximately 15 contacts with people of similar age and two contacts with people in the 75 and older age group. In reality, a person between the ages of 20 and 29 who works at a senior center would have distinctly different contact patterns than one who works as a bartender.

These models have great explanatory and predictive power but might not be well suited to examining the full variety of vaccination strategies—particularly ones that rely on the detailed structure of the contact network and large variations in contacts across people of similar age. This has led to potentially faulty recommendations about what the national vaccine strategy should be in the coming months by suggesting that directly vaccinating the vulnerable provides them with greater protection than vaccinating any other group. Our work found that, when the disease is modeled across a full, individual-based contact network that includes realistic variation in the amount of person-to-person contact, a strategy of vaccinating the most-active people might be superior in many respects—including in its ability to protect the vulnerable.

The relevant network science analysis can become mathematically involved, and, more often than not, it relies on interpreting the results of computer simulations. However, the key intuitions are simple to understand. There is a long and established literature about how pandemics spread across networks and how the spread can be stopped or mitigated through vaccination.[2] This work has seen limited application because researchers did not have good data on the large-scale structure of human contacts. Now, armed with a more realistic contact network, realized at the level of an entire city, researchers are in a position to overcome some of these limitations.

Choosing a Vaccination Strategy

There are a variety of ways to think about the objectives of a vaccination strategy, but common sense and human

> The best way to protect the vulnerable might be to prioritize vaccinating the critical workers and other active people who play the largest role in spreading the virus.

compassion suggest that the United States should focus on protecting the vulnerable.[3] Starting with that goal, it is then a small intellectual step to adopt a strategy of vaccinating the vulnerable first: in particular, the elderly and those with weakened immune systems. However, when the actual structure of human interactions is taken into account, this no longer appears to be the best way to proceed. The best way to protect the vulnerable might be to prioritize vaccinating the critical workers and other active people who play the largest role in spreading the virus. This strategy has the potential to reduce the spread of the virus so radically that everyone, including the vulnerable, will be safer. It might also permit the full reopening of the economy and a return to normal life much sooner than would be possible with a strategy focused on vaccinating the vulnerable first.

To understand this point, it is crucial to remember that this disease is not an external threat in which everyone is equally vulnerable, independent of other people. If everyone faced the same risk of exposure but some people were more likely than others to be harmed, it would make sense to vaccinate the vulnerable first, but COVID-19 is not like that. Humans are both the victims of the disease and the source of its threat. People get the virus from—and pass it to—one another. This means that the pandemic should be viewed as a dynamic phenomenon that spreads across a social network through person-to-person contacts. Thinking about disease transmission as something that happens across a contact network highlights the fact that those who become sick are also the source of new infections. As a result, people with many connections are both more likely to catch the disease and more likely to transmit it to others. This fact is the basis for the physical-distancing measures that have been enacted across the globe.

The central lesson is that an individual's role in spreading a disease very much depends on how many contacts a person has. A bit of simple math shows just how important this contact number is. Start by imagining a person with just one contact per day. That person has one chance to get the disease every day from that contact. Similarly, if the person does get infected, then they have one chance each day to pass the infection on to someone else. Compare the first person with someone who has two contacts per day. All else being equal, the second person is twice as likely to get the disease and, if infected, twice as likely to pass the infection to someone else. Because these probabilities compound in a simple way when the infected fraction of the population is relatively small, this second person is about 2 x 2 = 4 times as likely to end up spreading the disease, compared with the person with only one contact.[4]

Next, consider someone with not two, but 200 contacts per day (e.g., a bartender or an ER nurse). This person is 200 times more likely to be infected and 200 times more likely to pass this infection on, compared with a person with just one contact. That makes this person about 200 x 200 = 40,000 times more likely to end up spreading the disease, compared with the person with one contact, and 10,000 times more likely to end up spreading the disease than the person with two contacts. These high-contact people are extremely important in the spread of the virus throughout the population and should be vaccinated first in order to limit the spread of the virus. Loosely speaking, it is 10,000 times more effective in halting the virus to vaccinate the person with 200 contacts than it is to vaccinate the person with two contacts.[5]

The importance of high-contact people in the spread of the disease suggests a simple vaccination strategy:

Vaccinate the highest-contact people first. Although these might not be the most-vulnerable people, by vaccinating them, many additional infections should be prevented.

Any strategy should be evaluated against a set of criteria that can usually be stated in terms of minimizing some critical quantity: deaths, lost years, strain on health care system capacity, inequality of care, economic loss, etc. In the short term, these criteria are sometimes in conflict: Temporarily and partially shutting businesses hurts the economy but saves lives, and staying open might prove even more costly in the long run. Our preliminary analysis indicated that the vaccination of high-contact people does not present such a conflict. It has the potential to minimize the size, duration, deadliness, and economic impact of the pandemic compared with other vaccination strategies, making it the preferred approach according to any of these metrics.

