

Target Article

Complete Lives in the Balance

Samuel J. Kerstein, University of Maryland, College Park
Greg Bognar, New York University

The allocation of scarce health care resources such as flu treatment or organs for transplant presents stark problems of distributive justice. Persad, Wertheimer, and Emanuel have recently proposed a novel system for such allocation. Their “complete lives system” incorporates several principles, including ones that prescribe saving the most lives, preserving the most life-years, and giving priority to persons between 15 and 40 years old. This paper argues that the system lacks adequate moral foundations. Persad and colleagues’ defense of giving priority to those between 15 and 40 leaves them open to the charge that they discriminate unfairly against children. Second, the paper contends that the complete lives system fails to provide meaningful practical guidance in central cases, since it contains no method for balancing its principles when they conflict. Finally, the paper proposes a new method for balancing principles of saving the most lives and maximizing life-years.

Keywords: age discrimination; flu pandemic; organ transplant; rationing; resource allocation

The allocation of scarce health care resources presents stark problems of distributive justice (Kamm 1993; Daniels 2008; Nord 1999). When we must decide how to allocate beds in an intensive care unit, vaccinations during a flu pandemic, or organs for transplant, our choice can determine who lives and who dies. But there is little agreement on principles for such allocation. In life-and-death choices both familiar moral principles and intuitively plausible moral judgments quickly lead to conclusions that seem unacceptable. We are left with a number of competing principles and a motley of intuitions that are often inconsistent with one another.

It is important, therefore, to develop a coherent framework that can help us make these choices by combining the most plausible principles. Any acceptable “system” of principles must satisfy at least the following two conditions: First, its component principles must rest on secure moral foundations, and second, it must be able to provide practical guidance—especially in cases when different principles conflict and must be balanced. These conditions are necessary (but not sufficient) for any resource allocation scheme to be *legitimate*. A legitimate scheme can be seen as just and fair by those who are potentially affected by its decisions—that is, all of us—and it is capable of enjoying broad public support (Daniels 2008; Persad et al. 2009a).

Recently, an allocation system has been proposed by Govind Persad, Alan Wertheimer, and Ezekiel J. Emanuel. They call it the *complete lives system*.¹ In this paper, we argue that the complete lives system fails to satisfy both of the conditions just given: Some of its main component principles

lack adequate moral foundations, and it fails to provide meaningful guidance in a range of central cases. Having raised our objections, we discuss in more detail a problem associated with the second condition, namely, that of how to balance conflicting principles for the allocation of scarce life-saving health care resources. Finally, we propose a way of making progress toward resolving this vexing problem.

THE COMPLETE LIVES SYSTEM

The complete lives system aims to serve as a basis for just allocation of continually scarce life-saving interventions (429).² Examples of such interventions are organ transplants and vaccinations against new and deadly forms of flu. The system includes five principles: “youngest-first, prognosis, save the most lives, lottery, and instrumental value” (428).

Youngest-first directs us to give priority to younger people. Persad and colleagues, however, do not adopt a strict youngest-first principle, which would give priority for life-saving interventions to infants over older children and adults. Rather, they embrace a principle that prioritizes adolescents and young adults, that is, persons roughly between 15 and 40 years old, over infants and older adults (428). We shall call their principle “modified youngest-first.”

“Prognosis” and “save the most lives” are straightforward. According to prognosis, we ought to “save the most life-years” (425). This principle would give priority for a liver transplant to a patient who would live an additional 20 years over a patient who would live an additional 5 years

Address correspondence to Samuel J. Kerstein, University of Maryland, Department of Philosophy, College Park, MD 20742, USA. E-mail: kerstein@umd.edu

1. Persad and colleagues 2009a. All otherwise unattributed page references are to this paper.

2. Some bioethicists distinguish between “commodity scarcity,” that is, when some items, such as organs for transplant, are in limited supply, and “fiscal scarcity,” that is, when the items are available, but not used because of the expense (Morreim 1995, 47–51). Persad and colleagues do not invoke this distinction, but they appear to focus on developing a system for cases of commodity scarcity (423). In addition to issues concerning the just distribution of goods, fiscal scarcity raises ethical questions regarding how to interpret a physician’s duty to promote his or her patient’s best interests, as Morreim (1995) points out.

with the transplant. Of course, in some cases a principle of saving the most life-years yields different prescriptions than a principle of saving the most lives. We discuss this point later.

Persad and colleagues give limited roles to lottery and instrumental value principles in their system. They suggest that lotteries might be used to choose between “roughly equal” candidates for a life-saving intervention (428). Suppose, for example, that there is only one liver available for a transplant but two patients, each of whom, if he received it, would likely have his life extended for 10 years. If the difference between the patients relative to the other principles in the system is limited to one patient’s being to some small extent favored by the modified youngest-first principle, then Persad and colleagues suggest that it might be legitimate to determine by lottery who gets the organ.³

According to Persad and colleagues, “instrumental value allocation prioritizes specific individuals to enable or encourage future usefulness” (426). For example, this principle might imply that during a flu pandemic, medical staff necessary for the distribution of vaccine should be vaccinated before others. The complete lives system limits the application of the instrumental value principle to “some public health emergencies” (424, Table 1; 429).

How do Persad and colleagues derive the component principles of the system? They begin by presenting four categories of principles: those that promote or reward social usefulness, those that aim to treat people equally, prioritarian principles (which give priority to the worst off) and utilitarian principles (which aim to maximize benefits). They discuss two principles in each of these categories. They argue that some of these principles are *inherently flawed*, that is, “necessarily recognize” some morally irrelevant consideration (423). Other principles are *practically flawed*, that is, allow in practice some morally irrelevant consideration to affect allocation choices. Both sorts of principles should be rejected. In addition, there are *insufficient* principles—those that are based on some morally relevant consideration, but ignore others. Insufficient principles can be part of a “multi-principle” system, since the *joint* application of such principles can account for all morally relevant considerations.

In the rest of this section, we summarize and evaluate some of the arguments that Persad and colleagues give for classifying principles in each category as insufficient or inherently or practically flawed.

The category of promoting or rewarding social usefulness includes the principle of instrumental value and the principle of reciprocity. As we have seen, the former would give priority to those who are instrumental in promoting

3. They also say that lotteries might be used “to ensure that no individual—irrespective of age or prognosis—is seen as beyond saving” (428). So they are apparently open to the legitimacy, say, of holding a lottery for a single life-saving organ between a 70-year-old who, with the organ, stands to live an additional 5 years and a 20-year-old who, with it, stands to live an additional 30 years—namely, a lottery in which the 70-year-old would have only a very slight chance of receiving the organ.

some value, such as that of lives being saved. The latter principle would reward past promotion of such a value by, for example, giving priority to organ donors or research subjects. While the complete lives system incorporates instrumental value in limited contexts, it excludes reciprocity, at least partly on the grounds that it would be difficult and intrusive to ascertain who was worthy of being rewarded for past service (426).

