

The Ethics of Implementing Emergency Resource Allocation Protocols

Richard H. Dees, Jonathan Herington, Marianne Chiafery, Jessica C. Shand, Carl T. D'Angio, Chin-Lin Ching, and Margie Hodges Shaw

ABSTRACT

We explore the various ethical challenges that arise during the practical implementation of an emergency resource allocation protocol. We argue that to implement an allocation plan in a crisis, a hospital system must complete five tasks: (1) formulate a set of general principles for allocation, (2) apply those principles to the disease at hand to create a concrete protocol, (3) collect the data required to apply the protocol, (4) construct a system to implement triage decisions with those data, and (5) create a system for managing the consequences of implementing the protocol, including the effects on those who must carry out the plan, the medical staff, and the general public. Here we illustrate the complexi-

ties of each task and provide tentative solutions, by describing the experiences of the Coronavirus Ethics Response Group, an interdisciplinary team formed to address the ethical issues in pandemic resource planning at the University of Rochester Medical Center. While the plan was never put into operation, the process of preparing for emergency implementation exposed ethical issues that require attention.

In the very early days of the COVID-19 pandemic, as rumors surfaced of informal and chaotic rationing of mechanical ventilators in Northern Italy and New York City,¹ much attention was paid to creating criteria for the ethical and orderly allocation of ventilators.² Little attention, however,

Richard H. Dees, PhD, is Professor of Philosophy and Bioethics at the University of Rochester in Rochester, New York. richard.dees@rochester.edu

Jonathan Herington, PhD, is Assistant Professor of Philosophy and Bioethics at the University of Rochester.

Marianne Chiafery, DNP, PNP-C, MS, HEC-C, is Associate Professor Emerita of Nursing and of Health Humanities and Bioethics and a clinical ethics consultant at the University of Rochester Medical Center.

Jessica C. Shand, MD, MHS, FAAP, is Associate Professor of Pediatrics and of Health Humanities and Bioethics at the University of Rochester Medical Center.

Carl T. D'Angio, MD, is Professor of Pediatrics and of Health Humanities and Bioethics and Chief of the Division of Neonatology at the University of Rochester Medical Center.

Chin-Lin Ching, MD, FACP, is Associate Professor of Clinical Medicine at the University of Rochester Medical Center and Director of the Palliative Care Consultation Service and Cochair of the Ethics Consultation Service at Highland Hospital in Rochester, New York.

Margie Hodges Shaw, JD, PhD, HEC-C, is Associate Professor of Law and Bioethics, Director of the Division of Bioethics in the Department of Health Humanities and Bioethics, and an ethics consultant at the University of Rochester Medical Center.

was paid to the many ethical and practical challenges that would be faced by any medical center that tried to implement such protocols. Yet allocation plans are ethically worthless if they cannot be implemented effectively.

In this article, we explore the various ethical challenges that arise during the practical implementation of an emergency resource allocation protocol. We argue that to implement an allocation plan in a crisis, a hospital system must complete five tasks: (1) formulate a set of general principles for allocation, (2) apply those principles to the disease at hand to create a concrete protocol, (3) collect the data required to apply the protocol, (4) construct a system to implement triage decisions with those data, and (5) create a system for managing the consequences of implementing the protocol, including the effects on those who must carry out the plan, the medical staff, and the general public. Here we illustrate the complexities of each task and provide tentative solutions by describing the experiences of the Coronavirus Ethics Response Group (CERG), an interdisciplinary team formed to address the ethical issues in pandemic resource planning at the University of Rochester Medical Center (URMC). While the plan was never put into operation, the process of preparing for emergency implementation exposed ethical issues that require attention.

Each of these tasks requires ethical, not just practical, judgment. In particular, the implementation of a scarce resource allocation protocol raises two significant moral challenges. First, no implementation decision is value-free. Even detailed pre-pandemic triage plans contain lacuna in their ethical framework, and resolving these oversights requires moral judgment. Because such judgments are not value-free, they lack broader legitimacy without mechanisms for community review and input.³ Thus, implementation requires consultation with the community before any of the five tasks can be completed.

Second, implementation involves significant uncertainty. Faithfully translating pre-pandemic ethical frameworks into clinical (and nonclinical) criteria requires detailed knowledge of disease dynamics, the structure of the surges of cases that are likely to occur, and the effectiveness of interventions, all of which are deeply uncertain in the context of a novel disease. Navigating this uncertainty is not morally neutral. Deciding which possibilities to take seriously is not simply a matter of deciding which are more likely; it also includes an assessment of the impact that those possibili-

ties will have on the community as a whole and on particular communities. Weighing these possibilities requires moral judgment.

1. COMMUNITY REPRESENTATION

The inclusion of the community in both the formulation and the implementation of a ventilator protocol is morally crucial. Institutional legitimacy depends on the trust that the community places in it, which depends in part on whether the community feels that the institution is acting in its interest.⁴ Simply imposing a plan on a patient population disrespects people by treating them as mere subjects of actions from the medical center and not as agents who ought to have some input into their care. This consultation is especially important when the system may need to make policies that will override individual preferences. Since the principles of allocation will apply to the community, the community has a right to be involved in the process that creates them.

