The One and the Many

We've counted the viral spread across peoples; we need to count it within people.

BY SIDDHARTHA MUKHERJEE

In the third week of February, as the COVID-19 epidemic was still flaring in China, I arrived in Kolkata, India. I woke up to a sweltering morning—the black kites outside my hotel room were circling upward, lifted by the warming currents of air—and I went to visit a shrine to the goddess Shitala. Her name means "the cool one"; as the myth has it, she arose from the cold ashes of a sacrificial fire. The heat that she is supposed to diffuse is not just the fury of summer that hits the city in mid-June but also the inner heat of inflammation. She is meant to protect children from smallpox, heal the pain of those who contract it, and dampen the fury of a pox epidemic.

The shrine was a small structure within a temple a few blocks from Kolkata Medical College. Inside, there was a figurine of the goddess, sitting on a donkey and carrying her jar of cooling liquid—the way she has been depicted for a millennium. The temple was two hundred and fifty years old, the attendant informed me. That would date it to around the time when accounts first appeared of a mysterious sect of Brahmans wandering up and down the Gangetic plain to popularize the practice of *tika*, an early effort at inoculation. This involved taking matter from a smallpox patient's pustule—a snake pit of live virus—and applying it to the pricked skin of an uninfected person, then covering the spot with a linen rag.

The Indian practitioners of *tika* had likely learned it from Arabic physicians, who had learned it from the Chinese. As early as 1100, medical healers in China had realized that those who survived smallpox did not catch the illness again (survivors of the disease were enlisted to take care of new victims), and inferred that the exposure of the body to an illness protected it from future instances of that illness. Chinese doctors would grind smallpox scabs into a powder and insufflate it into a child's nostril with a long silver pipe.

Vaccination with live virus was a tightrope walk: if the amount of viral inoculum in the powder was too great, the child would succumb to a full-fledged version of the disease—a disaster that occurred perhaps one in a hundred times. If all went well, the child would have a mild experience of the disease, and be immunized for life. By the seventeen-hundreds, the practice had spread throughout the Arab world. In the seventeen-sixties, women in Sudan practiced *tishteree el jidderee* ("buying the pox"): one mother haggling with another over how many of a sick child's ripe pustules she would buy for her own son or daughter. It was an exquisitely measured art: the most astute traditional healers recognized the lesions that were likely to yield just enough viral material, but not too much. The European name for the disease, variola, comes from the Latin for "spotted" or "pimpled." The process of immunizing against the pox was called "variolation."

Lady Mary Wortley Montagu, the wife of the British Ambassador to Constantinople, had herself been stricken by the disease, in 1715, leaving her perfect skin pitted with scars. Later, in the Turkish countryside, she witnessed the practice of variolation, and wrote to her friends in wonder, describing the work of one specialist: "The old woman comes with a nut-shell full of the matter of the best sort of small-pox, and asks what vein you please to have opened," whereupon she "puts into the vein as much matter as can lie upon the head of her needle." Patients retired to bed for a

Measurement will help identify factors affecting the severity of COVID-19 cases.

ILLUSTRATION BY ALEXANDER GLANDIEN

couple of days with a fever, and, Lady Montagu noted, emerged remarkably unscathed. "They have very rarely above twenty or thirty in their faces, which never mark; and in eight days' time they are as well as before their illness." She reported that thousands safely underwent the operation every year, and that the disease had largely been contained in the region. "You may believe I am well satisfied of the safety of this experiment," she added, "since I intend to try it on my dear little son." Her son never got the pox.

In the centuries since Lady Montagu marvelled at the efficacy of inoculation, we've made unimaginable discoveries in the biology and epidemiology of infectious disease, and yet the COVID-19 pandemic poses no shortage of puzzles. Why did it spread like wildfire in Italy, thousands of miles from its initial epicenter, in Wuhan, while India appears so far to have largely been spared? What animal species transmitted the original infection to humans?

But three questions deserve particular attention, because their answers could change the way we isolate, treat, and manage patients. First, what can we learn about the "dose-response curve" for the initial infection—that is, can we quantify the increase in the risk of infection as people are exposed to higher doses of the virus? Second, is there a relationship between that initial "dose" of virus and the severity of the disease—that is, does more exposure result in graver illness? And, third, are there quantitative measures of how the virus behaves in infected patients (e.g., the peak of your body's viral load, the patterns of its rise and fall) that predict the severity of their illness and how infectious they are to others? So far, in the early phases of the COVID-19 pandemic, we have been measuring the spread of the virus *across* people. As the pace of the pandemic escalates, we also need to start measuring the virus *within* people.

