

Reconciling cost-effectiveness with the rule of rescue: the institutional division of moral labour

Shepley Orr · Jonathan Wolff

Published online: 18 April 2014

© The Author(s) 2014. This article is published with open access at Springerlink.com

Abstract Cost-effectiveness analysis suggests that a society should allocate its health care budget in order to achieve the greatest total health for its budget. However, in ‘rescue’ cases, where an individual’s life is in immediate peril, reasoning in terms of cost-effectiveness can appear inhumane. Hence considerations of cost-effectiveness and of rescue appear to be in tension. However, by attending to the division of labour in medical decision making it is possible to see how cost-effectiveness analysis and rescue-style reasoning are commonly combined in an ethically acceptable fashion.

Keywords Cost-effectiveness analysis · Rule of rescue · QALY · Health Care Resources

1 Introduction

How much should we spend to save a life? Even to ask this question appears heartless, and to betray a fundamental lack of humanity. When lives are at stake, how can we even think of money? Yet it is clear that in the context of health systems we can always do more. We can keep people alive for longer, or at least increase their chances of survival. When treatments are cheap, or, perhaps, where the numbers of people needing help are very small, the decision can be easy to make. Often, however, we find ourselves in circumstances where life saving or extending treatments are very expensive. To fund costly treatments on a large scale threatens to be damagingly expensive, driving expenditure from areas where it could do more total good. The promptings of

S. Orr · J. Wolff (✉)

Centre for Philosophy, Justice and Health, University College London, London, UK
e-mail: j.wolff@ucl.ac.uk

S. Orr

e-mail: s.orr@ucl.ac.uk

basic humanity can, therefore, come into worrying conflict with mundane, but vitally important considerations of cost-effectiveness.

In the literature on health care allocation the ‘humanitarian’ argument has often been couched in terms of a supposed ‘rule of rescue’ (Jonsen 1986; McKie and Richardson 2003; Sheehan 2007). Outside the health care system from time to time, we are faced with ‘rescue situations’: a child who has fallen down a well; miners trapped underground; a submarine languishing at the bottom of the ocean. In all of these cases, society responds as a matter of supreme urgency and apparently without reference to questions of cost. Action is required, and it is attempted. Few economists would dare enter the public debate to raise questions about whether the rescue is cost-effective, diverting resources from a more effective use.

In the world of health care allocation, we rarely meet situations that are directly analogous to rescue cases. But nevertheless there are many examples where a life is in immediate peril and steps could be taken that have a chance of making a difference. Emergency surgery, intensive care, antibiotics for severe infections, transplants or artificial organs, highly active anti-retro-viral therapy (HAART) and life-preserving cancer drugs all can make the difference between life and death. In health care, it seems, rescue, in some form, is routine. Should we put cost-effectiveness analysis to one side in such cases, or would it be morally irresponsible to do so? In this paper, we will argue that the standard methodology in health technology assessment already provides a plausible balance between cost-effectiveness and rescue. We do not claim that any standard methodology at present provides exactly the right balance, but rather that something like an institutional division of moral labour is a permissible method for achieving the correct balance between humanity and cost-effectiveness. But to make this case, we need first to review the standard debate, and to consider what we believe to be an unsuccessful attempt to achieve a reconciliation, in order to be in a position to present our own understanding.

2 Cost-effectiveness analysis in health

The standard approach to health technology assessment, especially as practiced in the UK, under the auspices of the National Institute of Health and Care Excellence (NICE¹) is widely known and so here we can be brief. The question faced by the health service is how to get the best ‘value’ out of a fixed health budget, or at least a budget that is fixed within a single budgetary cycle. In the first instance, ‘value’ is to be understood merely as ‘health’, although this is controversial and an issue to which we will return. In sum, then, the Department of Health and National Health Service seeks to achieve the greatest health for its budget.

