HEALTH

This Overlooked Variable Is the Key to the Pandemic

It's not R.

Updated at 1:17 p.m. ET on October 1, 2020

There's something strange about this coronavirus pandemic. Even after months of extensive research by the global scientific community, many questions remain open.

Why, for instance, was there such an enormous death toll in northern Italy, but not the rest of the country? Just three contiguous regions in northern Italy have 25,000 of the country's nearly 36,000 total deaths; just one region, Lombardy, has about 17,000 deaths. Almost all of these were concentrated in the first few months of the outbreak. What happened in Guayaquil, Ecuador, in April, when so many died so

quickly that bodies were abandoned in the sidewalks and streets?* Why, in the spring of 2020, did so few cities account for a substantial portion of global deaths, while many others with similar density, weather, age distribution, and travel patterns were spared? What can we really learn from Sweden, hailed as a great success by some because of its low case counts and deaths as the rest of Europe experiences a second wave, and as a big failure by others because it did not lock down and suffered excessive death rates earlier in the pandemic? Why did widespread predictions of catastrophe in Japan not bear out? The baffling examples go on.

I've heard many explanations for these widely differing trajectories over the past nine months—weather, elderly populations, vitamin D, prior immunity, herd immunity—but none of them explains the timing or the scale of these drastic variations. But there *is* a potential, overlooked way of understanding this pandemic that would help answer these questions, reshuffle many of the

current heated arguments, and, crucially, help us get the spread of COVID-19 under control.

By now many people have heard about R0—the basic reproductive number of a pathogen, a measure of its contagiousness on average. But unless you've been reading scientific journals, you're less likely to have encountered *k*, the measure of its dispersion. The definition of *k* is a mouthful, but it's simply a way of asking whether a virus spreads in a steady manner or in big bursts, whereby one person infects many, all at once. After nine months of collecting epidemiological data, we know that this is an *overdispersed* pathogen, meaning that it tends to spread in clusters, but this knowledge has not yet fully entered our way of thinking about the pandemic—or our preventive practices.

The now-famed R0 (pronounced as "r-naught") is an *average* measure of a pathogen's contagiousness, or the mean number of susceptible people expected to become infected after being exposed to a person with the disease. If one ill person infects three others

on average, the R0 is three. This parameter has been widely touted as a key factor in understanding how the pandemic operates. News media have produced multiple explainers and visualizations for it. Movies praised for their scientific accuracy on pandemics are lauded for having characters explain the "all-important" R0. Dashboards track its real-time evolution, often referred to as R or Rt, in response to our interventions. (If people are masking and isolating or immunity is rising, a disease can't spread the same way anymore, hence the difference between R0 and R.)

Unfortunately, averages aren't always useful for understanding the distribution of a phenomenon, especially if it has widely varying behavior. If Amazon's CEO, Jeff Bezos, walks into a bar with 100 regular people in it, the average wealth in that bar suddenly exceeds $1 billion. If I also walk into that bar, not much will change. Clearly, the average is not that useful a number to understand the distribution of wealth in that bar, or how to change it. Sometimes, the mean is not the message.

Meanwhile, if the bar has a person infected with COVID-19, and if it is also poorly ventilated and loud, causing people to speak loudly at close range, almost everyone in the room could potentially be infected—a pattern that's been observed many times since the pandemic begin, and that is similarly not captured by R. That's where the dispersion comes in.

There are COVID-19 incidents in which a single person likely infected 80 percent or more of the people in the room in just a few hours. But, at other times, COVID-19 can be surprisingly much less contagious. Overdispersion and super-spreading of this virus are found in research across the globe. A growing number of studies estimate that a majority of infected people may not infect a single other person. A recent paper found that in Hong Kong, which had extensive testing and contact tracing, about 19 percent of cases were responsible for 80 percent of transmission, while 69 percent of cases did not infect another person. This finding is not rare: Multiple studies from the beginning have

suggested that as few as 10 to 20 percent of infected people may be responsible for as much as 80 to 90 percent of transmission, and that many people barely transmit it.