Ensuring representation of community members in the process also serves a basic moral goal of equity and inclusion. Hospital ethics committees are usually made up mostly of hospital professionals, usually physicians and nurses. Persons of color and persons living in poverty, who may well be disproportionately affected by pandemic disease, are typically underrepresented in such groups. Providing a mechanism to include the values and perspectives of these groups helps to ensure that the range of moral considerations included will not be too narrow and that the protocol will not reinforce systemic bias and exclusion.

In addition to addressing the moral imperative of inclusion, ensuring community representation serves several important instrumental goals. Any protocol—even the most well-intentioned plan—that is simply imposed on a patient population could lead to resentments that would undermine the ability of a medical system to operate for the good of that community. For just this reason, transparency is an important public health value,⁵ and it is essential for creating trust. Many underrepresented, underserved, and marginalized communities lack trust in the medical system owing to long-standing social and health inequities, many of which were laid bare by the COVID-19 pandemic itself. In Rochester in particular, the relationship between the Black community and the healthcare system—and the URMC

in particular—is especially fraught.⁶ That deeper mistrust cannot be solved in the midst of a pandemic, but trying to ameliorate that lack of trust is essential.

While the 2015 New York State (NYS) Ventilator Allocation Guidelines, which became the basis of the CERG's work, were created with considerable community involvement,⁷ they assumed that a shortage would be created by an influenza pandemic, not a novel coronavirus, so the CERG knew that some modifications would be required (see secs. 2 and 3 below). We wanted any such changes to represent the needs and values of a diverse cohort of Rochester community members. From its beginning, the CERG included standing community members of the URMC Ethics Committee, and we collaborated with the medical center's Office of Equity and Inclusion to establish a working group for community engagement. Members of this working group included hospital employees who regularly engage community and health advocacy groups, as well as members of underrepresented community and health advocacy groups not directly employed by the hospital. Because the committee had to be created quickly, its membership was not constituted systematically but relied on existing relationships with community partners. While this is not a perfect deliberative process and introduces questions about the legitimacy of self-selecting community members, in the weeks before a surge was expected the community partners' existing understanding of the hospital's operations and values was thought to be important. Nevertheless, their input into the discussions was significant and important (see sec. 3). In addition, members of the community were to be incorporated into ventilator triage teams (see sec. 4) to ensure fairness and transparency, requiring further familiarity with the hospital. Decisions to withdraw ventilator support would be made not by detached and unresponsive groups of doctors making arbitrary judgments but by a diverse group of people who would be applying—and would be seen to be applying—a procedure as free from bias as possible.

In addition, the legitimacy of a protocol requires effective communication with the public as a whole. If the ventilator allocation protocol had been implemented, the public would need to understand the plan to trust that it was being fairly and competently implemented. Critically, patients and families admitted to the hospital would need to understand that the circumstances

would constrict the kinds of decisions that patients were allowed to make. In ordinary circumstances, patients have the autonomy to decide which treatments they wish to receive, and they are generally able to insist on receiving any treatment that might be effective. Basic informed consent requires that patients entering the hospital understand the conditions under which they are being admitted. In a crisis, mechanical ventilators, which are ordinarily offered to anyone who might benefit from them, would become analogous to donor organs: they could be offered only to the patients who can most benefit from them. But the allocation protocol is even more exacting: some patients would have their ventilators reallocated to others if they were not recovering fast enough. Such a reallocation would almost certainly result in that patient's death. Hence, newly admitted patients and their families need to understand that potentially life-saving treatment might be withdrawn if reallocation of a ventilator would result in the most lives saved. Worse yet, saving other patients may require those reallocations to occur rapidly, and the resultant removal may result in a rapid death, without allowing the families time with the patients before they die. Families need to be aware of this possibility before their loved ones are placed on ventilators.

Second, the community as a whole needs to understand the conditions under which health care is provided. Emergency standards of care, including a ventilator allocation protocol, are controversial. Ensuring community input into the deliberation and implementation process should promote confidence in the process and in the final protocol. Whether in fact such steps will gain the confidence of the community is another matter, but the moral obligation of the bodies determining policy is to be transparent and to include voices from the community.

2. SELECTING GENERAL PRINCIPLES FOR PANDEMIC TRIAGE

Besides engaging the community, the first task of a community-informed allocation protocol is to determine the precise ethical principles of allocation. Those principles have received considerable attention in the literature;⁸ that attention is warranted since the rest of the tasks will be guided by the principles a hospital system intends to instantiate. Yet that discussion tends to focus only on the initial allocation of ventilators

and not on the more difficult task of deciding when a ventilator needs to be reallocated.⁹ Given time constraints experienced during the initial weeks of a pandemic, rehashing the intricate arguments for basic allocation principles is not feasible. In Rochester, initial data projections for worst-case scenarios suggested that the hospital might need to implement emergency standards of care in as soon as two weeks, so the CERG decided simply to accept the 2015 NYS Task Force on Life and Laws Ventilator Allocation Guidelines.¹⁰ Among the most developed and systematic ventilator allocation guidelines nationally, the 2015 NYS guidelines were widely praised, and as an entity in NYS, any requirements from the State Department of Health would likely follow them.