In order to do this a common currency of health, across different conditions, is needed. It is all very well, for example, to say that for a certain sum of money we could conduct x number of heart transplants or y number of hip replacements, but, it seems, we need some sort of common currency to be able to say which of the

¹ NICE has been renamed twice, reflecting changes in responsibilities, since its foundation in 1999, as the National Institute of Clinical Excellence, although the acronym has always remained the same. Its current name dates from April 2013.

two procedures will provide us with the greatest health. To accomplish this task in the context of health care resource allocation the idea of the ‘Quality Adjusted Life Year’ has been devised (see [Cookson and Culyer 2010](#) for useful discussion). The ingenuity of this measure is that it offers a way in which improvements in health that provide no additional life expectancy can be compared with treatments that increase life expectancy. A year of life in full health is given a value of 1.0 QALY. Detriments to health are theorized as yielding life-years of less than 1.0 QALY. Suppose a year of life of someone who suffers severe pain and mobility problems but has no other condition is given a QALY value of 0.6. This would mean that restoring them to full health will provide an increment of 0.4 QALYs a year for the remainder of the person’s expected life (here we ignore the important question of time discount, as it is not germane to the issues of this paper). Providing a hip replacement will improve health, but not normally extend life, while providing a heart transplant will typically both improve health and extend life. QALY values are elicited from the public by a variety of questionnaire methods to capture their preferences, and by these means all therapies can, in principle, be given a QALY value. Not only that but when the cost of the procedure is known it becomes possible to calculate the cost-per-QALY of different treatments. With this information, and with the assumption that the aim of the health care system is to maximize health, cost-effectiveness analysis becomes the matter of seeking to achieve as many QALYs as possible from a fixed budget ([Cookson and Culyer 2010](#)). And it is easy to see why this approach can lead to results that appear heartless. If a therapy is extraordinarily expensive, or of more reasonable cost yet effective only for a short period, or in a limited fashion, or in a small proportion of cases, then its cost-per-QALY gained will be high, and quite possibly over the current cut-off point in the UK of between £20,000 and £30,000. If maximizing health gain can be achieved by saving or extending lives then such treatments will be justified by cost-effectiveness analysis. But in the sad cases where the cost-per-QALY for a particular treatment regime is above the threshold, saving or extending lives will be deemed not cost-effective. On strict cost-effectiveness grounds medical rescues would not always be attempted ([Harris 1987](#)).

3 The rule of rescue

There seems little doubt that in non-medical rescue cases, we have a very strong aversion to the idea that some sort of cost-effectiveness analysis should be undertaken before action is undertaken. One argument that a similar approach should be taken in medical cases suggests that there is a strong analogy between non-medical rescue cases and the medical rescues we have sketched above. However, it is certainly possible to question how close the analogy is ([Cookson et al. 2008](#)). Typical rescue cases have a number of features. First, it is known that there is a particular person, often identified, who is in peril. Second, the threat is normally one of death. Third, without action death is expected soon, normally within days. Fourth, if the rescue is successful full health will be restored. Fifth, the situation is unusual. All of these create a picture in which it is common for people to have the further intuition that a rescue should be attempted without reference to its cost.

It is interesting to reflect on whether our intuitions would remain as strong if these conditions did not apply. Suppose, first that it is believed with a relatively high degree of certainty that someone will die tomorrow, but it is not known who it is. Would we consider the situation equally pressing? Well, not in all cases. For example, it is known in the UK that tomorrow, like every day, a handful of people will die in traffic accidents. Yet transport safety is an area in which a cost-effectiveness analysis is conducted before safety measures are introduced. Secondly, if the peril was not death but serious injury, say the loss of an arm, would rescue still be called for? Perhaps cases differ; certainly agonising pain has been used in an example where it is suggested a rescue is morally required (Scanlon 1998: 235) but it is surely possible to take the contrary view. Third, if the threat was not immediate but over a longer period, such that we could instead enact rescue tomorrow, or the next day, or the next, would we react the same way? Certainly it would lose the sense of urgency. Fourthly, suppose that rescue would not lead to complete recovery but a life of poor health. Does this change our intuitions? Finally, suppose that rescues were not rare events, but occurred every day, or even several times a day (road deaths) or perhaps many times a day, as they do in intensive care units around the country. The repeated situation may have a numbing effect, and it becomes increasingly difficult to ignore the budgetary impact, which would drain resources from other programmes. We will argue that the degree to which rescue situations are routine is what makes these types of case subject to some sort of deliberate, rational planning at an earlier stage.