This highly skewed, imbalanced distribution means that an early run of bad luck with a few super-spreading events, or clusters, can produce dramatically different outcomes even for otherwise similar countries. Scientists looked globally at known early-introduction events, in which an infected person comes into a country, and found that in some places, such imported cases led to no deaths or known infections, while in others, they sparked sizable outbreaks. Using genomic analysis, researchers in New Zealand looked at more than half the confirmed cases in the country and found a staggering 277 *separate* introductions in the early months, but also that only 19 percent of introductions led to more than one additional case. A recent review shows that this may even be true in congregate living spaces, such as nursing homes, and that multiple introductions may be necessary before

an outbreak takes off. Meanwhile, in Daegu, South Korea, just one woman, dubbed Patient 31, generated more than 5,000 known cases in a megachurch cluster.

Unsurprisingly, SARS-CoV, the previous incarnation of SARS-CoV-2 that caused the 2003 SARS outbreak, was also overdispersed in this way: The majority of infected people did not transmit it, but a few super-spreading events caused most of the outbreaks. MERS, another coronavirus cousin of SARS, also appears overdispersed, but luckily, it does not—yet—transmit well among humans.

This kind of behavior, alternating between being super infectious and fairly noninfectious, is exactly what *k* captures, and what focusing solely on R hides. Samuel Scarpino, an assistant professor of epidemiology and complex systems at Northeastern, told me that this has been a huge challenge, especially for health authorities in Western societies, where the pandemic playbook was geared toward the flu—and not without reason, because pandemic

flu *is* a genuine threat. However, influenza does not have the same level of clustering behavior.

We can think of disease patterns as leaning deterministic or stochastic: In the former, an outbreak's distribution is more linear and predictable; in the latter, randomness plays a much larger role and predictions are hard, if not impossible, to make. In deterministic trajectories, we expect what happened yesterday to give us a good sense of what to expect tomorrow. Stochastic phenomena, however, don't operate like that—the same inputs don't always produce the same outputs, and things can tip over quickly from one state to the other. As Scarpino told me, "Diseases like the flu are pretty nearly deterministic and R0 (while flawed) paints about the right picture (nearly impossible to stop until there's a vaccine)." That's not necessarily the case with super-spreading diseases.

Nature and society are replete with such imbalanced phenomena, some of which are said to work according to the Pareto principle, named after the

sociologist Vilfredo Pareto. Pareto's insight is sometimes called the 80/20 principle—80 percent of outcomes of interest are caused by 20 percent of inputs—though the numbers don't have to be that strict. Rather, the Pareto principle means that a small number of events or people are responsible for the majority of consequences. This will come as no surprise to anyone who has worked in the service sector, for example, where a small group of problem customers can create almost all the extra work. In cases like those, booting just those customers from the business or giving them a hefty discount may solve the problem, but if the complaints are evenly distributed, different strategies will be necessary. Similarly, focusing on the R alone, or using a flu-pandemic playbook, won't necessarily work well for an overdispersed pandemic.

Hitoshi Oshitani, a member of the National COVID-19 Cluster Taskforce at Japan's Ministry of Health, Labour and Welfare and a professor at Tohoku University who told me that Japan focused on the overdispersion impact from early on, likens

his country's approach to looking at a forest and trying to find the clusters, not the trees. Meanwhile, he believes, the Western world was getting distracted by the trees, and got lost among them. To fight a super-spreading disease effectively, policy makers need to figure out why super-spreading happens, and they need to understand how it affects everything, including our contact-tracing methods and our testing regimes.