The NYS guidelines rest on one simple principle: saving as many lives as possible—where the measure of a saved life is whether the patient survives to hospital discharge (rather than, say, survival after one year or after five years). Contrary to other prominent protocols, post-discharge life-expectancy considerations that might disfavor the elderly, people with disabilities, or those with chronic diseases are excluded.¹¹ The only consideration is how likely it is that a patient will recover with ventilator support and survive to leave the hospital. Adult patients, pediatric patients, and neonatal patients are assessed in different ways, but all patients with similar chances of survival should be treated similarly. Among the patients with roughly equal chances of survival, determinations of who should either receive a ventilator or be withdrawn from a ventilator should be decided by a lottery. However, the guidelines do allow that within the groups that have a similar chance of survival, pediatric and neonatal patients can be favored in the lottery.¹² The logic behind this preference is that the very young have not had a fair chance to live a life. Other things being equal, then, they should be given the chance to do so.¹³ The guidelines did not, however, give any preference to anyone else, including pregnant patients.¹⁴

3. IMPLEMENTING THE PRINCIPLES FOR COVID-19

Once moral goals are established, hospital systems must select a plan to operationalize them in the context of the specific mass casualty and resource scenario. Simple moral maxims, even “save the most lives,” are not readily in-

terpretable into criteria for prioritization of different patients, and there is considerable indeterminacy in how to realize such goals in the context of a clinically novel, rapidly emerging pandemic disease. As we discovered, even with the detailed implementation plan provided in the 2015 NYS guidelines, there were a number of lacuna and uncertainties that required value judgments to resolve.

To minimize the role that implicit biases can play, the guidelines required that measures of disease severity should be based on laboratory and clinical indicators that do not require extensive clinical interpretation. For that reason, the guidelines measure severity of illness for adults using the Sequential Organ Failure Assessment (SOFA), which scores each of six organ systems, mostly using lab values, on a scale of 0–4, where a 4 indicates significant organ failure. Lower scores, then, indicate a patient who is relatively less ill. For pandemic flu, patients with SOFA scores ≤ 7 would have the highest priority (coded Red). If ventilators remained, patients with SOFA scores 8–11 would receive them (coded Yellow), and only then would ventilators be assigned to patients with SOFA scores 12 or higher (coded Blue). Within each category, patients were considered to have similar chances of survival, so patients within a category would be allocated a ventilator by lottery if there were insufficient machines for all the eligible patients.

Once allocated, patients would be given a “sufficient trial on the ventilator to determine whether the patient was benefiting from the treatment.”¹⁵ But they must be reassessed on a regular basis. Since patients who do not improve can tie up one ventilator for weeks, they may prevent several other patients who might recover quickly from using that same ventilator. In the unmodified NYS plan, patients are reassessed after 48 hours, again at 120 hours, and then every 48 hours thereafter. Crucially, if patients did not improve significantly (i.e., if they have less than a three-point improvement in their SOFA score), their priority status would be downgraded, and if unventilated patients have a higher priority, then they would lose their ventilator. While this protocol was extremely detailed relative to the other pre-pandemic plans available, during implementation we identified three problems that required value judgments to resolve.

First, COVID-19 is not influenza, and the particular clinical guidelines in the New York influenza plan were never meant to apply to other

diseases. With any novel disease, however, such a task will be based, at best, on limited data and educated guesswork.¹⁶ The first major task, then, was to adapt the guidelines to the clinical progression of COVID-19 by adjusting clinical thresholds and reassessment windows to realize the goal of saving the most lives. In our experience, initial data quickly showed that if the hospital experienced a surge of COVID-19 patients who needed ventilators, then the unmodified NYS protocol would result in very few survivors. At intubation, almost all of the COVID-19 patients were coded Red, the highest priority. But at 48 hours, almost all of them had deteriorated, so they would have shifted to Yellow or Blue and would have been replaced on ventilators with incoming patients who had a higher Red coding. The result would have been ventilator “churning,” where almost no patients would be allowed sufficient time on a ventilator to survive. Thus, for COVID-19, the unmodified guidelines requirement that patients maintain ventilator support only if they quickly improve, included as a way to steward resources to maximize the number of lives saved, would counterproductively result in high mortality and few eventual survivors.

The reason is simple. For influenza patients requiring mechanical ventilation, the expectation is that the disease is nearly at its worst when a patient requires respiratory support, and a relatively short trial of intubation will either result in improvement or prove futile. For COVID-19, however, patients often need ventilatory support in earlier phases of respiratory failure, and they will continue to get worse for days or weeks before any improvements are seen.

To save the most lives, then, the ventilator allocation protocol had to be adjusted. In the early days of the pandemic, however, there were few clinical data, so the adjustments could not be guided by observational studies or sophisticated modeling (some of which have subsequently been performed¹⁷). Data specialists at the URMC had already created a system that automatically collated lab records to calculate SOFA scores—a process that became easier later as those specialists created a novel system to collect the SOFA scores on every ventilated patient in the hospital (see sec. 4). The CERG used the available SOFA data to crudely test different time frames for reevaluations and develop an alternative protocol. We looked at delaying the first evaluation until 72 hours and at a system in which no improvement was expected at 72 hours, so that patients

were scored in exactly the same way at 72 hours as at intubation. But such a protocol, like the original one, would have removed ventilators from almost all patients before they had a reasonable chance to recover. We also looked at setting the initial reevaluation at 5 days, but even with that much time, expecting improvements compared to the assessments at intubation would have removed ventilators. However, the preliminary data suggested that if we did not expect patients to improve between intubation and day 5, we could simply evaluate them at day 5 using the same criteria as those used at intubation. By doing so, most of the patients who ultimately survived would not have their ventilators removed prematurely. After five days, patients would be reassessed every two days, but since COVID-19 patients often recovered slowly, they only had to show modest improvements (one point of SOFA score) to avoid being downgraded.