Most epidemiologists, given the paucity of data, have been forced to model the spread of the new coronavirus as if it were a binary phenomenon: individuals are either exposed or unexposed, infected or uninfected, symptomatic patients or asymptomatic carriers. Recently, the Washington *Post* published a particularly striking online simulation, in which people in a city were depicted as dots moving freely in space—uninfected ones in gray, infected ones in red (then shifting to pink, as immunity was acquired). Each time a red dot touched a gray dot, the infection was transmitted. With no intervention, the whole field of dots steadily turned from gray to red. Social distancing and isolation kept the dots from knocking into one another, and slowed the spread of red across the screen.

This was a bird's-eye view of a virus radiating through a population, seen as an "on-off" phenomenon. The doctor and medical researcher in me—as a graduate student, I was trained in viral immunology—wanted to know what was going on *within* the dots. How much virus was in *that* red dot? How fast was it replicating in *this* dot? How was the exposure—the "touch time"—related to the chance of transmission? How long did a red dot remain red—that is, how did an individual's infectiousness change over time? And what was the severity of disease in each case?

What we've learned about other viruses—including the ones that cause AIDS, SARS, and smallpox—suggests a more complex view of the disease, its rate of progression, and strategies for containment. In the nineteen-nineties, as researchers learned to measure how much H.I.V. was in a patient's blood, a distinct pattern emerged. After an infection, the virus count in the blood would rise to a zenith, known as "peak viremia," and patients with the highest peak viremia typically became sicker sooner; they were least able to resist the virus. Even more predictive than the peak viral load was the so-called set point—the level at which someone's virus count settled after its initial peak. It represented a dynamic equilibrium that was reached between the virus and its human host. People with a high set point tended to progress more rapidly to AIDS; people with a low set point frequently proved to be "slow progressors." The viral load—a continuum, not a binary value—helped predict the nature, course, and transmissibility of the disease. To be sure, every virus has its own personality, and H.I.V. has traits that make viral load especially revealing: it causes a chronic infection, and one that specifically targets cells of the immune system. Yet similar patterns have been observed with other viruses.

And, immunologically, that's not surprising. If your system is able to combat viral replication with some efficiency—owing to your age, your genetics, and other indices of immune competence—you'll have a lower set point. Could a lower initial exposure, as with children treated with *tika*, also lead to a lower set point? Faced with a smaller challenge, the immune system could have a greater chance of controlling the pathogen. In contrast, if you're inundated with multiple high-dose exposures, the swiftly replicating invader could gain ground that the immune system might be hard-pressed to reconquer.

An ingenious study on the relationship between the intensity of viral exposure and infectivity in human beings comes from a team at the Fred Hutchinson Cancer Research Center and the University of Washington, in Seattle. In 2018, an epidemiologist and statistician named Bryan Mayer joined a group of physicians and biologists who were researching a problem that seemed, on its face, almost impossible to tackle. Mayer, who is in his mid-thirties, is soft-spoken and precise: he uses words carefully, and speaks in long, slow sentences. "Even as a graduate student, I was interested in the idea of a *dose* of a virus or a pathogen," he told me. "But the problem is that the initial dose is often impossible to capture, because you only know a person is infected after he or she has *been* infected." Most infectious diseases can only be viewed in a rear-view mirror: by the time a patient becomes a patient, that critical moment of transmission has already passed.

But the researchers found an unusual resource: a cohort of new mothers and their children in Kampala, Uganda. A few years earlier, a pediatrician named Soren Gantt and a team of doctors examined these women, and asked them to provide oral swabs for a year. Then they measured how much the women shed a virus called HHV-6, which is usually spread through oral secretions to an infant after birth, and which causes fever and a red whole-body rash. It was now possible to investigate how the amount of virus-shedding—the "dose" of exposure—affected the likelihood of

a newborn infant becoming infected. Gantt, Mayer, and their colleagues had devised a way to eavesdrop on the dynamics of the transmission of a human viral infection from the very start. "Our data confirmed that there's a dose-response relationship in viral transmissions for HHV-6," Mayer told me. "The more virus you shed, the more likely you are to infect others." He'd managed to turn around the rearview mirror of epidemiology.

There's another aspect of transmission and disease, however: the host immune response. Viral attack and the immune system's defense are two opposing forces, constantly at odds. The Russian immunologist Ilya Metchnikoff, working in the early nineteen-hundreds, described the phenomenon as "the struggle"—or *Kampf*, in German editions of his work. Metchnikoff imagined an ongoing battle between microbe and immunity. The *Kampf* was a matter of ground gained or lost. What was the total "force" of the microbial presence? What host factors—genetics, prior exposure, baseline immune competence—were limiting the microbial invasion? And then: was the initial equilibrium tipped toward the virus, or toward the host?

