

## CTO

In the frightening, early days of the COVID-19 pandemic, doctors across the world struggled with a myriad set of challenges. First, what would be the clinical characteristics of this new infection? I remember receiving and committing to memory a hand-written set of notes that reportedly came from an ICU doctor in Seattle, where an initial wave of cases in America had occurred. Passed from colleague to colleague in its nth degree of duplication, it was a basic cheat sheet of clinical information about the illness before any real medical publication was available to us. What drugs, if any, could we use to treat it? At that point we had only educated guesses. What would we do if we were overwhelmed by the feared, coming wave of critically-ill patients? Visions of ICU patients overflowing into hospital corridors, tangles of tubes and wires, squawking alarms and ventilators all filled our thoughts. Plans were rapidly made with expert guidance at all levels – health system wide, facility wide, departments down to individuals, all to assess this potential. And as plans were disseminated and reviewed, one specific question emerged as a crystallization of our fear in the face of this world tragedy: If and when the coming wave of patients had exhausted our finite reserve of mechanical ventilators, how would we decide who would get the last ones? If we had to remove the ventilator from one patient in order to help someone else live, would we? Could we?

The field of medical ethics had an extensive literature on the subject known as the “The ethical allocation of scarce medical resources.” Most well-known in this field, because of its publication in Life magazine, had been the infamous experience of the 1961 “Seattle God Committee.” At the time, while hemodialysis as a life-saving treatment for endstage renal disease had been developed, there was neither mass production of the dialysis machines and related equipment, nor approved funding for their wide-spread use. As a result, doctors then faced the dilemma of being able to treat just a handful of patients with the limited supply of prototype equipment while being unable to save the lives of the great majority. By what criteria would they decide who would receive this resource? The solution at the time was to create a multi-disciplinary, anonymous committee that would evaluate the cases and judge them according to several criteria. But most importantly, they decided to use judgment of a patient’s social worth as a key criterion. By this way of thinking, a father who supported a family of 4 through work as a corporate executive would have a higher degree of social worth than a single man working as a laborer. Assuming similar medical criteria otherwise, the executive would get dialysis and live and the laborer would die. The scariness of the prospects of judging social worth, especially as revealed to the American public in Life magazine by Shana Alexander, ultimately led to congress’ involvement in funding treatment of endstage renal disease through federal resources, what would come to be an individual Medicare benefit.

To be clear, the situation in 2019 was very different from 1961. Ventilators were not prototypes as dialysis machines once had been, and they numbered at least in the 100,000’s across the nation. Neither were they limited to one treatment center in Seattle, but available in every hospital with an ICU. But the prospect of running out of this finite resource was very real and

very scary to many of us – for in the end, it might well have come down to taking a machine away from one patient in order to give it to another.

Enshrined within the normative ethics of the doctor-patient relationship is the notion that the doctor is to be devoted to the care of the patient before him or her above almost everything else. His or her duty is to that patient, certainly much more than to other factors such as one's time or convenience or individual preferences. But even more importantly, in the general understanding of this relationship, the duty is from one individual doctor to one individual patient. In other words, the health of a patient down the hall being cared for by another doctor, or of the entire American public, though incredibly important, is not his/her primary duty. This is not to say that their health is not the primary duty of others, for it is. Resultingly, taking a ventilator from one's patient in order to give it to another doctor's patient would violate fundamental aspects of how we understand medical ethics. Consequently, in order to actually do this, one would need to invoke a paradigm shift – a reframing of fundamental ethics at the bedside. This shift can be accomplished by "Declaring crisis standards of care." Through such a declaration, it can be understood that a physician might have to act contrary to the health of one's own patient in order to serve the greater good. And because of the foreignness of such a notion in routine American medical ethics, the idea of such a declaration took on huge meaning in the Bioethics community. In our planning, these declarations could be made at the state level, the healthcare system, or the individual hospital.