The new protocol is clearly better for COVID-19 patients than the old one, but how well it would have worked to save lives is harder to say. Even with a year’s worth of data now, we still cannot tell how well the protocol would have worked in a real crisis. In truth, this problem is endemic. Every disease must be assessed for its particular characteristics, and any ventilator allocation plan will have to be adjusted to those characteristics. For any novel disease, like COVID-19, those assessments will have to be done on the fly in circumstances of uncertainty. Ethically, then, we must adapt protocol to fit the situation in which we find ourselves, but we must accept that any goals that we create in an allocation plan will be imperfectly met. We are then morally required to make what adjustments we can to the best of our abilities. The need for a constant reevaluation of the protocols is essential. One implication is that we are also morally required to collect the data necessary to make those adjustments and to improve the processes used (see sec. 4).

Second, the application of the NYS protocol needed modification because its application to nonadults was less well defined and required a number of clinical and value judgments. SOFA scores are valid only for adults, and the NYS Task Force did not find any of the then-existing pediatric scoring systems adequate. Instead, it merely recommended the use of “a structured decision-making process that carefully considered only specific clinical factors” that would rely on both objective lab results and physician judgment.¹⁸ For neonatal patients, assessments would rely

even more on physician judgments to determine the likelihood for survival.¹⁹ In both cases, a similar timeline for assessments was recommended, though what would count as “significant improvement” was to be determined by whatever criteria individual programs adopted.

Thus, while adult SOFA scores were already calculated in the medical records system, any pediatric or neonatal scoring system would have to be created *ex nihilo*. The CERG therefore designated a team to construct a set of criteria to assign both neonatal and pediatric patients to priority bands in ways that reflected the general principles of the NYS guidelines and that, we hoped, could be automated. Perfecting the scoring system, as well as the allocation protocols to go with them, was seen as less urgent since children infected with COVID-19 were much less likely at that time to experience severe illness. Nevertheless, in the interests of fairness to all patients, they needed to be included in the allocation system. If all lives are equally valuable, then the lives of children are not more valuable than adults, and they should be subjected to the same jeopardy.

Given the limitations of data about neonatal and pediatric survival and the limitation of data collection, the measures the pediatric team developed were less exact than SOFA scores, and they were not fine-grained enough to make measuring “improvements” from one assessment to the next possible. Neonatal and pediatric patients, then, were simply placed in a priority group based on their current state of health at each assessment point. This system had the effect of giving some priority to children, since it did not require their condition to improve over time. However, they were given no other preferences. Children in lower-priority groups would be randomized to be removed from ventilators in the same ways that adults would be, but they were less likely to be placed in lower-priority groups.

Third, the decision to favor children in this way led to a potential quirk. The NYS guidelines gave no preferences to pregnant patients, so in principle a pregnant patient with a viable fetus could be subject to having their ventilator reallocated if they were not making sufficient improvements. Doing so would also terminate a viable fetus, which would be a double tragedy. Offering an emergency caesarean section would allow the child to survive, but doing so early would mean the birth of a child who would herself require a ventilator to survive. If children have some priority, then the child would receive

the ventilator over another adult patient, and nothing would be gained by withdrawing the ventilator from the mother. For that reason, the CERG decided to give priority to pregnant patients who were between 24 and 28 weeks of pregnancy. Before 24 weeks, fetuses have a less than 50 percent chance of survival, even with a ventilator, and after 28 weeks they are likely to survive without it. This compromise best balanced the deep desires that families have to have children with the unrelenting demands that a ventilator allocation plan imposes on everyone. In essence, this decision allows two patients to survive with one ventilator.

The final modification of the NYS guidelines on which the CERG deliberated was whether to incorporate social justice considerations into the allocation protocol. The NYS guidelines were silent on the use of equity considerations as a tiebreaker, and the CERG had to decide whether that silence meant that those considerations ought to be excluded, or whether they could be used as tiebreakers if warranted by local conditions. Indeed, other implementation plans have specified social justice tiebreaker criteria for patients within a priority band.²⁰ Indeed, the most significant flaw in the principles behind the NYS guidelines—a flaw that has become more evident from simulation studies conducted in the two years since the beginning of the pandemic—is that they may perpetuate structural inequalities in health that already exist. People with disabilities and the chronically ill are likely to have more severe, more complicated illness when they are admitted into the hospital and so are less likely to survive when they require respiratory support. In addition, people who are poor or underserved are disproportionately more likely to have chronic health issues, and they are more likely to distrust the medical system. For that reason, they, too, are likely to be sicker when they arrive at the hospital and therefore less likely to survive at the point at which they require ventilation. As a result, these groups will not only have higher rates of death but also be disproportionately denied respiratory support by an algorithm that focuses exclusively on probability of survival. Of course, for that very reason, they are also less likely to survive even if they do receive ventilation. The goal of saving the most lives (even excluding considerations of life expectancy) thus appears to be in tension with efforts to rectify health injustices.