That raises the second question—does a larger viral "dose" result in more severe disease? It's impossible to erase from one's memory the image of Li Wenliang, the thirty-three-year-old Chinese ophthalmologist who sounded the alarm on the first COVID-19 cases, in his final illness; a photograph shows him crimson-faced, sweating, and struggling to breathe in a face mask, shortly before his death. Then there's the unexpected death of Xia Sisi, a twenty-nine-year-old doctor from Union Jiangbei Hospital of Wuhan, who had a two-year-old child and, the *Times* reported, loved Sichuan hot pot. Another Chinese health-care worker, a twenty-nine-year-old nurse in Wuhan, fell so critically ill that she started hallucinating; later, she would describe herself as "walking on the edge of death."

Could the striking severity of their disease—twenty- and thirty-year-olds with COVID-19 generally experience a self-limited, flu-like illness—be correlated with the amount of virus to which they were initially exposed? At least two E.R. doctors in the United States, both on the front lines of the pandemic, have also fallen critically ill; one of them, in Washington State, is only in his forties. To go by available data from Wuhan and Italy, health-care workers don't necessarily have a higher fatality rate, but do they suffer, disproportionately, from the most severe forms of the disease? "We know the high mortality in older people," Peter Hotez, an infectious-disease specialist and vaccine scientist at Baylor College of Medicine, told CNN. "But, for reasons that we don't understand, front-line healthcare workers are at great risk for serious illness despite their younger age."

Some suggestive research has been done with other viruses. In animal models of influenza, it's possible to precisely quantify exposure intensity, and mice who were given higher doses of certain influenza viruses developed a more severe form of the disease. Yet the degree of correlation between dose and disease severity varied widely from one strain of the flu to the next. (Curiously, in one study a higher initial load of respiratory syncytial virus, which can cause pneumonia, especially in young children, correlated *negatively* with severe disease—although another study suggests that the correlation is positive with toddlers, the most affected patient population.)

What sparse evidence we have about coronaviruses suggests that they may follow the pattern seen in influenza. In a 2004 study of the coronavirus that causes SARS, a cousin of the one that causes COVID-19, a team from Hong Kong found that a higher initial load of virus—measured in the nasopharynx, the cavity in the deep part of your throat above your palate—was correlated with a more severe respiratory illness. Nearly all the SARS patients who came in initially with a low or undetectable level of virus in the nasopharynx were found at a two-month follow-up to be still alive. Those with the highest level had a twenty- to forty-per-cent mortality rate. This pattern held true regardless of a patient's age, underlying conditions, and the like. Research into another acute viral illness, Crimean-Congo hemorrhagic fever, reached a similar conclusion: the more virus you had at the start, the more likely you were to die.

Perhaps the strongest association between the intensity of exposure and the intensity of subsequent disease is seen in measles research. "I want to emphasize that measles and COVID-19 are different diseases caused by very different viruses with different behaviors," Rik de Swart, a virologist at Erasmus University, in Rotterdam, cautioned when we spoke, "but in measles there are several clear indications that the severity of illness relates to the dose of exposure. And it makes immunological sense, because the interaction between the virus and the immune system is a race in time. It's a race between the virus finding enough target cells to replicate and the antiviral response aiming to eliminate the virus. If you give the virus a head start with a large dose, you get higher viremia, more dissemination, higher infection, and worse disease."

He described a study from 1994 in which researchers gave monkeys different doses of the measles virus and found that higher infection doses were associated with earlier peaks in viremia. In human beings, de Swart added, the best evidence comes from studies in sub-Saharan Africa. "If you acquire measles through household contacts, where the density and dose of exposure is the highest—you might be sharing a bed with an infected child—then you typically have a higher risk of developing more severe illness," he said. "If a child contracts the disease through playground or casual contact, the disease is usually less severe."

I discussed this aspect of infection with the Harvard virologist and immunologist Dan Barouch, whose lab is among those that are working toward a vaccine against SARS-CoV-2, the virus that causes COVID-19. He told me that ongoing studies with macaques are investigating the relationship between the initial dose of the SARS-CoV-2 viral inoculum and the amount of virus in lung secretions at a later time. He believes that there may be a correlation. "If we extended this logic to humans, we would expect a similar relationship," he said. "And, logically, the larger amount of virus should trigger more severe disease by prompting a brisker inflammatory response. But that is still speculative. The relationship between initial viral dose and severity remains to be seen."