Planning for the eventuality of a declaration of crisis standards ensued across California and the nation among healthcare administrators and bioethicists. As can be imagined, in the absence of specific laws about how to proceed, different institutions took different approaches. As can be understood, the utilitarian idea that a limited supply of ventilators should be used for the patients for whom the greatest good could be done emerged and was broadly accepted. Judging who the patients were for whom the greatest good could be done was what proved to be challenging. All recognized that the limited data thus far showed that the prognosis for survival was the worst for elderly patients, and for those with very significant medical comorbidities or disabilities like cancer, emphysema, or cirrhosis. Some facilities thus chose to use age as a prognostic factor in deciding who might live and who might die. Other facilities did not. Why? Those above me burdened with making these plans feared the repercussions from groups who might feel these to be unfair forms of medical discrimination. As a result, some institutions such as mine chose a completely objective set of criteria that were blind to age, ethnicity, religion, and name. In such, one could create a computer algorithm to score patients' severity of illness based purely on physiologic criteria (like oxygen level, lab values, vital signs, etc), figure out which patients had the worst prognosis, and shift resources from them to other patients. Such a scoring system could be done by a computer, evaluated by a small group, and decisions made. In each facility, this system would be overseen by someone known as the Crisis Triage Official (CTO.) Such was my fate to be asked to do this task at my hospital. In some facilities, the CTO alone would make decisions on how to use limited resources. In the plan for my healthcare system, I was to oversee five 3-person teams who would rotate doing this work in collaboration with me and our PhD Clinical Ethicist. Our teams gathered and drilled the dilemma as an army unit might plan an attack. We practiced using the algorithm on our current

ICU patients for the eventuality that we might need to actually use it imminently. In the darkest days of the surge at my facility, my colleagues and I reckoned once that we were within a day or two of a declaration of crisis standards.

But one problem entered my mind and would not leave. Running the numbers of the algorithm did not really require my skills as a clinician. Truth be told, it did not really require a human – a computer could do it all. As a matter of fact, the computer really **did** do it all, and our job was to cross check it and put its work into effect. The system was designed to relieve individual ICU doctors of the responsibility of having to remove life support from their own patients, a terrible dilemma for anyone to bear. The system shifted that responsibility onto our teams. Moreover, it then relieved our teams of the responsibility of making those decisions – a computer would do it. As a middle-aged physician, my career has been dedicated to using the techniques of generations of doctors before me towards the diagnosis and treatment of diseases. Those techniques fundamentally require the use of 4 of my senses – hearing, feeling, seeing, even smelling. To be sure, computers can greatly aide us in many of those tasks. I routinely rely on all kinds of computer processing to aide me in my work, but not to defer to them above me. In attempting to reconcile this discrepancy, I vowed to deviate from the protocol to examine any patients in question so that the decision to take a patient’s ventilator away was indeed personal to me, and predicated on all my skills as a physician – and that I could verify that any decision actually made clinical sense. In that way the computer would work for me and not vice versa.

But as the son of a Holocaust survivor, I bear an inheritance that both illuminates, and creates shadows in, my work as a physician. The lives of my patients stand before me in my guard, even as the lives of my exterminated ancestors bear silent witness to their own fate. The dead ones, like some Greek chorus, bid me protect the living ones – all who could, like them, die as victims. And though the sick Covid patients in the ICU were not **my** patients, I could not stand idly by. Some would argue that precisely because they were not my patients, I had no obligation to them. They said, “You are not their doctor – You are the CTO”. Nameless, faceless, they had been dehumanized for the purposes of the algorithm. To not assert my role in their illness was to abandon them at the most critical point. How could I possibly not use every faculty that I possessed in order to make such a momentous decision? How could I possibly assuage my conscience in the misdirected comfort that a computer had made the decision and not I? I know the ease with which millions of lives were taken through finely-tuned, automated processes in the Holocaust – the computer algorithms of their day - also designed to minimize individual guilt. No, if there were a patient who had to die in order to save someone else, I would own that terrible decision as CTO. The alternative – to surrender their lives to the authority of a machine – would be far more terrible.