A Realistic Human Contact Network

To assess the relative merits of a variety of vaccination strategies, we used a simple yet powerful computer simulation that captures (1) only the most-essential features of the strategies and (2) the dynamics of the disease spreading across a real human contact network.

The model is based on a city-scale network of human contacts that we derived from a data set of 2.2 billion mobile device location points that were compiled by UberMedia, mostly for use in marketing applications. These data were collected from a very large number of mobile device apps (approximately 150,000) where users had agreed to share their location information with the app developer and its partners. We used big-data methods to turn these many points into a plausible contact network that combined behavior from before the pandemic and after the beginning of physical distancing in March and April 2020 to produce a network that we believe is representative of movements and interactions of people at those times.[6]

We kept the modeling as simple as possible by considering a highly idealized representation of the disease called an SIR model. In this model (Figure 1), people are in one of three states: susceptible (S), infectious (I), or removed (R). Because there is no natural immunity to COVID-19, everyone begins in the susceptible state. If a susceptible person comes into contact with an infectious person, that person has some probability of becoming infectious as well. This probability is dependent on the time that they spend in contact. After some time (in this simple case, we used 14 days), infected people are removed (through recovery or death, between which we did not distinguish). Removed individuals are no longer susceptible to the disease and play no further role in the simulation. Although there is mixed evidence about the duration of immunity from both illness and vaccination, this is one of the many fine points of COVID-19 dynamics that the model does not include. This

FIGURE 1

Fixed Period SIR Transmission Model

NOTE: The I to R transition occurs after a fixed period of 14 days in this SIR transmission model.

model primarily explores short- to medium-term disease dynamics where loss of immunity is not a major concern.

This model differs from the classical SIR model found in introductory texts on epidemiology in that it does not assume an evenly mixed population (which would allow the model to be analyzed with ordinary differential equations). Instead, we used the contact network derived from mobile device data to model each contact explicitly in terms of both the individuals involved and the time that they spent in contact with others.[7]

Using this model, we simulated five vaccination strategies that are designed to be broadly representative of a class of real strategies. In each case (except the no vaccination case), we vaccinated 15 percent of the population with a perfect vaccine. These numbers are designed to correspond to the volume and efficacy of the vaccine during the first three months of vaccine availability. For the United States, this corresponds to administering a 90-percent effective vaccine to an average of 16.5 million people per month (out of a population of 330 million people) for those first three months.[8]

Testing Five Vaccination Models

We explored the following five cases:

- No vaccination: This is the base case, where the epidemic is allowed to run its course with no intervention.
- Low contact: The 15 percent of people vaccinated are those with the fewest contacts. This corresponds to vaccinating people who are already identified as high risk and who are able to limit contacts.
- Uniform: The 15 percent of the population is vaccinated at random.
- High-contact imperfect: One-half of the 30 percent of people with the most contacts are vaccinated.
- High contact: The 15 percent of people with the most contacts are vaccinated.

The progression of the five strategies for a single run of the simulation is shown in Figure 2. Table 1 contains key metrics averaged across 100 different runs.

The differences between these strategies are stark. The low-contact strategy looks much like the no vaccination strategy. The vaccinated people become immune, but because they are relatively unlikely to pass the disease on to others, vaccinating them does little to change the course of the pandemic. In this model, vaccinating the 15 percent of people with the fewest connections reduces the number of infections by only about 1 percent. The uniform strategy cuts the peak number of infected people nearly in half, and the total number of infected people by approximately one-third—which is a meaningful improvement. The high-contact vaccination strategies are quite different. The pure high-contact strategy, where the 15 percent of people with the most connections are vaccinated, almost immediately crushes the epidemic—making it impossible for the disease to spread and cutting the number of infections that occur after vaccination by 96 percent. Of course, such an approach is impossible in practice both because it is not possible to identify all of the people with large numbers of contacts and because not all high-contact people will agree to be vaccinated. But the high-contact imperfect strategy shows that, even if one-half of the high-contact people are missed, a strategy targeted at high-contact people is much more effective than an untargeted strategy—almost