But to make matters worse, we now know from studies of COVID-19 patients that SOFA scores systematically underestimate the rates of

survival of Black patients: Black patients have the same rates of survival as white patients who score two points lower on the SOFA scale.²¹ For Black patients, then, the goal of saving the most lives is undermined by the prominent use of SOFA scores in protocols like the one in New York.

These problems have become more evident, of course, as the pandemic has progressed. But the CERG was worried enough about them at the time to consider using some measure of disadvantage to modify the allocation protocol, as proposed by Douglas White and Bernard Lo.²² Like White and Lo, the CERG considered using the Area Deprivation Index (ADI),²³ which measures the relative deprivation of people living in particular census tracts, in the allocation protocol.

At the time, we rejected the use of the ADI to alter color categories for six reasons. First, doing so, we thought, would take us outside the parameters of the NYS guidelines and hence outside the CERG's charge. Second, and more importantly, using ADI scores to modify SOFA scores, we thought at the time, would lead to more deaths since patients with less chance of survival would be favored. Using it as a tiebreaker, on the other hand, would not violate the spirit of the guidelines, but it was also unlikely to alter the disproportionate effects of the protocol. Third, using the ADI would have presented some practical problems that would have been difficult to overcome in the short amount of time we had to complete the work. Incorporating the information would be a significant programming task, and doing so would therefore require a considerable expenditure of resources. Fourth, the Community Engagement subcommittee was especially worried that including ADI in the information conveyed to the triage teams could create opportunities for implicit bias since the teams could easily become aware of much about a patient by knowing their ADI. Fifth, the ADI was developed to be used in community-level interventions, not for individual allocations or predictions,²⁴ and so it did not seem to be an appropriate tool. Finally, the committee felt that while the problems of structural health inequities are significant and important, an allocation protocol did not seem like the best place to address them. The Community Engagement subcommittee suggested instead that a better use of resources would be to create outreach programs into underserved communities to help them seek healthcare earlier, before symptoms became too severe. Such measures do not address the larger injustices in healthcare,

of course, but they did attempt to acknowledge the problem and ameliorate the scale of the potential disparate impact.

4. DATA COLLECTION

The ability to assemble the data required by the protocols is essential. Without accurate and timely data with which to implement the prioritization, the best ethical principles are rendered moot. An important task in operationalizing the NYS guidelines was to automate the collection of SOFA scores on every ventilated patient.

Although ease of data collection may not be as important as explicitly ethical reasons when selecting principles for allocation, it does have significant ethical implications. First, a system that is easier to implement makes the automation just described possible. To implement the NYS guidelines requires a lot of data, but they are comparatively easy to assemble. Assessing other measures, like long-term survival,²⁵ would require doctors to make long-term prognoses that would require considerable clinical judgment, which, with a novel disease, would be little more than guesswork. Automation reduces (though it certainly does not eliminate) errors in data entry, and it significantly reduces the opportunities for implicit biases to play a role in the assignment of priority categories. Our colleagues at other institutions were forced to calculate SOFA scores by hand, which required hours of work from experienced clinicians to collect the requisite lab values for each patient who was on a ventilator or who might need a ventilator.

Such a process is understandably fraught with the possibility of error, and since it required each chart to be opened and examined, it created the possibility that biases—often unintended and unrecognized—could skew the scoring. Second, since calculating scores requires considerable time from people with considerable medical sophistication, an automated system freed those healthcare workers to perform other tasks, which might be crucial in a crisis, when so many providers are overworked. Third, automation helps make the triage process (see sec. 5) less personal and, we hope, less traumatic since the members of the triage teams would not be faced with detailed patient charts. While decisions to reallocate ventilators would inevitably create moral distress (see sec. 6), a more automated process would, we hope, make it more bearable.

To achieve this goal, the CERG was fortunate that the medical record system at the URM was already programmed to produce SOFA scores.

The URMC administration then granted the CERG access to data specialists in the Clinical and Translational Medicine Institute, who, working with specialists in the hospital data, created innovative and creative software applications that collected the SOFA scores for the use of the triage teams. In addition, they created a system that generated a list every 20 minutes for unventilated patients, randomized within each priority category, to determine who would get an available ventilator if more than one patient needed one and another list every 12 hours for ventilated inpatients, randomized within each priority group, that would be used to determine who would have their ventilators removed if there were not enough ventilators for higher-priority patients. The programming task was enormous, and it required considerable university resources, resources that seem well spent given the ethical benefits of the program. Building a similar data transfer for the neonatal and pediatric patients would have required an even greater effort, but the danger of a surge had passed by the time the neonatal and pediatric protocol was ready for programming, and so that programming work was never done.

To minimize biases, the data transmitted were stripped of patient identifiers except for medical record number and, as a check, birthday month and day (but not year). It also included data about the patient's bed location and an identifier for the ventilator the patient was using. To our knowledge, such programming work was not done in other medical centers, crucial though it would have been to implementing the protocol in the event that crisis standards of care had been implemented. Indeed, without such a system, allocation decisions would be ungainly and therefore untimely. Patients who would otherwise live could die waiting for decisions.