There may be many different reasons a pathogen super-spreads. Yellow fever spreads mainly via the mosquito *Aedes aegypti*, but until the insect's role was discovered, its transmission pattern bedeviled many scientists. Tuberculosis was thought to be spread by close-range droplets until an ingenious set of experiments proved that it was airborne. Much is still unknown about the super-spreading of SARS-CoV-2. It might be that some people are super-emitters of the virus, in that they spread it a lot more than other people. Like other diseases, contact patterns surely play a part: A politician on the campaign trail or a student in a college dorm is very

different in how many people they could potentially expose compared with, say, an elderly person living in a small household. However, looking at nine months of epidemiological data, we have important clues to some of the factors.

In study after study, we see that super-spreading clusters of COVID-19 almost overwhelmingly occur in poorly ventilated, indoor environments where many people congregate over time—weddings, churches, choirs, gyms, funerals, restaurants, and such—especially when there is loud talking or singing without masks. For super-spreading events to occur, multiple things have to be happening at the same time, and the risk is not equal in every setting and activity, Muge Cevik, a clinical lecturer in infectious diseases and medical virology at the University of St. Andrews and a co-author of a recent extensive review of transmission conditions for COVID-19, told me.

Cevik identifies "prolonged contact, poor ventilation, [a] highly infectious person, [and]

crowding" as the key elements for a super-spreader event. Super-spreading can also occur indoors beyond the six-feet guideline, because SARS-CoV-2, the pathogen causing COVID-19, can travel through the air and accumulate, especially if ventilation is poor. Given that some people infect others before they show symptoms, or when they have very mild or even no symptoms, it's not always possible to know if we are highly infectious ourselves. We don't even know if there are more factors yet to be discovered that influence super-spreading. But we don't need to know all the *sufficient* factors that go into a super-spreading event to avoid what seems to be a *necessary* condition most of the time: many people, especially in a poorly ventilated indoor setting, and especially not wearing masks. As Natalie Dean, a biostatistician at the University of Florida, told me, given the huge numbers associated with these clusters, targeting them would be very effective in getting our transmission numbers down.

Overdispersion should also inform our contact-tracing efforts. In fact, we may need to turn them

upside down. Right now, many states and nations engage in what is called forward or prospective contact tracing. Once an infected person is identified, we try to find out with whom they interacted afterward so that we can warn, test, isolate, and quarantine these potential exposures. But that's not the only way to trace contacts. And, because of overdispersion, it's not necessarily where the most bang for the buck lies. Instead, in many cases, we should try to work *backwards* to see who first infected the subject.

Because of overdispersion, most people will have been infected by someone who also infected other people, because only a small percentage of people infect many at a time, whereas most infect zero or maybe one person. As Adam Kucharski, an epidemiologist and the author of the book *The Rules of Contagion*, explained to me, if we can use retrospective contact tracing to find the person who infected our patient, and *then* trace the forward contacts of the infecting person, we are generally going to find a lot more cases compared with

forward-tracing contacts of the infected patient, which will merely identify *potential* exposures, many of which will not happen anyway, because most transmission chains die out on their own.

The reason for backward tracing's importance is similar to what the sociologist Scott L. Feld called the friendship paradox: Your friends are, on average, going to have more friends than you. (Sorry!) It's straightforward once you take the network-level view. Friendships are not distributed equally; some people have a lot of friends, and your friend circle is more likely to include those social butterflies, because how could it not? They friended you and others. And those social butterflies will drive up the average number of friends that your friends have compared with you, a regular person. (Of course, this will not hold for the social butterflies themselves, but overdispersion means that there are much fewer of them.) Similarly, the infectious person who is transmitting the disease is like the pandemic social butterfly: The average number of people they infect will be much higher than most of

the population, who will transmit the disease much less frequently. Indeed, as Kucharski and his co-authors show mathematically, overdispersion means that "forward tracing alone can, on average, identify at most the mean number of secondary infections (i.e. R)"; in contrast, "backward tracing increases this maximum number of traceable individuals by a factor of 2-3, as index cases are more likely to come from clusters than a case is to generate a cluster."