FIGURE 2

Comparison of Different Vaccination Strategies for a Single Simulation Run

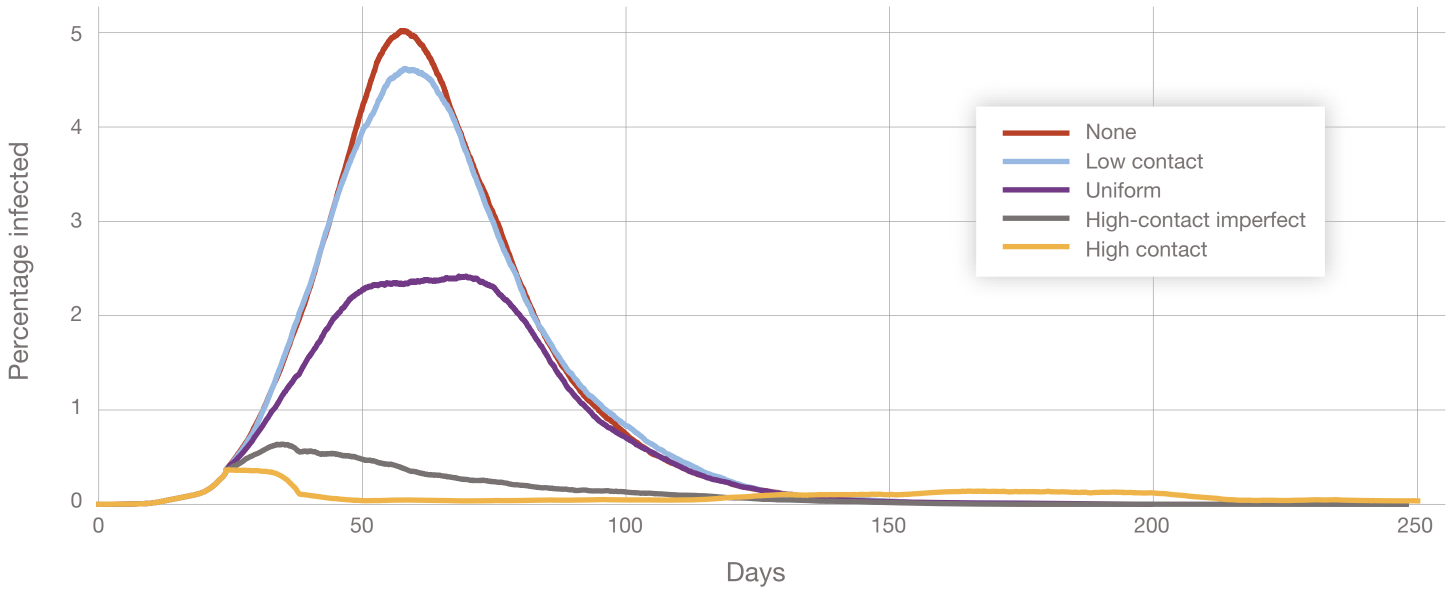

TABLE 1

Simulations Results, Averaged over 100 Runs, with 95-Percent Confidence Intervals

Vaccination Strategy	Total Infected After Intervention	Percentage Infected After Intervention	Peak Infected	Percentage Peak Infected
High contact	1,700 (± 300)	0.6 (± 0.1)	1,033 (± 6)	0.376 (± 0.002)
High-contact imperfect	7,300 (± 200)	2.66 (± 0.07)	1,530 (± 30)	0.56 (± 0.01)
Uniform	27,210 (± 80)	9.90 (±0.03)	7,300 (± 100)	2.66 (± 0.05)
Low contact	40,630 (± 70)	14.79 (± 0.03)	12,800 (± 100)	4.64 (± 0.05)
None	41,150 (± 80)	14.98 (± 0.03)	12,900 (± 100)	4.69 (± 0.05)

NOTE: Although we expect that the observed network characteristics will produce herd immunity at a rate that is significantly lower than would be expected with a population-based SIR model, this model is too simple to produce a reliable estimate of herd immunity level.

immediately halting the growth of the disease and causing it to decline even without additional vaccination or distancing. The high-contact imperfect strategy ends up cutting the number of people infected after vaccination by about 82 percent—a radically better outcome than the low-contact strategy that cuts new infections by only 1 percent.[9]

Figure 3 provides a more detailed view of how people at different levels of connectedness fare under the five models. In comparing the low-contact strategy (blue) to no vaccination (red), the only place where we saw any meaningful impact is in the low-contact category. The first blue bar (low-contact people infected in the low-contact case) is somewhat shorter than the first red bar (low-contact people infected in the no vaccination case). However, the other four pairs of red and blue bars are essentially identical. Vaccinating the low-contact people had almost no effect on the pandemic in the rest of the population.

This case can be contrasted with the high-imperfect case, where we vaccinated one-half of the top 30 percent of the population in terms of their contacts. This case corresponds to the gray bars in Figure 3. Although we only vaccinated people in the top category (more than 50 contacts), we saw dramatically better outcomes in every category. Even among the lowest-contact people, where no vaccination took place, we saw a drop of about 83 percent in infections.