To answer the third question—whether we can track a COVID-19 patient's viral

load in a way that helps us predict the course of the disease—we'll need more quantitative research into SARS-CoV-2 counts within patients. One unpublished German study has measured viral loads on oral swabs taken of both symptomatic and asymptomatic individuals. Initially, it was reported that patients who experienced no symptoms had slightly higher loads than those who fell ill. The results were curious. But at the time only seven patients had been studied. Sandra Ciesek, the director of the Institute of Medical Virology, in Frankfurt, who was running the study, told me that no significant differences between the two groups emerged as a larger patient population began to be sampled. "In swabs, we don't know of a correlation," she informed me. The problem with measuring viral loads in a swab is that it is "affected by preanalytic factors, such as the way in which the swab is taken," she added. Oral swabs are notoriously affected by small variations in how they're done. "But a correlation with severe disease may well be true for the viral load in blood." Joshua Schiffer, a clinical virologist at the Fred Hutchinson Center, and a co-author of the HHV-6 study, reports that more stringent nasal-swabbing methods for a range of respiratory viruses have yielded consistent, reliable viral-load counts, and that these loads have generally tracked well with disease symptoms and progression. In a paper published online by *The Lancet Infectious Diseases* in March, researchers at the University of Hong Kong and Nanchang University reported that viral loads in nasopharyngeal swabs from a group of patients with severe COVID-19 were sixty times higher, on average, than the loads among patients with a mild form of the disease.

As the virus continues to cyclone across the world, we will begin to find quantitative answers to these questions of how exposure intensity and subsequent viral loads relate to the clinical course of COVID-19. We will supplement the bird's-eye view with the worm's-eye view. How will these insights change the way we manage patients, hospitals, and populations?

Start with the relationship between exposure intensity and infection. Think, for a moment, of how we monitor those who work with radiation. Using radiation dosimetry, we quantify someone's total exposure, and we set limits on it. We already know how critical it is for doctors and nurses to limit exposure to the coronavirus by using protective equipment (masks, gloves, gowns). But for health-care workers on the front lines of the COVID-19 pandemic, especially in places where protective equipment is scarce, we might also keep track of total exposure, and put in place viral-dosimetry controls, so that one individual can avoid repeated interactions with some set of highly contagious patients.

Establishing a relationship between dose and disease severity could, in turn, affect patient care. If we could identify pre-symptomatic patients who were likely exposed to the highest doses of viruses—someone cohabitating or socializing with multiple sick family members (as with the close-knit Fusco family of Freehold, New Jersey, which has had four deaths), or a nurse exposed to a set of patients shedding large amounts of the virus—we might predict a more severe experience of the disease, and give them priority when it came to limited medical resources, so that they could be treated faster, earlier, or more intensively.

And, finally, the care of COVID-19 patients could change if we began to track virus counts. These parameters could be gauged using fairly inexpensive and easily available laboratory methods. Imagine a two-step process: first, identifying infected patients, and then quantifying viral loads in nasal or respiratory secretions, particularly in patients who are likely to require the highest level of treatment. Correlating virus counts and therapeutic measures with outcomes might result in different strategies of care or isolation.

The value of a quantitative approach applies to clinical studies as well. Clinical drug trials are typically more informative when run on subjects who aren't yet critical; once the subjects have reached that stage, any therapy might be too little, too late. And if the disease course in such patients is followed using viral-load metrics, rather than by tracking symptoms alone, the effect of a drug in different trials can be compared more easily and accurately.

We will also want to be able to identify people who have recovered from infection, have become immune to SARS-CoV-2, and are no longer contagious. Such people must meet two criteria: they must have a measured absence of viral shedding, and they must have signs of persistent immunity in their blood (something readily determined by an antibody test). As the Chinese discovered with smallpox in the twelfth century, such individuals—especially those who are health-care workers—are of particular value to medicine: barring any decay in immunity, they can generally tend to the sickest patients without getting sick themselves.

My clinical practice is in oncology. Measurement and enumeration are the mainstays of medicine for people in my field: the size of a tumor, the number of metastases, the exact shrinkage of a malignant mass after chemotherapy. We talk about "risk stratification" (categorizing patients according to health status) and the "stratification of response" (categorizing patients according to their response to treatment). I am able to spend half an hour or more with every patient to describe risk, explain how a remission is measured, and carefully devise a clinical plan.

A pandemic, by contrast, goes hand in hand with panic. Chaos reigns. Italian doctors are hanging I.V. drips on makeshift poles for patients lying on makeshift cots in makeshift wards. Measurement—viral-load testing—can seem like an improbable indulgence under such circumstances. But this crisis will require that we stratify and assess risk, and deploy dwindling resources in the most effective manner.

The word "epidemiology" is derived from "epi" and "demos"—"above the people." It is the science of aggregation, the science of the many. Yet it works most effectively when it moves in step with medicine, the science of the one. On the morning I visited the Shitala shrine in Kolkata, the goddess of bygone population-decimating epidemics was also serving as the personal goddess of a mother who had brought a child with a weeklong fever. To win the *Kampf* against COVID-19, it's essential to trace the course of the virus as it moves through populations. But it's equally essential to measure its course within a single patient. The one becomes the many. Count both; both count. ♦