Those principles that aim to treat people equally include *allocation by lottery* and the rule of *first come, first served*. Persad and colleagues argue that while the former is merely insufficient, and hence may be part of an acceptable allocation system, the latter should be rejected. First come, first served is practically flawed, they say, since even though it aims to treat everyone equally, it can be exploited by the wealthy, powerful, and well connected. Because of this, it allows morally irrelevant considerations to distort allocation.

Prioritarian principles include *sickest-first*—giving priority to those whose immediate future prospects are worse—and *youngest-first*—giving priority to younger people. Persad and colleagues argue that giving priority to the sickest is inherently flawed, since in true scarcity, who happens to be in the worst medical condition at the time is a morally arbitrary factor. It is unjust to expend resources on an acutely ill person at the expense of someone whose condition is less bad but will progressively get worse. They claim that what matters from the perspective of distributive justice is how well or badly one’s life goes as a whole, and not how one fares at one time (see also Nagel 1979, 120).

But this argument will be unconvincing to many. Perhaps the just distribution of some goods requires that we consider whole lives. For instance, it might not be unjust that people have less income at the beginning of their adult lives than at the height of their careers. Faring relatively worse at some periods of life may be adequately compensated by being better off at other times. But health does not seem to be the same sort of good. One thing illness might cause is pain and suffering. Many people would argue that the alleviation of severe, debilitating pain has special moral urgency. It would be invidious if a health care system did not treat you when you are in pain because you are considered to be “too well off overall” to have an urgent medical need, or if your health care professional argued that the pain you are experiencing now is “compensated” by full health at other times in your life. Whether just health care resource allocation should be concerned with how you fare at a time or throughout your whole life is, in our view, a much more complicated problem than Persad and colleagues suggest. It is precipitate to dismiss the principle of sickest-first as “inherently flawed.”⁴

But perhaps even more worrying are some of the implications of the argument for the doctor–patient relationship.

4. On the issue of who should be considered the worst off for the purposes of health care resource allocation, see Brock (2002). Note also that some of the points we make here can be reformulated to defend the principle of *first come, first served*.

For if the principle of sickest-first is inherently flawed, then apparently doctors need not care about who has the greatest medical need at the time, since, as the architects of the complete lives system argue, who is worst off at a time is a morally irrelevant factor. It hardly needs spelling out why most people would find the implications of this view unacceptable.⁵

Of course, Persad and colleagues would argue that expending resources on those who are sickest now will deprive those who might be even worse off in the future. This is true. But all this shows is that other, competing considerations are also relevant in allocating scarce health care resources across different times—it does not show that medical need at a time is not a morally relevant factor.

As we mentioned, in addition to sickest-first, Persad and colleagues categorize youngest-first as a prioritarian principle. The complete lives system incorporates a modified version of youngest-first, whose justification we discuss at length in the next section.

The final category of principles includes those that aim to maximize benefits. Persad and colleagues call these *utilitarian* principles. Their examples include a principle that directs us to maximize the number of lives saved and one that directs us to maximize the number of life-years saved. While these principles rest on morally relevant factors, they are insufficient, according to Persad and colleagues. By focusing on maximizing the quantity of some benefit, they ignore its distribution. For instance, by maximizing the number of lives saved, you ignore how long one has lived: The principle would tell you to be indifferent between saving a 70-year-old and a 20-year-old person. Similarly, the principle to maximize life-years saved would leave you indifferent between saving the 70-year-old and the 20-year-old if you can extend their lives only for 1 year.

What is peculiar about Persad and colleagues' discussion of the "utilitarian" principles is that these are not standard utilitarian principles. When utilitarians argue for maximizing benefits, they usually have in mind maximizing *utility*, where utility is a measure of well-being or quality of life. Utilitarians would argue that treating all people as having the same utility (as in the principle of maximizing the number of lives saved) or treating all years of life as having the same utility (as in the principle of maximizing life-years saved) are at best very imprecise approximations of their principle and at worst ignore what they really care about.

Of course, the measurement of utility is notoriously difficult, and there is little agreement on what well-being consists of. Nevertheless, health economists have developed sophisticated measures of health utilities, and these mea-

5. Persad and colleagues might point out that the complete lives system is intended to apply only to contexts of persistent scarcity when lives are at stake. But the public—which must be able to regard the allocation scheme as legitimate—might not distinguish between these contexts and others. Ignoring the present suffering of patients is likely to be hard for the public to accept *especially* in life-and-death cases, regardless of scarcity.

asures have been proposed as part of a utilitarian framework for the allocation of health care resources (e.g., McKie et al. 1998). One well-known example is the quality-adjusted life year, or QALY. Persad and colleagues deny a place to QALYs in health care resource allocation by arguing that "people, not QALYs, deserve equal treatment" (428). But of course this does not establish that the principles of quality of life maximization (whether in terms of QALYs, utilities, or some other measure) should have no place in allocation; all the argument shows is that this principle is not the only relevant one.

Indeed, quality-of-life considerations are conspicuously absent from Persad and colleagues' discussion. We can imagine, for example, that, relative to the principles in the complete lives system, two people are equal candidates for life-saving flu treatment, except that one person stands to live an additional 3 years, while the other would live only 1. According to the complete lives system, we ought to give the treatment to the one who would live longer. We ought to do so even if he would spend all of his additional life unaware of his surroundings, while the person not saved would have had a year of active engagement with his loved ones and his projects. Perhaps Persad and colleagues do believe after all that quality of life is morally irrelevant in the allocation of scarce health care resources. But this view would at least need an argument.

COMPLETE LIVES AND MODIFIED YOUNGEST-FIRST

The aim of the complete lives system is to promote complete human lives. In the allocation of scarce life-saving health care resources, we should enable people to live such lives, contend the system's developers. In this section, we raise some problems for this idea. In the next section, we show that the system fails to provide meaningful guidance in a whole range of central cases.

The notion of a "complete life" is central to Persad and colleagues' proposal. It is unfortunate, therefore, that they never tell us precisely what they mean by it. What they do say is compatible with different and mutually exclusive interpretations.