FIGURE 3

People Infected After Vaccinations by Connectedness

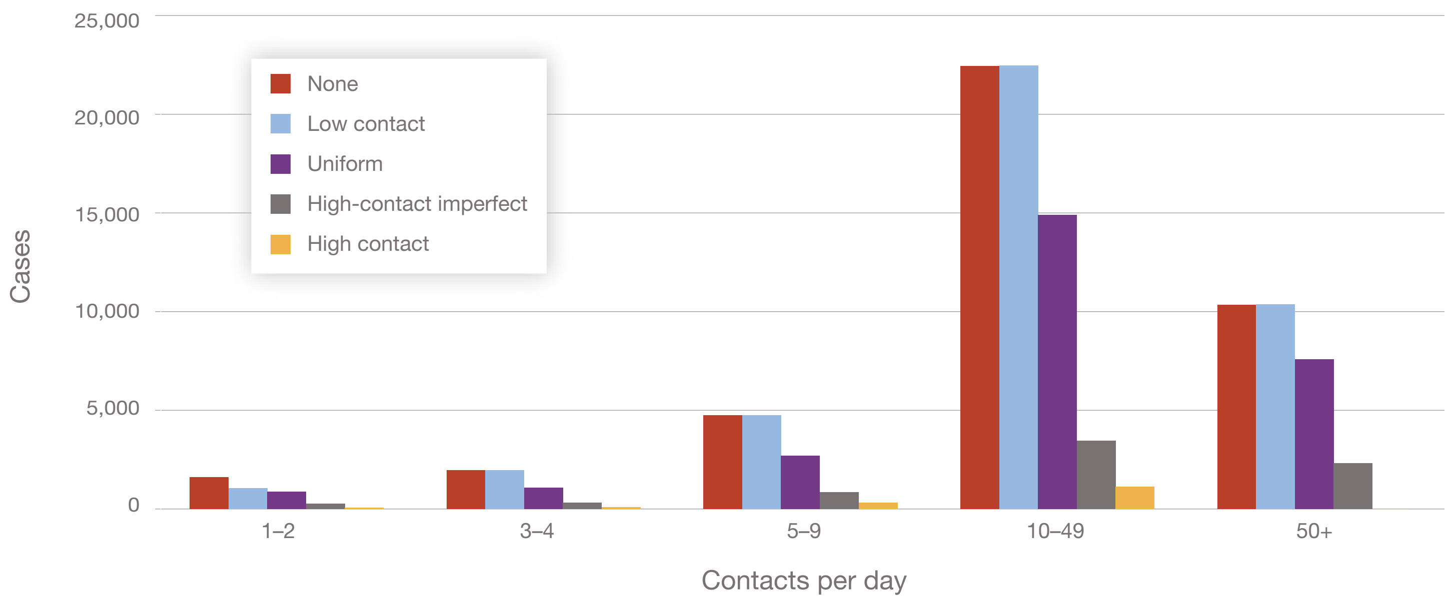

Protecting the Vulnerable by Vaccinating the Most Active

We generally assumed that the most-vulnerable (usually elderly) people have fewer contacts than the younger people who are at less risk. It is well established that older people and people with various other health conditions are at much greater risk of death if they contract COVID-19 (Yang et al., 2020). However, Figure 3 indicates that the high-imperfect strategy reduces the risk of contracting the disease dramatically across all contact categories—even those groups that are not vaccinated at all. An 83-percent reduction in cases among those who are most vulnerable can be expected to translate to a roughly 83-percent reduction in deaths among people in this group.

This is a better outcome than can be achieved by direct vaccination of the vulnerable. In theory, direct vaccination could produce a 90-percent reduction in deaths among the most-vulnerable 15 percent—but it is unrealistic to think that everyone in this group can be identified and vaccinated. Missing just a small percentage (or having some fraction decline the vaccine) would reduce effectiveness in the vulnerable group below that achieved with the high-imperfect strategy. This means that the high-imperfect strategy is at least as effective (and probably more effective) at protecting the vulnerable than direct vaccination.[10]

And, of course, the rest of the population sees similar dramatic improvements as well. Figure 3 shows that the vast majority of cases occur not among the least-connected people, but rather among those with 10 to 50 contacts per day. Although these more-active people are likely at less risk of death if they become infected, the number of infections among them means that the number of deaths could be roughly comparable. In reality, we noted similar numbers of COVID-19–related deaths among people in the age cohorts of 55–64, 65–74, 75–84, and 85 and over (National Center for Health Statistics, 2021). This is presumably because the younger cohorts are both more populous and more active, and thus more likely to be exposed. A vaccination strategy that focuses only on those 65 and older does little to address the comparable number of deaths in the 55–64 category. Preventing deaths in this younger category through direct vaccination would require more vaccine than would be available in the first phases of vaccination, but deaths in this category would be comparably reduced using a targeted strategy.

