OPENING THE DEBATE ON THE NEW SEPSIS DEFINITION

Defining Sepsis: A Case of Bounded Rationality and Fuzzy Thinking?

The Sepsis-3 Task Force recently offered a conceptual definition of sepsis as life-threatening organ dysfunction secondary to a dysregulated host response to infection (1). As with prior definitions, the task force acknowledged there is no “gold standard” by which to operationalize this definition. Nonetheless, they believed there was a need to offer some standard and, based on several domains of validity, reliability, and practicality, recommended sepsis be defined clinically as an increase of two or more Sequential Organ Failure Assessment (SOFA) points in the setting of infection. The task force reached its recommendations after vigorous debate; opinions were not unanimous, and the subsequent reaction in the broader community reflects many of the same points of contention. Debate is welcomed, but one disconnect is the certainty with which some views have been expressed, despite the widely acknowledged uncertainties plaguing sepsis.

Why We Care So Much About a Label for Sepsis

There is a strong, justifiable impetus to have discrete diagnostic criteria for sepsis. The clinician wants to assign a diagnosis to trigger specific treatments, discard alternate diagnoses, and inform patients and families; the hospital and payor want to know who to count in quality improvement initiatives; the trialist wants to know who to enroll in a study of a potential therapy; and the epidemiologist wants to know who to follow to track incidence and outcome. Each stakeholder needs a discrete label, “sepsis,” to solve her version of a classic logic statement: “IF (sepsis) {action x }, ELSE {action y }. ” This is a set theory problem: we need to generate a set called “sepsis” and a set called “not sepsis,” and place each patient “crisply” into his or her set.

It is easy to crisply assign “sepsis” to a patient presenting with fever, altered mental status, hypotension, and a chest X-ray showing lobar pneumonia. Similarly, we can assign “infected but NOT septic” to the patient who appears well except for dysuria and neutrophilia. But consider the trauma victim who is in the intensive care unit for 2 weeks with acute respiratory distress syndrome and acute renal failure and who now has moderate hypotension, borderline fever, and a report from her nurse that she is sluggish to rouse with increased purulent sputum on suctioning. Is she septic? We are not sure, and crisp assignment is not possible. When labels cannot be assigned with certainty, the logic statement cannot be solved without some accommodation.

A Satisficing Solution to a “High Stakes” Problem

Regardless of academic deliberations on uncertainty, something needs to be done. Patients will arrive tomorrow, and a lack of discrete labels (or the use of “wrong” labels) will engender fear that cases will be missed, and lives may be lost. Being forced to act despite imperfect information is a well-studied problem. The economist Herbert Simon received the Nobel Prize for his theory of bounded rationality, which states that optimal solutions are compromised by the complexity of the problem and our cognitive inability to solve it (2). Instead, so-called rational choices are bounded (limited) by behaviors that satisfice (a blend of satisfy and

suffice), either by “finding optimum solutions for a simplified world, or by finding satisfactory solutions for a more realistic world” (3). Sepsis appears to be no exception. For example, some of us decide we “know” who is septic, possibly via cognitive errors, such as representation bias, where we attribute too much weight to stereotypical cases when envisioning the problem (4). Others accept the uncertainty, but perhaps do not appreciate its nature, choosing solutions for the wrong problem.

When All We Have Is a Hammer

Medicine is very used to problem solving under uncertainty, and so we quickly reach for our “how to manage uncertainty” toolkit. These tools (e.g., sensitivity and specificity, receiver operating characteristics, or net reclassification indices) are based on probability theory. In other words, when told “we don’t know whether the patient is septic or not,” we convert the sentence to something similar: “What is the probability the patient has sepsis?” Framing uncertainty this way accommodates the fact that patients have not just two options (0 or 100% chance of being septic), but a variable chance from 0 to 100%. A crucial assumption is that the problem is exclusively one of imperfect information: with enough information, we can determine to which set the patient belongs. For example, to calculate sensitivity, one needs some mechanism to know a true positive (i.e., has sepsis), even if the diagnostic test or criterion under evaluation operates with imperfect knowledge regarding who is “truly” septic (and thus mislabels as septic some who are not). We then rank different diagnostic strategies based on how well they correctly label patients (i.e., weigh their probability) as belonging to one set or the other.

This is where we have a problem: the question is not just whether our patient is septic but also to what degree she is septic. In other words, the limit on crisp assignment to one set or another is not just that the probability of sepsis varies from 0 to 100%, but also that the degree to which she has sepsis varies from 0 to 100%. Probability and degree sound similar, but are not the same. Consider our trauma intensive care patient: we are unsure both about whether these new features are a result of sepsis (i.e., the probability that sepsis is present) and, even if blood cultures were to turn positive, whether this constellation of features is sepsis or just infection (i.e., the degree that sepsis is present). Unlike probability, more information does not necessarily reduce degree to two choices: even with perfect information, we could conclude the patient is “definitely somewhat” septic.

Because we are familiar with probability problems, we satisfice by reframing sepsis as a simpler problem than it is (crisp assignment limited only by imperfect information). Not only does this simplification result in potentially incorrect analyses, but it permeates the debate with the illusion that the solution we seek (a bit more information) is just beyond the horizon, emboldening advocates to feel more strongly for their position than perhaps they ought. Against this threat, I see two paths forward. The first borrows methods from social sciences to define sepsis despite these caveats.

The second borrows a trick from engineering to bypass defining sepsis altogether.