Finally, the collection of these data helps facilitate concurrent and future improvements to the triage allocation protocol. The CERG created a quality improvement team, kept strictly separate from the triage teams, which would have monitored the allocation outcomes of triage decisions. That team was tasked with overseeing how well the process was working, troubleshooting issues as they arise, and determining whether the triage decisions were in fact impacting certain communities more than others. The data that the CERG continued to collect through the subsequent waves of COVID-19 hospitalizations may also provide the means to construct better protocols in the future.

5. OPERATIONALIZING TRIAGE DECISION-MAKING

Like data collection, the ability to implement an allocation protocol depends on having a system in place in which triage decisions are made and transmitted to emergency departments and intensive care units (ICUs) to be carried out. Hospitals must create procedures that detail the exact process by which decisions will be made and, if possible, test those procedures. Trying to create them during a crisis would lead to confusion and chaos. Without such a system, time and lives will be lost making decisions. In this case, the details of the plan are less important from an ethical point of view than the existence of the plan itself.

Ethically, the system must be set up so that the teams can operate independently, without pressure from outsiders, including pressure from patient families, doctors, or administrators. Providing anonymized medical information helps avoid that pressure, as it also prevents the teams from manifesting any biases they may have. In addition, the process itself plays a key ethical role. In a period of shortage, chaos is likely, so a plan that is clear to everyone will mitigate the inevitable disruptions and ease some of the stress.²⁶ In addition, in a crisis, individual doctors may feel pressured to make decisions on the fly, without institutional support, and decisions would then be made haphazardly, without standards of fairness and without the clear or consistent application of any ethical principles. Besides the poor decisions that are likely in such a situation, it also places all of the moral burden on overworked and stressed physicians. The existence of the protocol and of the triage teams to implement it is, then, a crucial element in creating a fair system of allocation.

The CERG set up triage teams of six members, who would work in shifts to apply the allocation protocol. The teams were made up of volunteers from throughout the medical center who were not directly involved in emergency or ICU care and members of the community, so they would have no obvious conflicts of interest. Members of the community were given the training necessary to handle protected health data. Teams were trained in the use of the protocol, and trials were conducted of the system that would be used. The teams were big enough to allow every patient decision about priority categories to be double-checked by other team members. The

team would assign color codes to patients, and once assigned, the patients would be randomized by the computer to indicate the order in which ventilators would be reassigned from currently ventilated patients and moved to new patients with higher-priority scores. Then, those decisions, too, could be verified by team members. At the URMC, a center was set up to house the triage teams, and dry runs of the procedures were conducted so that the team members would be thoroughly familiar with their jobs and how the process would work. More training and simulations would have been done had the surge not subsided.

A process for communicating between the triage teams and both the emergency department and the ICUs was established and tested so that the teams would know when a ventilator became available to assign to a new patient and when a ventilator was needed by a new patient. In the latter case, the team would inform the treating team to initiate the withdrawal of the ventilator and begin comfort measures for that patient and to get the ventilator to the new patient as quickly as feasible. Minutes could make a difference in the life of an incoming patient. As noted earlier, families must be prepared that their loved ones may have their ventilators removed with virtually no notice.

A key element, then, are the liaison officers who would match medical record numbers and ventilator identifiers to patients and order the transfers. Only those officers, working separately from the triage teams, would have access to concrete patient information, and they would have no authority to alter the decision of the triage teams.

6. MORAL DISTRESS

The decision to withdraw ventilator support from a patient is morally fraught. When it is withdrawn, the patient would be expected to die. One of the effects of the early chaotic days of the pandemic is that healthcare workers were forced to make such decisions without any guidance whatsoever, and they were thereby forced to bear the whole burden of those choices. Being able to implement a reasonable plan does much, then, to relieve moral distress. Yet it does not eliminate it. To be morally sound, then, any allocation plan must consider and ameliorate the harms that are the result of its application. Obviously, the pa-

tients who lose their ventilators are the most affected by the plan. Although they cannot be saved without creating a greater loss of life, they can be given the best comfort care possible consistent with their wishes for end-of-life care. And their families can be provided with grief counseling and other services. Hospitals typically have teams of providers who attempt to fill these needs—though in a pandemic crisis, those teams will be stretched thin, so medical centers should plan ways to supplement their staffs.

Allocation decisions will also affect the providers. Healthcare workers will have to withdraw ventilators from patients whom they thought might recover and then watch them die. Such an experience is bound to be heart-wrenching, and it might be traumatic,²⁷ and the trauma can be compounded when it has to be repeated.²⁸ The natural responses to being given this task are sadness, depression, anger, and frustration. Even acting as a member of the triage teams, as instruments of the decisions that would lead to deaths, could be traumatic. Indeed, even the theoretical possibility of such decisions disturbed some members of the CERG enough that they withdrew their participation, and for a similar reason, others did not choose to join any of the triage teams.

All institutions have a moral obligation to protect the well-being of their employees who suffer doing their jobs, and crisis situations create special challenges. Crucially, the most effective way to discharge this duty may be to prevent the circumstances that create moral injury in the first place. Removing the burden of triage decision-making from frontline clinical workers and replacing it with an anonymized system of dedicated triage panels may help alleviate (or at least share) the burden of moral decision-making. Even so, some distress is inevitable, and providers who suffer distress must be offered organized debriefing to help them process the complex emotions they experience. The CERG identified a team with experience in addressing moral distress to work with frontline workers to organize these efforts.