Consider a related idea, formulated by some philosophers, that focuses on the concept of a life plan. People construct and revise their overarching plans and projects for their lives in the light of how long they expect to live and how much time they expect they will need to carry out their plans. In some views, a system for allocating scarce resources should aim to provide the opportunity to complete life plans. Justice in health care requires equalizing such opportunity.⁶

This is not an unattractive idea, but it does not seem to be what Persad and colleagues have in mind. For them, a complete life seems to consist in a given number of life years—which might vary depending on the typical life span in a given society (429)—rather than in having the

6. For an account that develops this idea, see Daniels (1988; 2008). See also Rawls (1971).

opportunity to carry out a life plan.⁷ Moreover, they argue that “youngest-first allocation directs resources to those who have had less of something supremely valuable—life-years” (425). This suggests that they hold the view that life-years have intrinsic value, independently of what opportunities they provide and what level of well-being they enable people to achieve. Having more life-years is valuable even if life provides few opportunities and contains very little well-being.

But perhaps Persad and colleagues treat life-years merely as a proxy for well-being, rather than as valuable in themselves. (Of course, this immediately raises the question whether age is an appropriate proxy for well-being in general or in the sorts of applications they have in mind.) They claim that “the complete lives system justifies preference to younger people because of priority to the worst-off” (429). In the proxy view, younger people are worse off by virtue of having had less well-being than older people. Nevertheless, it is hard to be sure whether this is their view, since, as we argued earlier, they give very little role to well-being and quality of life in their discussion of the component principles of the system. It is not clear, therefore, what exactly in their view justifies giving priority to the younger over the older.

Some might object that such priority amounts to unfair discrimination against older people, regardless of whether life-years are an indication of well-being or valuable in themselves. However, Persad and colleagues argue that the discrimination that results is *not unfair*. They offer the following justification for this claim. Letting age determine who should get priority does not discriminate between people in the same way as prioritizing them by race or sex would do. Because everyone ages, age discrimination does not violate the requirement of treating people as equals. While discrimination by race and sex allocates burdens and benefits among different lives, discrimination by age allocates burdens and benefits among different life-stages. Even though some people will benefit and others will not, everyone can potentially receive the benefits. So this sort of discrimination is not unfair.⁸

There is, however, a complication. Persad and colleagues argue that the very young should not get the same priority as older children, adolescents, or young adults. In their view, the probability of receiving an intervention should gradually rise from a low base after birth until early adulthood, diminish until late middle age, and then begin to drop more steeply (see Figure 1). Earlier in this paper we

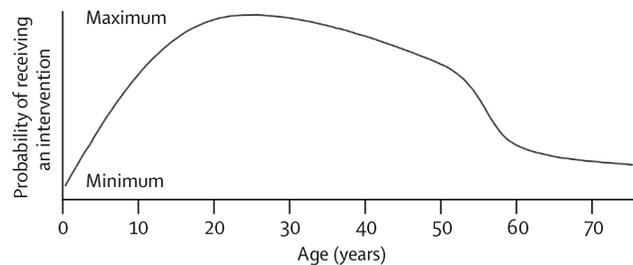


Figure 1. Age-based priority for receiving scarce medical interventions under the complete lives system. Source: Persad and colleagues (2009a), 428.

called this weighting scheme the modified youngest-first principle.⁹

In order to defend this principle, some explanation is needed why the very youngest are not given priority.¹⁰ Persad and colleagues borrow an argument that was put forward by Ronald Dworkin (1993). Dworkin argues that the death of an older child or a young adult is more tragic than the death of a young child or an infant because of the “investment” that has been made in the older person. In Persad and colleagues’ interpretation, this involves education and parental care that would be “wasted” if the young person were denied a complete life.

Thorough evaluation of this argument is beyond the scope of this paper. But the argument is problematic in the context of the complete lives system. For one thing, it seems arbitrary to think of “investment” in a person as limited to formal education and parental care. Why, for example,

9. There is also a difficulty with interpreting what makes a life “complete.” The term suggests that a life can be either complete or fall short of that. But this weighting scheme suggests that what we have is a range: Lives can be more or less complete depending on how much one has lived already. In this interpretation, the completeness of a person’s life may depend on how much she has lived *in comparison to* others: It may be (more) complete when compared to some people, and incomplete when compared to others.

10. In an attempt to support the modified youngest-first principle, Persad and colleagues cite empirical evidence of people’s preferences regarding whom to save. But it is not clear that this evidence favors modified youngest-first over youngest-first. For example, they claim (428) that empirical research by Tsuchiya, Dolan, and Shaw (2003) bolsters modified youngest-first. But this claim is questionable, to say the least. Tsuchiya and colleagues presented to subjects a scenario in which people of five different ages (5, 20, 35, 55, and 70 years old) will die in a few days without treatment. They asked the subjects to assume that each person, if treated, would go on to live a normal life span. They then gave the subjects the task of ranking the five age groups in terms of the order in which they would give them treatment. Seventy-six percent of the subjects ranked age 5 first, 13% ranked age 20 first, 10% age 35 first, 0% ranked 55 first, and 1% ranked age 70 first (Tsuchiya et al. 2003, Table 4, 693). These results seem far better suited to support youngest-first than modified youngest-first. The latter would, of course, give priority to 20-year-olds over 5-year-olds. For further discussion of the ambiguity of empirical evidence regarding age preferences, see Bogner 2008.

7. At one point, Persad and colleagues do say that the fulfillment of long-term plans requires a complete life (428). But they do not seem to endorse the notion that a complete life for a person *consists in* having the opportunity to fulfill a life plan.

8. They borrow this argument from Daniels (1988). Giving priority to younger people might, however, create further inequalities between the sexes. Since women tend to live longer, rationing by age would disadvantage them disproportionately. On these issues, see Jecker (1991).

should not the experience and training that a diplomat or business leader gets on the job also count as societal investment in her? If the greater degree of societal investment in a 20-year-old over an infant gives us reason to prioritize saving the 20-year-old, then it seems that the greater degree of societal investment in a 40-year-old over a 20-year-old would give us reason to prioritize saving the 40-year-old. But in the complete lives system, the 20-year-old would have priority (see Figure 1). To this objection, Persad and colleagues might reply that more investment would be “wasted” in the case of the 20-year-old, since the 40-year-old has already given back to society. But this reply is unconvincing. The societal investment in the 40-year-old (e.g., a surgeon) might be much larger than that in the 20-year-old (e.g., a student), and, as a result of a lengthy training period, she might not yet have had much occasion to produce returns.