Measuring Validity in the Absence of a Gold Standard

The task force already took steps down this road. Assuming no gold standard, we adopt the rubric used in social sciences to define abstract concepts that lack gold standards, which requires considering multiple domains of validity. One domain by which the SOFA score was evaluated was predictive validity (5). Predictive validity is not prediction of sepsis, as sepsis cannot be known. Rather, it is prediction of events that we can measure and that are more common after sepsis. For example, because sepsis is life-threatening, we assume death is more common in sepsis than in other infected patients. And we can measure death. Criteria among infected patients that identify patients who die more frequently define a set of infected patients who, in turn, are more likely to be septic. Two things are worth noting. First, because measures here such as sensitivity and specificity reflect prediction of outcomes associated with sepsis, and not sepsis itself, we should not expect perfect scores. To do so is to conclude that every sepsis case is lethal, and that no one who is infected but not septic can ever die. Second, this is one of multiple validity domains. Different diagnostic criteria may fare differently across domains, and different stakeholders will prefer different criteria, depending on the domains in which they excel (6).

Some Fuzzy Thinking

Tests based on probability theory quantify ignorance. However, the problem with sepsis is not just ignorance, but vagueness (how septic is she?). We can probe vagueness with fuzzy logic. Fuzzy logic was developed as an extension of control theory to address instances in engineering when classic logic-based rules (such as those faced by our stakeholders) could not be executed because set membership could not be assigned crisply (7). For example, a rule requiring separate actions for hot versus warm water needs criteria to divide hot from warm water. But some water feels simultaneously somewhat warm and somewhat hot; where is the break-point? Fuzzy logic allows fuzzy rather than crisp set membership rules (e.g., a range of intermediate temperatures), where the degree of membership in each set can be graded. In turn, the rule can be evaluated across the fuzzy range, and in some instances, the range can be narrowed (de-fuzzified), depending on rule performance. The key is that fuzzy logic bypasses the requirement to first define whether a given temperature is hot or warm, or whether a person is septic or not.

The term fuzzy logic is not common in the sepsis field. But, whether by intent or accident, the case of procalcitonin is illustrative. Many researchers tested whether procalcitonin is a good diagnostic test for infection or sepsis, all using probability-based tools (8). But others stopped asking that question. Instead, they tested whether an antibiotic-prescribing rule based on procalcitonin was better than one without procalcitonin (9). The rule stated simply, and fuzzily, that higher procalcitonin values might be more common in patients with bacterial infection, and vice versa. Armed with that information, physicians could choose

to withhold antibiotics for patients who might have looked infected, but in whom procalcitonin was low. Several, but not all, of these studies suggest the procalcitonin-based rule yields superior outcomes, but that is not the point. The point is that no one knows who was actually infected, and they did not need to. Rather, the approach implemented a rule to manage infection despite vagueness regarding the degree to which each individual could be declared infected or not.

Moving Forward with Intentional Compromise

It is neither realistic nor necessary to expect a perfectly discrete label for a condition bounded by uncertainty. The two paths I discussed both require a form of surrender, in that they stop claiming sepsis is currently “knowable,” and hence can be counted, with certainty. But they also free us to make progress with what can be measured. Our challenge, which no doubt begins with some sober introspection, is to fully embrace, and quantify, the uncertainty. ■

Author disclosures are available with the text of this article at www.atsjournals.org.

Acknowledgment: The Sepsis-3 Task Force was sponsored by the European Society of Intensive Care Medicine and the Society of Critical Care Medicine, and its report was endorsed by a number of professional societies, including the American Thoracic Society. The author was a voting member of the task force but received no financial support for his participation.

Derek C. Angus, M.D., M.P.H.
Department of Critical Care Medicine
University of Pittsburgh School of Medicine
Pittsburgh, Pennsylvania

References

1. Singer M, Deutschman CS, Seymour CW, Shankar-Hari M, Annane D, Bauer M, Bellomo R, Bernard GR, Chiche JD, Coopersmith CM, *et al*. The third international consensus definitions for sepsis and septic shock (sepsis-3). *JAMA* 2016;315:801–810.
2. Simon HA. A behavioral model of rational choice. *Q J Econ* 1955;69:99–118.
3. Simon HA. Rational decision making in business organizations. *Am Econ Rev* 1979;69:493–513.
4. Kahneman D, Tversky A. Subjective probability: a judgment of representativeness. *Cognit Psychol* 1972;3:430–454.
5. Seymour CW, Liu VX, Iwashyna TJ, Brunkhorst FM, Rea TD, Scherag A, Rubenfeld G, Kahn JM, Shankar-Hari M, Singer M, *et al*. Assessment of clinical criteria for sepsis: for the third international consensus definitions for sepsis and septic shock (sepsis-3). *JAMA* 2016;315:762–774.
6. Angus DC, Seymour CW, Coopersmith CM, Deutschman CS, Klompas M, Levy MM, Martin GS, Osborn TM, Rhee C, Watson RS. A framework for the development and interpretation of different sepsis definitions and clinical criteria. *Crit Care Med* 2016;44:e113–e121.
7. Zadeh LA. Fuzzy sets. *Inf Control* 1965;8:338–353.
8. Simon L, Gauvin F, Amre DK, Saint-Louis P, Lacroix J. Serum procalcitonin and C-reactive protein levels as markers of bacterial infection: a systematic review and meta-analysis. *Clin Infect Dis* 2004;39:206–217.
9. Christ-Crain M, Jaccard-Stolz D, Bingisser R, Gencay MM, Huber PR, Tamm M, Müller B. Effect of procalcitonin-guided treatment on antibiotic use and outcome in lower respiratory tract infections: cluster-randomised, single-blinded intervention trial. *Lancet* 2004;363:600–607.

Copyright © 2016 by the American Thoracic Society