7. CONCLUSION

Fortunately, none of the procedures the CERG developed had to be implemented. Yet the COVID-19 pandemic put a spotlight to the ethical issues that arise in the allocation of scarce

resources like ventilators and the need to create ethically sensitive protocols for such allocations. But little attention has been paid to the numerous ethical issues that arise when a system tries to implement such a protocol. Our experiences at the URMC demonstrate that these implementation issues are not trivial, and they require considerable resources. Faced with a novel disease, considerable effort was required to translate principles into concrete plans. But more importantly, an effective allocation plan requires sophisticated computer programming so that the plan can be implemented smoothly and as accurately as possible. And resources must be available to deal with the emotional aftermath of the decisions that are made.

An ethical implementation thus requires medical centers to devote resources to this task, and some of those resources should be invested well before the need for resource triage becomes apparent. While future pandemics will have their unique features, many elements of the implementation plans can be created ahead of time. Many of these efforts can—and should—be made long before the next pandemic hits. Investing in translational science and creating a computer environment that can be used by triage teams is work that can be done even before the exact parameters of an allocation are known. But centers also need to plan enough that they have teams available to perform all of these tasks should the need arise. Importantly, medical centers also need to work with their communities to begin to address the inequalities in access to healthcare and in health itself that led to the differences in effects of COVID-19 on different socioeconomic and racial groups.²⁹

Of course, medical centers also need to create plans that will try to avoid the need for allocation protocols in the first place. Some resources should be spent to ensure that hospitals have adequate ventilators and personal protective equipment. Even more importantly, we as a society should be devoting resources to the public health infrastructure to control epidemics, both here and abroad, long before they pose the kind of threat that COVID-19 has posed. But that argument is for a different article.

NOTES

1. J. Horowitz, "Italy's Health Care System Groans under Coronavirus—a Warning to the World," *New York Times*, 12 March 2020, <https://www.nytimes.com/2020/03/12/world/europe/12italy-coronavirus>

-health-care.html; M. Rothfeld, S. Sengupta, J. Goldstein, and B.M. Rosenthal, "13 Deaths in a Day: An 'Apocalyptic' Coronavirus Surge at an N.Y.C. Hospital," *New York Times*, 25 March 2020, <https://www.nytimes.com/2020/03/25/nyregion/nyc-coronavirus-hospitals.html?searchResultPosition=4>.

2. See, e.g., E.J. Emanuel et al., "Fair Allocation of Medical Resources in the Time of COVID-19," *New England Journal of Medicine* 382 (2020): 2049–55; A.L. Maguire et al., "Ethical Challenges Arising in the COVID-19 Pandemic: An Overview from the Association of Bioethics Program Directors Task Force," *American Journal of Bioethics* 20, no. 7 (2020): 15–27; R. Truog, C. Mitchell, and G. Daley, "The Toughest Triage—Allocating Ventilators in a Pandemic," *New England Journal of Medicine* 382 (2020): 1973–75; and D. White and B. Lo, "A Framework for Rationing Ventilators and Critical Care Beds during the COVID-19 Pandemic," *JAMA* 323 (2020): 1773–4.

3. The importance of engaging with the community is emphasized by the crisis plans developed by New York, Maryland, and the University of California. See New York State (NYS) Task Force on Life and the Law, *Ventilator Allocation Guidelines* (2015), https://www.health.ny.gov/regulations/task_force/reports_publications/docs/ventilator_guidelines.pdf; Maryland ASR Framework Team, *Maryland Framework for the Allocation of Scarce Life-Sustaining Medical Resources in a Catastrophic Public Health Emergency* (2017), <http://bioethicstoday.org/wp-content/uploads/2020/03/Daugherty-Maryland-framework-PH-emergency-2017.pdf>; and A. Rajczi, J. Daar, A. Kheriaty, and C. Dastur, "The University of California Crisis Standards of Care: Public Reasoning for Socially Responsible Medicine," *Hastings Center Report* 51, no. 5 (2020): 30–41. Maryland places special emphasis on the importance of continuing engagement with the community. See V. Brown, "Centering Social Justice for Covid-19 Resources and Research," *Hastings Center Report* 51, no. 5 (2021): 51–3.

4. L. Gilson, "Trust and the Development of Health Care as a Social Institution," *Social Science and Medicine* 56 (2003): 1453–68.

5. See J.F. Childress et al., "Public Health Ethics: Mapping the Terrain," *Journal of Law, Medicine, and Ethics* 30 (2002): 170–8; and R. Upshur, "Principles for the Justification of Public Health Intervention," *Canadian Journal of Public Health* 93, no. 2 (2002): 101–3.

6. Indeed, an incident at the beginning of the pandemic exposed many of the fault lines of the intersection of healthcare, law enforcement, and violence experienced by its Black community. In March 2020—just as the CERG was beginning its work—Daniel Prude died in police custody after suffering a PCP-induced mental breakdown. Body camera video released in September 2020 showed that his asphyxiation was caused by the use of a spit hood and physical restraints used by the officers. See Office of the

Attorney General, Special Investigation and Prosecutions Unit, Report on the Investigation into the Death of Daniel Prude, 2021, https://ag.ny.gov/sites/default/files/oag_report_-_prude.pdf. For a physician's perspective on the case, see D.A. Paul, "The Death of Daniel Prude—Reflections of a Black Neurosurgeon," *New England Journal of Medicine* 383 (10 December 2020): e135.