In any case, the prioritarian argument Persad and colleagues invoke for the youngest-first principle undermines the *modified* youngest-first principle. As we have noted, one way that they justify giving priority to the younger is on the grounds that they are worse off than the older in terms of years lived. An infant is obviously worse off than an adolescent in these terms. So if benefits ought to go to the worse off, then they should go to the infant. But according to modified youngest-first, of course, it is the adolescent who should get priority.

Persad and colleagues face yet another problem in embracing the modified youngest-first principle. This principle is incompatible with the argument that they use to defend age discrimination. In order to see why, consider an old person who is denied some life-saving intervention. She cannot argue that she is being treated unfairly, since as a young person she enjoyed (or would have enjoyed if she had been in need) the benefits of an arrangement that gives priority to the young. In other words, she has no legitimate complaint that she is denied a life-saving resource. Consider now a very young child who, according to the complete lives system, is denied a life-saving intervention because priority is given to older children and young adults. It seems that she does have a legitimate complaint (one that someone can advance on her behalf): After all, she has not benefited from an arrangement that gives priority to young adults. In fact, she is being denied a life-saving resource for the sake of those who have had more of that “supremely valuable” thing—life-years. She neither has enjoyed nor ever will enjoy the benefits of the arrangement. She is not even potentially compensated.

In sum, the prioritarian view that younger people ought to get priority on the grounds of being worse off and the modified youngest-first principle undermine one another. Moreover, the argument that age discrimination is not unfair since everyone can expect to live through the same ages is unavailable for those who accept the modified youngest-first principle.

The moral foundations of the complete lives system are much less secure than they might initially seem. But perhaps it is nevertheless a workable approximation to a morally

sound procedure of the sort needed to solve urgent practical problems. In the next section, we consider whether the practical guidance it provides can compensate for its theoretical shortcomings.

DOES THE COMPLETE LIVES SYSTEM PROVIDE PRACTICAL GUIDANCE?

Persad and colleagues call the complete lives system a “coherent multiprinciple framework” that “has been developed to justly allocate persistently scarce life-saving interventions” (429). However, as we show next, it fails to help us in reaching an allocation decision in a variety of instances. To be sure, as its developers point out, the complete lives system is not an algorithm. We do not consider this to be a shortcoming, for we doubt that anyone can pinpoint a set of clear-cut steps that, if carried out in a specified order, will always yield a just allocation. Just allocation sometimes requires painstaking weighing of competing principles and context-sensitive judgment. What we do consider to be a shortcoming, however, is that in a wide range of cases—indeed, just the sort of cases in which an allocation system should give us guidance—the complete lives system proves unhelpful.

We illustrate this point with the help of examples that are streamlined in the service of brevity and clarity. We assume in them that we have greater certainty than we would in fact have regarding patients’ prognosis. Moreover, we suppose that each patient has the same social usefulness. The principle of instrumental value thus fails to have practical implications in the examples. We also assume that society does not owe any of the patients any less than it otherwise would as a result, for example, of their bearing moral responsibility for their medical condition, or any more than it otherwise would as a result, for example, of their having performed some great service to society in the past. Finally, for the reasons discussed earlier, we assume that if saved, each patient would have a high and roughly equal quality of life.

Consider first a case in which we have three 18-year-old patients who will soon die unless they receive transplants. We have one heart and one set of lungs available. If we give the whole heart/lung combination to the first patient, she will live until 70—which, we shall assume, is sufficient for a complete life. If, in contrast, we give the heart to the second and the lungs to the third patient, they will live for 2 years each. Prognosis prescribes that we give the heart/lung combination to the first patient, for that is the way to maximize life-years. However, to act in accordance with the principle of saving the most lives we would obviously have to give the heart to the second and the lungs to the third. Here we have a conflict between prognosis and maximizing the number of lives saved: Should we sacrifice two lives for the sake of a complete one? The complete lives system offers no guidance for how to proceed (see Gandjour 2009).

Next imagine that, through a multiple transplant, we can either save one 20-year-old for 4 years or two

55-year-olds for 2 years each. Since either way we preserve the same number of life-years, prognosis does not tip the scale in favor of saving the one or saving the two. The modified youngest-first principle favors saving the 20-year-old, for she is worse off in terms of the extent to which she has lived a complete life. However, the principle of saving the most lives would obviously imply that we should save the two 55-year-olds. The complete lives system leaves us with no clear idea of what we are required to do.

Finally, suppose that we are at an outpost in the midst of a flu pandemic and we have only enough medicine to treat either a 20-year-old who will then live for 5 years or an infant who will then live for 80 years. The complete lives system gives us no help in determining how to distribute the medicine justly. Prognosis requires giving it to the infant, while modified youngest-first demands that we give it to the 20-year-old.¹¹

As we mentioned, Persad and colleagues suggest that in cases where there are “roughly equal” candidates for life-saving interventions, it is legitimate to conduct a lottery. But within our cases do we have such candidates? For example, is the 18-year-old who, if given a heart and lungs, will live for another 52 years roughly equal to the two other 18-year-olds who, if given organs, will live for another 2 years each? We are unsure how Persad and colleagues would answer that question. But even if their view is that the candidates within our cases are roughly equal, it is not clear what sort of lottery we should conduct. In the case of the 18-year-old patients, should each one get a 50% chance of being saved, or should the one who needs a heart/lungs combination receive a 1/3 chance while each of the others gets a 2/3 chance?

These cases are simplified, but ones like them are likely to confront us in practice. Threats of avian and swine flu pandemics remind us that priorities need to be set regarding the distribution of scarce medicine. We must decide whether we will prioritize the treatment of those between 15 and 40 years over the treatment of infants or vice versa, for example. The dosage of medicine (e.g., Tamiflu) necessary to treat an infant can be less than one-half of what is necessary to treat an adult (Centers for Disease Control and Prevention 2009). Now suppose we make the simplifying assumption that the medicine is equally effective in infants and in adults in preventing flu deaths and that the flu poses equal mortality risks to these age groups. The complete lives system leaves us without guidance in this situation. Prognosis and maximizing the number of lives saved would favor prioritizing the infants. We could save twice as many of them and secure many more life-years if we did so. However, modi-

11. The numbers, of course, are merely illustrations. Depending on the precise weights assigned to different ages on the priority curve (Figure 1), the larger benefit to the infant may outweigh the increased chance of the 20-year-old of receiving the intervention. Evidently, however, there are always going to be cases in which modified youngest-first remains in conflict with the other principles, whatever the precise weights are.

fied youngest-first would have us prioritize people between 15 and 40.