Limitations of the Model

This illustrative model lacks many of the features of the real world, and these limitations should be borne in mind when thinking about the applicability of the results. The interaction network is based on a large sample of human movement, but it is not a fully random sample. The movement profiles were built using some very broad assumptions that might have a meaningful impact on the realism of the network structure. Although we used an established epidemiological model, it is one of the simplest, and we made broad, simplifying assumptions to represent vaccination strategy. Most notably, that vaccination occurs in a single instantaneous application rather than dynamically. The availability of vaccines and the rate at which people can be inoculated could have a strong impact on the effectiveness of a given strategy. There is also the problem of willingness to accept the vaccine—if many people are resistant to vaccination, and those people are connected, parts of the network could harbor disease for an extended period. The model is applied to the interaction network of a single city, without interaction with other places that might be managing vaccination differently. The model also does not differentiate people by age or risk profile because our interaction network is based on anonymous mobile device data—this greatly limited what we can say about who the most- (and least-) active people are and therefore the actual impact of each strategy on mortality. All of these issues and many more are limitations of this model and require exploration with more-detailed models and observations.

However, the model does use an interaction network that is based on real movement and captures a much more holistic picture of human geospatial interaction than is commonly available for models of this sort. For this reason, we believe that it provides meaningful qualitative guidance for vaccination strategy. Delivering a targeted vaccination on this scale has the potential to quell the pandemic and eliminate resurgence even after lifting distancing measures.

Discussion

This result shows that the debate between protecting the vulnerable and stopping the spread of the virus might be a false choice. Vaccinating the most-vulnerable people will provide those people with 90-percent protection but will leave them (and everyone else) in the midst of a raging pandemic. That leaves many vulnerable people at risk, both those vulnerable people who have still not been vaccinated and the 10 percent of the vaccinated in whom the vaccine did not provide full protection because of efficacy rates of current vaccines. A targeted strategy, on the other hand, has the potential to halt the epidemic quickly. It would provide protection to the whole community (including the vulnerable) by more rapidly and efficiently cutting down the most-predominant transmission paths. If a targeted vaccination strategy can disrupt the spread of COVID-19 to the point where it no longer spreads out of control, the vulnerable will be largely safe, simply because they are unlikely to be exposed in the first place. People with fewer contacts can then be vaccinated in a more routine fashion and life can begin to return to normal.

An analogy can be made to mask-wearing. Similar to the way that one wears a mask not so much to protect oneself as to protect others, this analysis indicates that vaccination should be seen in the same way—people should be vaccinated not so much to protect them, but to protect others from them.

Thus, health officials should consider prioritizing not the most vulnerable but rather the most dangerous: those who are most likely to both contract and spread the disease. Health care workers are clear vaccine targets in this view. Not all critical workers meet this standard, but those who have routine and unavoidable contact with members of the public certainly do. This might include such groups

> The debate between protecting the vulnerable and stopping the spread of the virus might be a false choice.

as grocery store workers and rideshare drivers. This might also include people who do critical jobs in close contact not with members of the public but with other workers. This category might include workers in meatpacking plants and other industrial settings, where full distancing is not practical. It might also include people who live and work in close group quarters, including nursing homes and prisons, with emphasis on the employees who come and go from these facilities. More controversially, it could be beneficial to vaccinate low-compliance groups, such as college students and other young adults. Although they are not at particularly high risk from the disease, they can pose a huge risk to others if they are not willing or able to follow distancing guidelines. Although this model does not provide guidance about how to identify the most-active people, it makes it clear that those who fall in this category should be a major focus of the vaccination effort, because the gains from removing such people from the transmission chain are immense.

One way to identify high-contact people would be to combine vaccination with contact tracing by asking people who test positive to list those with whom they came into contact during the past one to three days. These contacts would then be tested with a rapid antibody test to see if they have already had the disease—a procedure that can save precious vaccine doses for those who do not yet have antibodies (Bubar et al., 2020). Those who do not have antibodies would then be vaccinated if they agree to be vaccinated. This procedure would be extremely efficient at vaccinating high-contact people because people with many contacts would be much more likely to be named than those with few contacts. The practicality of taking such an approach on a national scale would need to be determined, but, in theory, it could provide a useful supplement to vaccinating those who are known to have high numbers of contacts because of their work or living arrangements.

It should also be noted that each time the virus moves between hosts, there is a small chance of selecting for an even more harmful mutated form of the virus. This provides yet another reason to protect the vulnerable by stopping the spread. Loosely speaking, 83-percent fewer transmission events means an 83-percent lower probability of producing a faster spreading, more lethal, or vaccine-resistant strain.

The network view involves shifting perspective one step up the causal chain, from trying to protect people from the effects of the disease to attacking the spread of the disease and bringing it under sufficient control that the vast majority of people will be safe until enough vaccine can be produced to protect everyone. This simple exploration supports the idea that the United States might be able to bring the disease under control in just a few months if the first several million vaccine doses are used well. A targeted strategy could shift forward by months the timeline for reopening schools and businesses and rebuilding the social ties that are such an important part of life.