Even in an overdispersed pandemic, it's not pointless to do forward tracing to be able to warn and test people, *if* there are extra resources and testing capacity. But it doesn't make sense to do forward tracing while not devoting enough resources to backward tracing and finding clusters, which cause so much damage.

Another significant consequence of overdispersion is that it highlights the importance of certain kinds of rapid, cheap tests. Consider the current dominant model of test and trace. In many places, health authorities try to trace and find forward contacts of

an infected person: everyone they were in touch with since getting infected. They then try to test all of them with expensive, slow, but highly accurate PCR (polymerase chain reaction) tests. But that's not necessarily the best way when clusters are so important in spreading the disease.

PCR tests identify RNA segments of the coronavirus in samples from nasal swabs—like looking for its signature. Such diagnostic tests are measured on two different dimensions: Are they good at identifying people who are not infected (specificity), and are they good at identifying people who are infected (sensitivity)? PCR tests are highly accurate for both dimensions. However, PCR tests are also slow and expensive, and they require a long, uncomfortable swab up the nose at a medical facility. The slow processing times means that people don't get timely information when they need it. Worse, PCR tests are so responsive that they can find tiny remnants of coronavirus signatures long after someone has stopped being contagious, which can cause unnecessary quarantines.

Meanwhile, researchers have shown that rapid tests that are very accurate for identifying people who do *not* have the disease, but not as good at identifying infected individuals, can help us contain this pandemic. As Dylan Morris, a doctoral candidate in ecology and evolutionary biology at Princeton, told me, cheap, low-sensitivity tests can help mitigate a pandemic even if it is not overdispersed, but they are particularly valuable for cluster identification during an overdispersed one. This is especially helpful because some of these tests can be administered via saliva and other less-invasive methods, and be distributed outside medical facilities.

In an overdispersed regime, identifying *transmission events* (someone infected someone else) is more important than identifying *infected individuals*. Consider an infected person and their 20 forward contacts—people they met since they got infected. Let's say we test 10 of them with a cheap, rapid test and get our results back in an hour or two. This isn't a great way to determine exactly who is sick out of that 10, because our test will miss some positives,

but that's fine for our purposes. If everyone is negative, we can act as if nobody is infected, because the test is pretty good at finding negatives. However, the moment we find a few transmissions, we know we may have a super-spreader event, and we can tell all 20 people to assume they are positive and to self-isolate—if there are one or two transmissions, there are likely more, exactly because of the clustering behavior. Depending on age and other factors, we can test those people individually using PCR tests, which can pinpoint who is infected, or ask them all to wait it out.

Scarpino told me that overdispersion also enhances the utility of other aggregate methods, such as wastewater testing, especially in congregate settings like dorms or nursing homes, allowing us to detect clusters without testing everyone. Wastewater testing also has low sensitivity; it may miss positives if too few people are infected, but that's fine for population-screening purposes. If the wastewater testing is signaling that there are *likely* no infections, we do not need to test everyone to find every last

potential case. However, the moment we see signs of a cluster, we can rapidly isolate everyone, again while awaiting further individualized testing via PCR tests, depending on the situation.

Unfortunately, until recently, many such cheap tests had been held up by regulatory agencies in the United States, partly because they were concerned with their relative lack of accuracy in identifying positive cases compared with PCR tests—a worry that missed their population-level usefulness for this particular overdispersed pathogen.

To return to the mysteries of this pandemic, what *did* happen early on to cause such drastically different trajectories in otherwise similar places? Why haven't our usual analytic tools—case studies, multi-country comparisons—given us better answers? It's not intellectually satisfying, but because of the overdispersion and its stochasticity, there may not be an explanation beyond that the worst-hit regions, at least initially, simply had a few unlucky early super-spreading events. It wasn't just pure luck:

Dense populations, older citizens, and congregate living, for example, made cities around the world more susceptible to outbreaks compared with rural, less dense places and those with younger populations, less mass transit, or healthier citizenry. But why Daegu in February and not Seoul, despite the two cities being in the same country, under the same government, people, weather, and more? As frustrating at it may be, sometimes, the answer is merely where Patient 31 and the megachurch she attended happened to be.