The architects of the complete lives system might embrace the indeterminacy of their view regarding these examples and ones like them. They might maintain that as long as all applicable principles included in the system are taken into account and somehow balanced in reaching an allocation decision in a particular case, that decision will be just.¹² But this response is not plausible. For, according to it, the system would, for example, be consistent with giving prognosis 5 or even 10 times the weight of maximizing the number of lives saved. Yet it seems clearly unjust to allocate a scarce life-saving resource to one person who, with its help, would live an additional 82 years rather than to 8 persons who would live an additional 10 years each.

Another response to the examples open to Persad and colleagues is to appeal to the idea that if, in a particular case, an allocation is favored by a majority of the applicable principles included in the complete lives system, then that allocation is just. But they would presumably not want to appeal to this idea. For it would imply that we should save two people who have already had complete lives—two 95-year-olds, say—for 5 additional years each, rather than save one 20-year-old for 9 years.

Note that our criticism of the complete lives system’s practical effectiveness is not predicated on the notion that in every case a plausible system must entail that one particular allocation alone would be just. We grant that context-sensitive employment of such a system’s principles might sometimes yield a range of permissible allocations. It might lead us to conclude, say, that it would be acceptable either to give the last remaining intensive care bed to A or to have a lottery in order to decide between A and B.

But the complete lives system includes principles that prescribe divergent courses of action in a variety of cases. Each one of three principles—maximizing the number of lives saved, prognosis, and modified youngest-first—sometimes conflicts with the other two. When and how lottery might be used is also left unspecified. As the case of pandemic flu planning illustrates, these conflicts occur in just the sort of allocation scenarios in which we most need guidance. But the complete lives system fails to provide it.

ALLOCATION SYSTEMS AND BALANCING

The practical ineffectiveness of Persad and colleagues’ proposal leaves us with an important lesson. In order to develop a just system for the distribution of persistently scarce, life-saving resources, we need to undertake the arduous task of specifying how to balance allocation principles when they yield conflicting prescriptions. Of course, we also need to determine which principles should figure into allocation decisions in the first place. As our criticisms of the foundations of the complete lives system suggest, we doubt whether there is sufficient warrant to include a principle of modified youngest-first. Indeed, although we cannot

12. Persad and colleagues suggest this sort of response in 2009b.

discuss our reasons here, we are also skeptical whether we should include any principle that itself demands priority for the youngest (Bognar and Kerstein 2010).

However, we do agree that a system for just allocation must balance some principle akin to prognosis with some principle akin to maximizing the number of lives saved. A principle akin to prognosis that we defend elsewhere prescribes extending the lives of persons: beings who have certain psychological capacities, including the capacities to set ends and to form, act on, and revise plans for attaining them (Bognar and Kerstein 2010). Another principle might prescribe extending life, but only when its quality is above a certain threshold. These principles do not demand that we use scarce resources to prolong lives regardless of their quality. But we shall put aside considerations regarding the precise shape that such principles should take and consider how we might balance between the defeasible imperatives to save as many lives as we can and to extend life as much as possible.¹³ Such reflection, no matter how helpful, would constitute only one step toward developing a system for scarce, life-saving resource allocation. In order to develop such a system, we would need to take other steps—for example, provide moral foundations for the principles we are striving to balance and determine which other principles must be included. For present purposes, we simply assume that the two imperatives we are discussing (or ones akin to them) do rest on secure moral foundations.

Balancing between the defeasible imperatives to save as many lives as we can and to extend life as much as possible needs to occur, of course, because allocating resources in order to maximize lives saved does not always maximize life-years saved and vice versa. Suppose, for example, that we have to choose between saving one person for 11 years or five people for 2 years each. Prognosis would favor saving the one, while maximizing the number of lives saved would favor saving the five.

Here is a proposal for balancing these principles. We begin by determining the proportion between the values relative to each principle that are manifested in the sets of persons who are in competition for the resources. (For the sake of simplicity, here we focus on two sets.) The value relative to prognosis is the number of additional life-years made possible, while the value relative to life-saving is the number of lives saved. The set that contributes the higher value to the proportion relative to a principle is “favored” on that principle. We then determine which proportion relative to each principle is greater. We preserve the set of persons that is favored by the proportion that yields the higher number.

Our example will help to illustrate the procedure. We must choose between saving one person for 11 years and saving five people for 2 years each. The one person has a higher value relative to prognosis, but the group of five has

a higher value relative to life-saving. Regarding prognosis, the proportion between the values possessed by the one versus the group is 11/10 (11 years versus 5×2 years). Thus, the one person is favored. In contrast, the proportion between the values possessed by the group and the one regarding life-saving is 5/1 (5 lives saved versus 1 life saved). On this principle, the group is favored. The second proportion is equivalent to a number (5) that is greater than that yielded by the first proportion (1.1). So, according to this method, we should save the group of five persons.¹⁴

We offer this proposal as a baseline, intended as a starting point for further investigation. This investigation should not be limited to philosophical inquiry, but should also include empirical research on trade-offs people are willing to make between saving and extending life. Balancing policies (for instance, policies that help to determine how scarce flu vaccine gets distributed in a pandemic) should emerge in part from deliberation by publicly accountable officials, aided by public input. Persad and colleagues would likely agree: They argue that it is important that allocation schemes be legitimate. Using empirical evidence on people’s moral judgments about trade-offs is one means to take in the service of legitimacy. Our baseline proposal provides a starting point both for empirical research and for public deliberation.

Although our proposal has intuitively plausible implications in a variety of cases, it does generate controversial results in others. To use another schematic example, suppose we could save one person for 30 years or two people for 7 years each. Regarding prognosis, the proportion between the values is 30/14 (= 2.14) in favor of the one, while regarding life-saving the proportion is 2/1 (= 2.0) in favor of the two. So the procedure would entail that we save the one person for 30 years. But some think we should save the two. Although people value both life-saving and life extension, they seem to give more weight to the former. Health economists have begun to carry out empirical research on the relative weights that people assign to saving and extending life (Nord 1999; Nord et al. 1999). Such studies might go some way in helping us to improve the proposed scheme.

But the prospect of appealing to them raises an additional question for bioethicists and philosophers: What are the roles that need to be played in the development of a legitimate allocation system by empirical studies on the trade-offs people are willing to make between health-related goods, philosophical work on distributive principles, and public deliberation? We cannot try to resolve this question here. But we believe that ethical defenses of distributive principles and their balancing procedures should serve as constraints on both public deliberation and the use of preference data. Only in this way can we ensure that the outcome of public deliberation and the use of people’s preferences do not merely reflect prevalent prejudices and

13. By a defeasible imperative, we mean simply an imperative that can legitimately be overridden by some other principle in an allocation system. If an imperative to preserve the most lives were categorical, in contrast to defeasible, then, according to it, any allocation that did not maximally preserve lives would be wrong.