Appendix A. Mobile Device–Based Contact Network

We used volunteered mobile device data provided by UberMedia, which aggregates such data for commercial purposes. We used the data to construct an approximate contact network for the city of Portland, Oregon, both before physical distancing (late February and early March 2020) and after physical distancing began (late March and early April 2020). We began with a collection of 2.2 billion pings from mobile devices (mostly smartphones) that were using one of more than 150,000 apps that have asked users to allow sharing of location data with the app developer and its partners. Each ping is associated with an anonymized device identifier and includes both the time and the geographic location of the device with a nominal precision of three meters. Reporting from these devices is quite uneven, so we developed a procedure to composite data from across about a month to build up a general movement pattern for about 250,000 of the 1.7 million individuals in Portland. We restricted the analysis to devices that had a common evening location within the Portland metropolitan area and appeared in the data on at least ten days in both the predistancing and postdistancing periods.

The details of the process by which we built the contact network are complex and beyond the scope of this paper, but in general terms, the process involved two steps: producing a set of stops for each device, then determining when devices were stopped near one another during overlapping times to produce contacts. We used four types of stops. Common evening locations (CELs) were estimated by UberMedia by identifying the location where each device was most often found between 7:00 p.m. and 7:00 a.m. Not every device had such a location, and we excluded those that either did not have a CEL or where the CEL was outside of the Portland metropolitan area. Common daytime locations (CDLs), also estimated by UberMedia, reflected a similar location where a device was often found between 8:00 a.m. and 5:00 p.m. Devices that did not have a CDL were retained in the data set because many people do not have the kind of jobs that keep them in one location all day. We then identified devices that entered a radius of 10 meters from the center of a set of about 9,000 known business locations in Portland. If the device pinged within that business, we assumed that it was there for at least five minutes, although if it pinged repeatedly, we measured the length of the visit. Finally, we identified places where a device moved by less than 10 meters during repeated pings, tracking locations and durations where devices stayed in one place. This had the effect of removing devices that were in motion (e.g., car navigation apps) and capturing activity taking place in a wide variety of locations (e.g., homes, parks).

With these stops determined, we made some broad assumptions in order to compile them into a composite movement pattern that compensated for the fact that the volunteered mobile device data represented only a sample of actual activity. First, we assumed that each person spends 12 hours at home each day—from 8:00 p.m. to 8:00 a.m. Second, we assumed that each person who has a common daytime location spends eight hours at that location each day—from 9:00 a.m. to 5:00 p.m. Finally, we took about two weeks' worth of the other stops and collapsed them down to a single day while retaining the time of day for each stop. We calibrated the number of days to use in a way that produced an average number of contacts per person that corresponded to activity estimates developed by the Network Dynamics and Simulations Science Laboratory at Virginia Tech (Network Dynamics and Simulation and Science Laboratory, undated). Finally, we divided the Portland metropolitan area into 10-meter squares and looked for instances in which two devices were in the same square during overlapping times. These overlaps were aggregated upward to produce a list of potential contacts between people—instances where people stopped within 10 meters of one another for some amount of time. Once these contacts were established, all positional information was stripped from the contact list to prevent deanonymization.

This contact network is far from perfect. Many people are home for more or less than 12 hours each day. Not everyone works eight hours a day, and not all of that work happens between 9:00 a.m. and 5:00 p.m. People do not use their mobile devices at statistically random times—so the times when they appear in the data cannot be assumed to be a random sample of their activity. We have no way of knowing what people were doing when they stopped within 10 meters of one another, whether they were wearing masks, etc. However, the network does represent an empirically derived view of the whole movement pattern of a city. Even with its flaws, it provides an unprecedented look at the overall pattern of human contact and gives an objectively based sense of the shape of the human physical contact network.

Appendix B. Network Science and Herd Immunity

Network scientists have known for many years that the relative efficacy of vaccinating high-contact nodes, as opposed to vaccinating nodes at random, depends very strongly on the shape of the network. Because the real network of person-to-person contacts has been historically very difficult to obtain, these scientists have often turned to the study of simple, stylized networks. For example, in a seminal paper, Pastor-Satorras and Vespignani (2002) showed that, if the network takes the form of a scale-free network (a type of stylized network with a small number of very highly connected nodes and a much larger number of less-connected nodes), then a targeted vaccination strategy can stop the pandemic once 16 percent of the population has been immunized.[11] Remarkably, if, instead, nodes are vaccinated uniformly at random, then the pandemic will not be halted unless a large majority of the population is immunized. As the population (N) tends to infinity, the proportion of the population needed to be immunized tends to 100 percent. This is a strong result—not only is the targeted vaccination strategy more effective than a uniform strategy, the uniform strategy can be said to completely fail. In the uniform case, vaccination only safeguards the individual, and there is no communal benefit to vaccinating any given node. Interestingly, when Pastor-Satorras and Vespignani considered the same problem but for a different stylized network known as a Watts-Strogatz network, the results were completely different. The two vaccination strategies (targeted and uniform) had similar efficacies, and both caused the spread to halt once 39 percent of the population had been inoculated (Table B.1).