7. NYS Task Force, *Ventilator Allocation Guidelines*, see note 3 above, sec. 2.

8. Emanuel et al., "Fair Allocation," see note 2 above; Maguire et al., "Ethical Challenges," see note 2 above; Truog, Mitchell, and Daly, "Toughest Triage," see note 2 above; White and Lo, "Framework for Rationing Ventilators," see note 2 above; S.D. Goold, "The COVID-19 Pandemic: Critical Care Allocated in Extremis," *American Journal of Bioethics* 20, no. 7 (2020): 6–8; and R. Rhodes, "Justice and Guidance for the COVID-19 Pandemic," *American Journal of Bioethics* 20, no. 7 (2020): 163–6.

9. See R. Truog, "Ventilator Allocation Protocols: Sophisticated Bioethics for an Unworkable Strategy," *Hastings Center Report* 51, no. 5 (2021): 56–7.

10. NYS Task Force, *Ventilator Allocation Guidelines*, see note 3 above.

11. See, e.g., G.M. Piscitello et al., "Variation in Ventilator Allocation Protocols by US States during the Coronavirus Disease 2019 Pandemic: A Systemic Review," *JAMA Network Open* 3, no. 6 (2020): e2012606; and M.L. Gaurke et al., "Life-Years and Rationing in the Covid-19 Pandemic: A Critical Analysis," *Hastings Center Report* 51, no. 5 (2021): 18–29.

12. NYS Task Force, *Ventilator Allocation Guidelines*, see note 3 above, p. 105.

13. *Ibid.*

14. *Ibid.*, 168.

15. *Ibid.*, 48.

16. We now have good reason to believe that SOFA scores do not do a good job predicting survival for COVID-19 patients or for ICU patients in general. Indeed, retrospective studies show that it does worse in predicting survival of COVID-19 patients than proposals that rely more on comorbidities or even than those that simply use age. See J.J. Jezmir et al., "Performance of Crisis Standards of Care Guidelines in a Cohort of Critically Ill COVID-19 Patients in the United States," *Cell Reports Medicine* 2 (2021): 100376; S.V. Bhavani et al., "Simulation of Ventilator Allocation in Critically Ill Patients with COVID-19," *American Journal of Respiratory and Critical Care Medicine* 204 (2021): 1224–7; and R.A. Raschke et al., "Discriminant Accuracy of the SOFA Score for Determining

Probable Mortality of Patients with COVID-19 Pneumonia Requiring Mechanical Ventilation," *JAMA* 325 (2021): 1469–70.

17. See, e.g., Bhavani et al., "Simulation of Ventilator Allocation," see note 16 above.

18. NYS Task Force, *Ventilator Allocation Guidelines*, see note 3 above, pp. 86 and 138.

19. *Ibid.*, 179–81.

20. H. Schmidt, "Vaccine Rationing and the Urgency of Social Justice in the Covid-19 Response," *Hastings Center Report* 50, no. 3 (2020): 46–9.

21. D.C. Ashana et al., "Equitably Allocating Resources during Crises: Racial Differences in Mortality Prediction Models," *American Journal of Respiratory and Critical Care Medicine* 204 (2021): 178–86.

22. At the advice of White and Lo, the ADI was incorporated into the early protocol of the University of Pittsburgh and defended in D. White and B. Lo, "Structural Inequities, Fair Opportunity, and the Allocation of Scarce ICU Resources," *Hastings Center Report* 51, no. 5 (2021): 42–7. See University of Pittsburgh School of Medicine, *Allocation of Scarce Critical Care Resources during a Public Health Emergency* (2020), https://ccm.pitt.edu/sites/default/files/UnivPittsburgh_ModelHospitalResourcePolicy_2020_04_15.pdf.

23. A.J.H. Kind and W.R. Buckingham, "Making Neighborhood-Disadvantage Metrics Accessible—the Neighborhood Atlas," *New England Journal of Medicine* 378 (2018): 2456–8.

24. J.K. Hick, D. Hafling, M.K. Wynia, and E. Toner, "Crisis Standards of Care and COVID-19: What Did We Learn? How Do We Ensure Equity? What Should We Do?," *National Academy of Medicine Perspectives*, 30 August 2021.

25. White and Lo, "Framework for Rationing Ventilators," see note 2 above.

26. A. Milliken et al., "Addressing Challenges Associated with Operationalizing Crisis Standards of Care Protocols for the Covid-19 Pandemic," *NEJM Catalyst Innovations in Care Delivery*, 12 August 2020.

27. R. Williams, J. Brundage, and E. Williams, "Moral Injury in Times of Crisis," *Journal of Health Service Psychology* 46 (2020): 65–9; C. Fourie, "Moral Distress and Moral Conflict in Clinical Ethics," *Bioethics* 29 (2015): 91–7.

28. E.G. Epstein and A.B. Hamric, "Moral Distress, Moral Residue, and the Crescendo Effect," *Journal of Clinical Ethics* 20 (2009): 330–42.

29. Researchers at the UPMC are currently using the data collected during the pandemic to investigate these questions.