14. A fully developed weighing scheme would have to be sensitive to the uncertainty of a choice regarding both the number of persons preserved and the duration of their preservation. This is a further complication that we set aside.

lead to inconsistent policy choices. At the same time, different societies may be willing to make different trade-offs between different principles, and, within the appropriate ethical limits, sensitivity to these differences would be a desirable feature of any allocation system for scarce life-saving resources.

Someone might claim that if a system emerges from a “fair procedure,” it thereby gains all the moral justification it needs. One might, for example, hold that no matter how weak (or lacking in coherence) the philosophical arguments in favor of a system might be, if the majority of those potentially affected by it offer their informed, voluntary endorsement of it in a vote, then the system is morally sound. We disagree. We believe that an allocation system can be morally unacceptable even if it has been embraced through such a procedure. The system might, for example, fail to take sufficient account of the interests of some minority among its stakeholders. Although we will not defend the point here, we doubt whether there exists a practically realizable formal procedure such that one could plausibly claim that every allocation system that emerged from it would be above moral reproach.

What about the idea that it is more tragic if a young adult dies than if an infant does? Suppose, for the sake of simplicity, that only two principles—namely, prognosis and maximizing the number of lives saved—come into play in a choice between giving scarce, life-saving flu treatment to six infants or to three young adults. The infants would have priority over the young adults, if we assume as we did earlier in this paper that the treatment is equally effective in both groups and that the dosage for infants is less than half of what it is for adults. Giving priority to infants would maximize the saving of both lives and life-years. Some, like Persad and colleagues, who are attracted to the idea would find this result counterintuitive.

Our way of avoiding this result would be to distinguish between human beings who have and those who lack certain psychological capacities, including, for example, those to set ends and to form, act on, and revise plans for attaining them. We might then privilege preserving and extending the lives of those who have the capacities (i.e., “persons”) over the lives of those who lack them, including infants. A basis for such privileging would be the notion that by virtue of possessing these capacities beings have special worth or dignity. Appealing to this basis would not commit us to the view that beings who lack the capacities lack intrinsic value altogether. They might nevertheless have such value, just not as much as those who possess the capacities.

Careful reflection would need to precede a determination of how much priority to give to persons. Many of us would presumably reject a priority scheme that entailed that we save an adult who would go on to live 1 additional year with the capacities in question rather than saving 20 young children each of whom lacks the capacities now but would develop them in a few years and possess them for many decades. But then what should the priority scheme be (if indeed there should be one)? We do not try to answer

this question here. But it is just the sort of question that, we hope, our baseline proposal might prompt.

In setting out our baseline balancing proposal, we do not take ourselves to have been constructing, let alone defending, a whole system for scarce, life-saving resource allocation. First, in order to provide practical guidance, a system would need to include more principles than the ones we have discussed in this section. To cite just one example, it would have to incorporate a principle that prescribes how to proceed when, relative to all of the other principles in the system, patients are equally strong candidates for a resource. Second, like any principles in a defensible system, ones prescribing that we preserve persons or life-years for persons would need secure moral foundations. We have not, of course, provided such foundations here, although we try to show elsewhere that a broadly Kantian notion of the dignity of persons is capable of grounding both of the principles (Bognar and Kerstein 2010; Kerstein 2010).

CONCLUSION

In this paper, we raised objections to the complete lives system of allocating scarce life-saving medical interventions, developed by Govind Persad, Alan Wertheimer, and Ezekiel J. Emanuel. We argued that their proposal lacks secure moral foundations and fails to provide meaningful guidance when its component principles conflict. In our view, the balancing of competing principles is the most difficult part of any allocation system. We made a proposal for how to deal with one sort of conflict that often arises between principles. We emphasized that it is a baseline proposal, one that should be developed further by both empirical and conceptual work. Reflection on proposals like ours is necessary if we are to make progress toward answering the troubling and urgent question of how to decide who is to live and who is to die. ■

REFERENCES

- Bognar, G. 2008. Age-weighting. *Economics and Philosophy* 24: 167–189.
- Bognar, G. and S. J. Kerstein. 2010. Saving lives and respecting persons. *Manuscript*.
- Brock, D. W. 2002. Priority to the worse off in health care resource prioritization. In *Medicine and social justice: Essays on the distribution of health care*, eds. R. Rhodes, M. P. Battin, and A. Silvers, 362–372. New York: Oxford University Press.
- Centers for Disease Control and Prevention. 2009. *Emergency use authorization of Tamiflu (oseltamavir): Fact sheet for health care providers*. July 14. Available at: <http://www.cdc.gov/h1n1flu/eua/pdf/tamiflu-hcp.pdf> (accessed July 20, 2009).
- Daniels, N. 1988. *Am I my parents' keeper? An essay on justice between the young and the old*. New York: Oxford University Press.
- Daniels, N. 2008. *Just health: Meeting health needs fairly*. New York: Cambridge University Press.
- Dworkin, R. 1993. *Life's dominion: An argument about abortion, euthanasia, and individual freedom*. New York: Knopf.

- Gandjour, A. 2009. Ethical criteria for allocating health-care resources. *Lancet* 373: 1425.
- Jecker, N. 1991. Age-based rationing and women. *Journal of the American Medical Association* 266: 3012–3015.
- Kamm, F. M. 1993. *Morality, mortality. Volume I: Death and whom to save from it*. New York: Oxford University Press.
- Kerstein, S. J. 2010. Dignity and preservation of personhood. In *Humiliation, degradation, dehumanization: Human dignity violated*. Paulus Kaufmann et al. ed. Berlin: Springer.
- Nagel, T. 1979. *Mortal questions*. New York: Cambridge University Press.
- McKie, J., J. Richardson, P. Singer, and H. Kuhse. 1998. *The allocation of health care resources: An ethical evaluation of the 'QALY' approach*. Aldershot: Ashgate.
- Morreim, E. H. 1995. *Balancing act*. Washington, DC: Georgetown University Press.
- Nord, E. 1999. *Cost-value analysis in health care: Making sense out of QALYs*. Cambridge: Cambridge University Press.
- Nord, E., J. L. Pinto, J. Richardson, P. Menzel, and P. Ubel. 1999. Incorporating societal concerns for fairness in numerical valuations of health programmes. *Health Economics* 8: 25–39.
- Persad, G., A. Wertheimer, and E. J. Emanuel. 2009a. Principles for the allocation of scarce medical interventions. *Lancet* 373: 423–431.
- Persad, G., A. Wertheimer, and E. J. Emanuel. 2009b. Ethical criteria for allocating health-care resources—Authors' reply. *Lancet* 373: 1424–1425.
- Rawls, J. 1971. *A theory of justice*. Cambridge, MA: Harvard University Press.
- Tsuchiya, A., P. Dolan, and R. Shaw. 2003. Measuring people's preferences regarding ageism in health: Some methodological issues and some fresh evidence. *Social Science and Medicine* 57: 687–696.