The two strategies contrast with the random (Erdős–Rényi) network assumption that underlies the basic SIR models that form the foundation of many population-based models. This network assumes that all nodes have about the same number of links (specifically, that the node degrees are normally distributed) and that the links have no particular structure, with each node linking to other nodes at random. In such a network, herd immunity is expected to begin when about 1 – 1 / R0 nodes are vaccinated, where *R0* is the number of additional cases produced by the average case when no control measures are being taken. This number has been estimated at about three in the case of COVID-19, so herd immunity would be expected at around 67-percent vaccination.

These results imply that, in order to assess the efficacy of a vaccination strategy, it is crucial to understand the network in which the disease is spreading. The same strategy could lead to wildly different results, depending on the network. The key takeaway then is that the network is extremely important for the dynamics of the contagion, both with and without any vaccination. In dealing with the COVID-19 pandemic, it then becomes extremely important to understand the structure of the real-world network of actual person-to-person connections. It is not sufficient to approximate the complexities of the real world with stylized networks.

TABLE B.1

Required Vaccination Percentage to Halt the Spread of the Disease

Vaccination Program	Scale-Free	Watts-Strogatz
Uniform vaccination	100 percent (as N → ∞)	39 percent
Targeted vaccination	16 percent	39 percent

SOURCE: Pastor-Satorras and Vespignani, 2002.

Notes

[1] See Figure 4 in Buckner, Chowell, and Springborn, 2020.

[2] See, for example, Alain Barrat, Marc Barthélemy, and Alessandro Vespignani, *Dynamical Processes on Complex Networks*, New York: Cambridge University Press, 2008.

[3] See Matrajt et al., 2020, for COVID-19, Medlock and Galvani, 2009, for influenza.

[4] By multiplying the probabilities this way, we made an approximation that is valid when the contacts are short in duration and the transmission rate is not too large.

[5] For the sake of simplicity, we used simple degree (contact) count in these examples. In reality, of course, the nature of the contact also matters a great deal. Longer contacts carry more risk than shorter ones, unmasked indoor contacts carry more risk than masked outdoor contacts, etc. In our model, we weighted these contacts by duration but not by mask wearing.

[6] See Appendix A for more information about the data and the derived contact network.

[7] Even with the elaboration of an explicit contact network, the SIR framework remains a major simplification of the dynamics of the disease. For that reason, readers should not take the numerical values or timescales displayed by the model literally. Although simple variations of disease duration, transmissibility, incubation period, and the like would change the timescales of the model, a simulation with sufficient realism to explore policy options in more detail would require a more complex modeling approach.

[8] The efficacy of Pfizer's vaccine was reported to be 90 percent (Thomas, Gelles, and Zimmer, 2020). Uniformly vaccinating 15 percent of the population is equivalent to vaccinating 16.7 percent of the population with a vaccine that is 90-percent effective.

[9] Because of the highly simplified nature of this model, we do not address hospitalization rates and capacities here. Properly addressing these factors would require the construction of a more elaborate model—one that is also harder to interpret and understand.

[10] For this simple calculation, we assume that a 90-percent effective vaccine produces one of two outcomes: immunity (90 percent) or no protection (10 percent). The reality of vaccine efficacy is more nuanced, with a vaccine sometimes providing partial protection. This might allow vulnerable people who become ill to have a better chance of recovery than they would if they had not been vaccinated. In this case, the math becomes more complicated, and the preservation of lives outside of the most-vulnerable group would likely need to be considered in evaluating each strategy. Because this model does not directly track ages or risk categories, it is not able to make these comparisons, but this analysis generally indicates that the strategy of vaccinating the active would remain superior in terms of lives saved across all age and risk categories.

[11] Note that this is an idealized case where every single one of the 16 percent of nodes with the most connections are vaccinated, and the vaccine is 100-percent effective. Although real-world networks might bear some resemblance to a scale-free network, and real vaccines can be upward of 90-percent effective, this case should be taken as illustrative of a general principle rather than a realistic outcome.

References

Barrat, Alain, Marc Barthélemy, and Alessandro Vespignani, *Dynamical Processes on Complex Networks*, New York: Cambridge University Press, 2008.