Overdispersion makes it harder for us to absorb lessons from the world, because it interferes with how we ordinarily think about cause and effect. For example, it means that events that result in spreading and non-spreading of the virus are asymmetric in their ability to inform us. Take the highly publicized case in Springfield, Missouri, in which two infected hairstylists, both of whom wore masks, continued to work with clients while symptomatic. It turns out that no apparent infections were found among the 139 exposed clients (67 were directly tested; the rest

did not report getting sick). While there is a lot of evidence that masks are crucial in dampening transmission, that event *alone* wouldn't tell us if masks work. In contrast, studying transmission, the rarer event, can be quite informative. Had those two hairstylists transmitted the virus to large numbers of people despite everyone wearing masks, it would be important evidence that, perhaps, masks aren't useful in preventing super-spreading.

Comparisons, too, give us less information compared with phenomena for which input and output are more tightly coupled. When that's the case, we can check for the presence of a factor (say, sunshine or Vitamin D) and see if it correlates with a consequence (infection rate). But that's much harder when the consequence can vary widely depending on a few strokes of luck, the way that the wrong person was in the wrong place sometime in mid-February in South Korea. That's one reason multi-country comparisons have struggled to identify dynamics that sufficiently explain the trajectories of different places.

Once we recognize super-spreading as a key lever, countries that look as if they were too relaxed in some aspects appear very different, and our usual polarized debates about the pandemic are scrambled, too. Take Sweden, an alleged example of the great success or the terrible failure of herd immunity without lockdowns, depending on whom you ask. In reality, although Sweden joins many other countries in failing to protect elderly populations in congregate-living facilities, its measures that target super-spreading have been stricter than many other European countries. Although it did not have a complete lockdown, as Kucharski pointed out to me, Sweden imposed a 50-person limit on indoor gatherings in March, and did not remove the cap even as many other European countries eased such restrictions after beating back the first wave. (Many are once again restricting gathering sizes after seeing a resurgence.) Plus, the country has a small household size and fewer multigenerational households compared with most of Europe, which further limits transmission and cluster possibilities. It kept schools fully open without distancing or masks,

but only for children under 16, who are unlikely to be super-spreaders of this disease. Both transmission and illness risks go up with age, and Sweden went all online for higher-risk high-school and university students—the opposite of what we did in the United States. It also encouraged social-distancing, and closed down indoor places that failed to observe the rules. From an overdispersion and super-spreading point of view, Sweden would not necessarily be classified as among the most lax countries, but nor is it the most strict. It simply doesn't deserve this oversize place in our debates assessing different strategies.

Although overdispersion makes some usual methods of studying causal connections harder, we can study failures to understand which conditions turn bad luck into catastrophes. We can also study sustained success, because bad luck will eventually hit everyone, and the response matters.

The most informative case studies may well be those who had terrible luck initially, like South Korea, and

yet managed to bring about significant suppression. In contrast, Europe was widely praised for its opening early on, but that was premature; many countries there are now experiencing widespread rises in cases and look similar to the United States in some measures. In fact, Europe's achieving a measure of success this summer and relaxing, including opening up indoor events with larger numbers, is instructive in another important aspect of managing an overdispersed pathogen: Compared with a steadier regime, success in a stochastic scenario can be more fragile than it looks.

Once a country has too many outbreaks, it's almost as if the pandemic switches into "flu mode," as Scarpino put it, meaning high, sustained levels of community spread even though a majority of infected people may not be transmitting onward. Scarpino explained that barring truly drastic measures, once in that widespread and elevated mode, COVID-19 can keep spreading because of the sheer number of chains already out there. Plus,

the overwhelming numbers may eventually spark more clusters, further worsening the situation.