Correspondence

Response to Open Peer Commentaries on “Complete Lives in the Balance”

Samuel J. Kerstein, University of Maryland, College Park
Greg Bognar, New York University

We are grateful to our commentators not only for showing us ways in which our views need to be refined and improved, but also for making points that further the daunting but necessary debate about how to allocate scarce, life-saving resources justly. We believe that the prospects for progress in this debate improve dramatically when positions are defended as they are by Govind Persad, Alan Wertheimer, and Ezekiel J. Emanuel (2009) and by our commentators: on the basis of argument rather than rhetoric.

In our view, one of the most important questions in this debate concerns the appropriate division of labor between bioethicists, philosophers, health economists, and other empirical researchers. How can bioethicists and philosophers contribute to the task of working out legitimate allocation systems? In our discussion of the complete lives system, we argued that they can help identify and clarify relevant ethical principles. These principles, in turn, can serve as constraints on public deliberation and the use of preference data. We are in agreement with Persad and colleagues (2009) in rejecting the idea that it suffices to follow a “fair procedure” to choose justly between different allocations.

We soon part ways from them, however. In any allocation system, ethical principles will often come into conflict with one another, and we believe that bioethicists and philosophers can and should make recommendations for balancing conflicting principles. An allocation system needs to be able to provide guidance. We argued that in a range of central cases the complete lives system fails to do so. In their reply, Persad and colleagues (2010) say that the complete lives system “identifies the relevant principles and counsels against using only one principle or including flawed principles” (46). But this is hardly sufficient to resolve conflicts: even if you know which principles you should *not* use, you still have to find a way for balancing those which you *should* use. Persad and colleagues may take Ole Norheim’s (2010) view, namely that if you have all the relevant principles, then “it is not a fundamental problem that the correct balance between them must be worked out in each specific

case” (60). But our complaint was that the complete lives system lacks the resources to tell us how to work out that balance. In our view, this remains one of its fundamental shortcomings.

To be sure, Persad and colleagues try to meet our challenge by showing what the complete lives system would recommend in one of our representative cases. In this example, we can save either two 18-year-olds who will live an additional two years each, or one 18-year-old who will live an additional 52 years. They say that “only one of the eighteen-year olds has a chance at a complete life; the other two do not. We should thus give the organs to the person who can live to 70” (46). They reach this conclusion by appealing to “the core idea of a complete life” (46). But they do not show how their conclusion follows from the application of the principles of modified youngest-first, prognosis, save the most lives, and so forth. We think that is because, ironically, the principles that constitute the complete lives system fail to capture the core idea of a complete life. Persad and colleagues seem in effect to be embracing a new principle—namely, one that tries to balance between save the most lives and prognosis by giving priority to the latter when acting in accordance with it involves saving at least one person who would live a complete life. We are not convinced that this balancing principle is plausible. It implies that we should give priority to saving one person who would live a complete life over several others of the same age who each could otherwise continue to live for a substantive period, but short of a complete life.

In their reply, Persad and colleagues also appeal to the Rawlsian notion of reflective equilibrium, which they believe allows them to combine considered judgments about the youngest-first principle and the “investment modification,” resulting in what we called the modified youngest first principle. But they do not rebut our argument that the prioritarian justification they explicitly embrace for the youngest-first principle undermines modified youngest-first. In our view, they are left with a choice between giving up this justification of youngest-first

Address correspondence to Samuel J. Kerstein, University of Maryland, Department of Philosophy, College Park, MD 20742, USA. E-mail: kerstein@umd.edu

and abandoning modified youngest-first. They would presumably do the former. But if so, what would remain to ground the youngest-first component of modified youngest-first? Our objections are not met by simply proposing a principle that mixes an investment modification into youngest-first. Moreover, we should ask whose considered judgments are supposed to play a role in their appeal to reflective equilibrium. For we are not alone in our hesitance to accept that “investment” and age are appropriate factors in rationing. (We return to the investment argument below.) We are still not sure what role people’s judgments and public deliberation play in the complete lives system.

In contrast, we give an important role to people’s judgments and public deliberation in our “comparative proportion procedure” (we thank Paul Menzel (2010) for the apt name, which we find a bit less of a mouthful than Erik Nord’s (2010) “proportional trade-off between priority factors”). The procedure could be used in conflicts between the principles of maximizing the number of lives saved and maximizing the number of life-years saved (prognosis). The procedure was intended as a baseline proposal, to be fine-tuned on the basis of people’s willingness to trade off these (and other) moral factors. For we are in agreement with Nord: if people in different countries and societies hold different value judgments about the relative importance of different factors, a legitimate allocation system should, within moral constraints, be sensitive to such societal values.

We did not, however, intend to give a role to age in our procedure—contrary to Nord’s reconstruction of it. But Nord is correct in pointing out that the comparative proportion procedure gives more weight to the number of people saved, relative to the number of life-years saved, than the conventional QALY procedure. It seems to us that this is an advantage of the procedure. For we hypothesized that people tend to view saving *more people* to be morally more important than saving fewer people for a *longer time*. Menzel suggests that this may be not because of the higher *ratio* of number of lives saved in one outcome than in another, but simply because more people are saved. In other words, our choice in trade-off situations is not a matter of proportions. But insofar as we are concerned with balancing conflicting principles, it is natural to think of the trade-offs involved in terms of the proportions of the gains along different “dimensions.”

John McMillan and Tony Hope (2010) raise important further issues. We find ourselves in agreement with many of their points. McMillan and Hope emphasize the role of uncertainty and the probabilistic nature of many of the judgments that would have to be made in applying an allocation system. In our article, we acknowledged the role of uncertainty but set it aside to keep the discussion manageable. We agree that a fully developed system of principles needs to address this issue. McMillan and Hope also point out the ambiguity of the notion of a “life saved.” Since the best any intervention can do is to extend life, you can never really “save” a person’s life. But we think we can deal with

the ambiguity if we remember that what matters is not the time for which your life is extended, but what you can do with that period: complete your life plan, organize your affairs, spend time with your loved ones, and so on. The point is that saving a life is to extend it for a *morally significant time*. This helps us avoid the sort of counterexamples where our procedure seems to imply that you should extend the life of a very large number of people by a few seconds, rather than extending the lives of a few people for a longer period. What counts as a morally significant extension of life is context-dependent. But does this imply that no counterexamples can be cooked up against the comparative proportion procedure? No, it doesn’t. However, since our procedure is intended as a starting point, we don’t think this is an insurmountable problem.