Bubar, Kate M., Kyle Reinholt, Stephen M. Kissler, Marc Lipsitch, Sarah Cobey, Yonatan H. Grad, and Daniel B. Larremore, "Model-Informed COVID-19 Vaccine Prioritization Strategies by Age and Serostatus," medRxiv, December 7, 2020.

Buckner, Jack H., Geraldo Chowell, and Michael R. Springborn, "Dynamic Prioritization of COVID-19 Vaccines When Social Distancing is Limited for Essential Workers," medRxiv, updated December 17, 2020.

Cox, Christopher, "The Vulnerable Can Wait. Vaccinate the Super-Spreaders First," *Wired*, November 10, 2020. As of January 21, 2021: https://www.wired.com/story/covid-19-vaccine-super-spreaders/

Matrajt, Laura, Julia Eaton, Tiffany Leung, and Elizabeth R. Brown, "Vaccine Optimization for COVID-19: Who to Vaccinate First?" medRxiv, updated December 15, 2020.

Medlock, Jan, and Alison P. Galvani, "Optimizing Influenza Vaccine Distribution," *Science*, Vol. 325, No. 5948, 2009, pp. 1705–1708.

National Center for Health Statistics, "Weekly Updates by Select Demographic and Geographic Characteristics: Provisional Death Counts for Coronavirus Disease 2019 (COVID-19)," webpage, Centers for Disease Control and Prevention, updated January 21, 2021.

Network Dynamics and Simulation and Science Laboratory, "Synthetic Data Products for Societal Infrastructures and Proto-Populations: Data Set 2.0," Blacksburg, Va.: Virginia Polytechnic Institute and State University, NDSSL-TR-07-003, undated.

Pastor-Satorras, Romualdo, and Alessandro Vespignani, "Immunization of Complex Networks," *Physical Review E*, Vol. 65, No. 3, 2002.

Thomas, Katie, David Gelles, and Carl Zimmer, "Pfizer's Early Data Shows Vaccine Is More than 90% Effective," *New York Times*, updated December 15, 2020.

Tingley, Kim, "Who Should Get a Covid-19 Vaccine First?" *New York Times Magazine*, November 5, 2020.

Yang, Wan, Sasikiran Kandula, Mary Huynh, Sharon K. Greene, Gretchen Van Wye, Wenhui Li, Hui Tai Chan, Emily McGibbon, Alice Yeung, Don Olson, Anne Fine, and Jeffrey Shaman, "Estimating the Infection-Fatality Risk of SARS-CoV-2 in New York City During the Spring 2020 Pandemic Wave: A Model-Based Analysis," *Lancet Infectious Diseases*, October 19, 2020.

About This Perspective

Recent models of coronavirus disease 2019 vaccination have tended not to take into account the person-to-person contact structure that results in the disease's spread. In this Perspective, we describe research in which a model of a contact network—derived from a data set of 2.2 billion mobile device location points—was used to run five different vaccine models, each of which varied by the number of contacts of those vaccinated. We anticipate that this might be useful for anyone determining the best vaccination model for a certain population.

RAND Social and Economic Well-Being is a division of the RAND Corporation that seeks to actively improve the health and social and economic well-being of populations and communities throughout the world. This research was conducted in the Community Health and Environmental Policy Program within RAND Social and Economic Well-Being. The program focuses on such topics as infrastructure, science and technology, community design, community health promotion, migration and population dynamics, transportation, energy, and climate and the environment, as well as other policy concerns that are influenced by the natural and built environment, technology, and community organizations and institutions that affect well-being. For more information, email chep@rand.org.

Funding

Funding for this Perspective was provided by gifts from RAND supporters and income from operations.

About the Authors

Timothy R. Gulden is a senior policy researcher at the RAND Corporation. His research focuses on modeling complex systems in the context of data, as well as more general policy analysis including cost/benefit analysis, organizational design, and cybersecurity policy. He has a Ph.D. in public policy.

Gavin S. Hartnett is an information scientist at the RAND Corporation. A theoretical physicist by training, Hartnett conducts policy-oriented research into a wide variety of topics, including artificial intelligence, machine learning, blockchains, and quantum computing. He has a Ph.D. in physics.

Raffaele Vardavas is a senior mathematician at the RAND Corporation. His expertise is primarily in constructing and analyzing epidemic models for the spreading of infectious disease using both deterministic and stochastic methods and using both population-level and agent-based models. He has a Ph.D. in physics.

David Kravitz is an adjunct data scientist at the RAND Corporation. He has worked on a large variety of projects at RAND, including with the City of Santa Monica on an initiative to help those experiencing homelessness; creating missile defense software to simulate attacks from foreign adversaries; and on numerous projects involving network and text analysis of social media. He has a Ph.D. in algorithms, combinatorics, and optimization.

For more information on this publication, visit www.rand.org/t/PEA1068-1.

PE-A1068-1