As Kucharski put it, a relatively quiet period can hide how quickly things can tip over into large outbreaks and how a few chained amplification events can rapidly turn a seemingly under-control situation into a disaster. We're often told that if Rt, the real-time measure of the average spread, is above one, the pandemic is growing, and that below one, it's dying out. That may be true for an epidemic that is not overdispersed, and while an Rt below one is certainly good, it's misleading to take too much comfort from a low Rt when just a few events can reignite massive numbers. No country should forget South Korea's Patient 31.

That said, overdispersion is also a cause for hope, as South Korea's aggressive and successful response to that outbreak—with a massive testing, tracing, and isolating regime—shows. Since then, South Korea has also been practicing sustained vigilance, and has demonstrated the importance of backward tracing.

When a series of clusters linked to nightclubs broke out in Seoul recently, health authorities aggressively traced and tested tens of thousands of people *linked to the venues*, regardless of their interactions with the index case, six feet apart or not—a sensible response, given that we know the pathogen is airborne.

Perhaps one of the most interesting cases has been Japan, a country with middling luck that got hit early on and followed what appeared to be an unconventional model, not deploying mass testing and never fully shutting down. By the end of March, influential economists were publishing reports with dire warnings, predicting overloads in the hospital system and huge spikes in deaths. The predicted catastrophe never came to be, however, and although the country faced some future waves, there was never a large spike in deaths despite its aging population, uninterrupted use of mass transportation, dense cities, and lack of a formal lockdown.

It's not that Japan was better situated than the United States in the beginning. Similar to the U.S. and Europe, Oshitani told me, Japan did not initially have the PCR capacity to do widespread testing. Nor could it impose a full lockdown or strict stay-at-home orders; even if that had been desirable, it would not have been legally possible in Japan.

Oshitani told me that in Japan, they had noticed the overdispersion characteristics of COVID-19 as early as February, and thus created a strategy focusing mostly on cluster-busting, which tries to prevent one cluster from igniting another. Oshitani said he believes that "the chain of transmission cannot be sustained without a chain of clusters or a megacluster." Japan thus carried out a cluster-busting approach, including undertaking aggressive backward tracing to uncover clusters. Japan also focused on ventilation, counseling its population to avoid places where the three C's come together—crowds in closed spaces in close contact, especially if there's talking or singing—bringing together the science of overdispersion with the recognition of

airborne aerosol transmission, as well as presymptomatic and asymptomatic transmission.

Oshitani contrasts the Japanese strategy, nailing almost every important feature of the pandemic early on, with the Western response, trying to eliminate the disease "one by one" when that's not necessarily the main way it spreads. Indeed, Japan got its cases down, but kept up its vigilance: When the government started noticing an uptick in community cases, it initiated a state of emergency in April and tried hard to incentivize the kinds of businesses that could lead to super-spreading events, such as theaters, music venues, and sports stadiums, to close down temporarily. Now schools are back in session in person, and even stadiums are open—but without chanting.

It's not always the restrictiveness of the rules, but whether they target the right dangers. As Morris put it, "Japan's commitment to 'cluster-busting' allowed it to achieve impressive mitigation with judiciously chosen restrictions. Countries that have ignored

super-spreading have risked getting the worst of both worlds: burdensome restrictions that fail to achieve substantial mitigation. The U.K.'s recent decision to limit outdoor gatherings to six people while allowing pubs and bars to remain open is just one of many such examples."

Could we get back to a much more normal life by focusing on limiting the conditions for super-spreading events, aggressively engaging in cluster-busting, and deploying cheap, rapid mass tests—that is, once we get our case numbers down to low enough numbers to carry out such a strategy? (Many places with low community transmission could start immediately.) Once we look for and see the forest, it becomes easier to find our way out.