Let us return to investment modification. Persad and colleagues claim that we have misunderstood the grounds they offer for modifying a youngest-first principle. They base their modification on the “social and personal investment that people are morally entitled to have received at a certain age” (2009, 428). We find this claim puzzling. If we take it at face value, their priority curve (40) seems to imply that the growth in societal investment that a person is entitled to as she moves from 4 to 6 years of age is greater than the growth in societal investment that she is entitled to as she moves from 10 to 12. (The slope of the curve is steeper from 4 to 6 than from 10 to 12.) But why would that be?

Alexander Friedman (2010) joins Persad and colleagues in challenging our criticism of the investment argument as a basis for the modified youngest first principle. Friedman claims that as a basis for the principle “we should consider the investments made by people in *their own lives*,” including in terms of “emotional growth, development of plans and dreams, [and] fostering of relationships” (58). Contrary to Persad and colleagues’ intention, the investment argument supports the view that 40-year-olds are to be prioritized over 20-year-olds, we suggested. And Friedman’s rationale suffers from the same difficulty: 40-year-olds have made more investments in their own lives than have 20-year-olds, so his modified investment argument also suggests that 40-year-olds should get priority. Friedman might try to avoid this result by insisting that 40-year-olds should not get priority because they’ve had more “return” on their investments. But then he would have to explain how he could avoid embracing the conclusion that, contrary to modified youngest-first, 20-year-olds should not get priority over 10-year-olds. For have not 20-year-olds received more return on their investment than 10-year-olds?

Menzel finds something intuitively attractive in the modified youngest-first principle—in particular in the priority it gives to saving a 10-year-old over saving a 4-year-old. Menzel mentions some abilities that he believes 10-year-olds, but not 4-year-olds, possess, including that to be proud of what they have accomplished with others outside of the home and of their relationships with siblings, parents, and grandparents. But one of us, the father of a 3-year-old daughter, is convinced that preschoolers routinely

have these abilities as well. Perhaps Menzel is correct in suggesting that a 10-year-old, but not a 4-year-old, can “have conception of a ‘life well lived,’ and of things she aspired to that she wished she had done better” (50). But we simply do not see why this ability would entitle the 10-year-old to higher priority.

Here we must acknowledge Sadath Sayeed’s (2010) point that many people do not see why someone’s having the capacity to set and pursue ends would entitle him to priority over another who lacks this capacity, as we suggest it might. Our position obviously requires defense that we do not undertake in our article. Our main aim there was to suggest a way for those who, upon reflection, hold that saving an infant should have lower priority than saving an adolescent to preserve this judgment without having to embrace the idea, implied by Persad and colleagues, that saving a 15-year-old should have lower priority than saving a 25-year-old. Sayeed also raises the question of whether we believe that a 10-year-old with Trisomy 21 and moderate cognitive impairment has the capacities requisite for possessing special worth or dignity. We do indeed.

Another point of clarification is needed in connection with the commentary by Vawter and colleagues (2010). They point out that, according to Minnesota project participants, maximizing life-years of persons might sometimes be unfair—for example, if it led to helping some with longer expected life spans (40-year-old white women from high income areas) rather than others with shorter expected life spans (40-year-old men of color from low-income areas). But as we make explicit in our article, we do not try to construct a whole system of allocation principles. In our view, such a system would have to take account of the role of racial injustice in the development of life span disparities. That being said, we believe there is something *morally misguided* in an allocation system that, for example, would give a 90-year-old woman from impoverished socioeconomic circumstances who, with some scarce life-saving treatment, would live an additional five years precisely the same chance of receiving the treatment as a 20-year-old woman from equally impoverished socioeconomic circumstances who, if treated, would live another sixty years.

One final comment. As it is no doubt evident from the contributions to this debate, a legitimate allocation system will be complex, multifaceted, and, as we put it in our original article, implementing it in practice will require “painstaking weighing of competing principles and

context-sensitive judgment.” Friedman worries about a basic issue in connection to this—namely, that no matter how we refine any system of principles, we will always end up with incoherence: our judgments about different allocations will be intransitive. In his view, intransitivity is a pervasive feature of our ethical world.

It is evidently beyond the scope of this reply to address this issue. As a theoretical problem, we are not convinced by the threat of intransitivity. But we agree it is likely to be a practical problem: both because there is no guarantee that people’s judgments will be coherent, and because of the context-dependence of the notion of a life saved that we discussed above. In the complex world of resource allocation, transitivity and coherence are best seen as aspirations: We should strive to avoid intransitivity and other forms of incoherence, but we should not be paralyzed if we can’t. ■

REFERENCES

- Friedman, A. 2010. Complete lives, incomplete theories. *American Journal of Bioethics* 10(4): 58–60.
- McMillan, J., and T. Hope. 2010. Balancing principles, QALYs and the straw men of resource allocation. *American Journal of Bioethics* 10(4): 48–50.
- Menzel, P. 2010. Complete lives, short lives, and the challenge of legitimacy. *American Journal of Bioethics* 10(4): 50–52.
- Nord, E. 2010. Balancing relevant criteria in allocating scarce life saving interventions. *American Journal of Bioethics* 10(4): 56–58.
- Norheim, O. F. 2010. Priority to the young or to those with least lifetime health? *American Journal of Bioethics* 10(4): 60–61.
- Persad, G., A. Wertheimer, and E. J. Emanuel. 2009. Principles for the allocation of scarce medical interventions. *The Lancet* 373: 423–431.
- Persad, G., A. Wertheimer, and E. J. Emanuel. 2010. Standing by our principles: Meaningful guidance, moral foundations, and multi-principle methodology in medical scarcity. *American Journal of Bioethics* 10(4): 46–48.
- Sayeed, S. 2010. Assessing the modified youngest-first principle and the idea of non-persons at the bedside: A clinical perspective. *American Journal of Bioethics* 10(4): 52–54.
- Vawter, D., J. E. Garrett, K. Gervais, A. Prehn, and D. DeBruin. 2010. Dueling ethical frameworks for allocating health resources. *American Journal of Bioethics* 10(4): 54–56.