From this, we can see that medical rescue cases do not share all of the features of paradigm non-medical rescues. On this ground, perhaps, one might simply dismiss the idea of the rule of rescue in medical contexts. Yet it is precisely in medical contexts that it was introduced, by Jonsen, who objected to what he called the ‘Bentham in a box’ approach of standard cost-effectiveness analysis (Jonsen 1986). And we can see his point. Although medical rescue cases are different from paradigm non-medical cases, they have enough in common that we can raise the question of whether we should respond as we do in other emergencies. When a road traffic accident victim is wheeled into the emergency room, or a patient develops a life-threatening infection, could it be right to conduct a cost-effectiveness analysis before proceeding? Of course, it will be important to conduct an analysis of clinical effectiveness to ensure that money is not being spent on futile treatment, which can also be to the detriment of the patient, but to make a calculation of cost-effectiveness, even when public money is being spent, would appear grotesque. So it appears that even in the current system some form of ‘rescue reasoning’ is already engaged. We will return to this important point shortly, but first we will consider some of the possible motivations for the rule of rescue.

4 Apparent motivations for the rule of rescue

There is no doubt that the rule of rescue has strong intuitive power in paradigm non-medical cases, and we believe that there will be traces, or more, of this ‘rescue intuition’ carrying over into at least some of the medical cases, even when they do not share all the features of the paradigm cases. To understand where and why these intuitions apply, and whether they should be relied upon, it will be helpful to explore their basis.

If, as is sometimes supposed, they can be best understood as a form of irrationality then there would be reason to dismiss them. So, for example, psychologists often mention an ‘identified victim effect’ where we treat the claims of those victims we can identify with greater urgency than those we cannot identify (McKie and Richardson 2003). This may explain why it is often easier to fund regimes of treatment than prevention, where the potential victims are not known, and indeed will never be known if the regime goes ahead. However, it would be wrong to assume in advance that the identified victim effect is an error of reasoning, for that is exactly what is at issue.

One line of support for the rule of rescue is inspired by reflection on a case originally presented by William Godwin: suppose you are faced with the necessity of making a choice between pulling your mother or someone of great public importance from a fire (assuming your mother has no particular public role) (Godwin 1793/1976). What should you do? Some forms of consequentialist reasoning suggest that you should save the public figure—in Godwin’s example Archbishop Fenelon—as this would lead to greater overall good. Others have suggested that even though you generally have a duty to do the greatest good, in this case your special relation to your mother allows you to override your general duty, and hence you can reason: ‘Morality generally requires me to save the Archbishop, but because I have a special relation to my mother, I am permitted in this instance to save her instead.’ Bernard Williams, however, has suggested in examples of this type that this reasoning includes ‘one thought too many’ (Williams 1981). Someone with a sound ethical sense would simply save his or her mother, without going through a process of reasoning in which that conclusion is derived from more general considerations. The general point can be formulated in terms of sympathy: ethical motivation is to a large part the result of sympathy and there are cases where the ethically right thing to do is to act out of sympathy without engaging reason. Holding this view does not commit you to the view that there is never a place for calm reason, but simply that there are situations where reasoning about the requirements of morality itself seems to betray a lack of clear moral judgement.²

A different type of view, variants of which can be found in Kant (1785/1998), Kamm (1998, 2007) and Scanlon (1998), is that the problem with the reasoning just summarized is not that reasoning is involved, but rather it is the wrong type of reasoning. Morality is constituted not by calculation of consequences but by duties. These duties are justified or generated in different ways by different theorists, but on some theories include a duty to rescue those in immediate peril. How convincing this is will depend on the details of the particular theory, which is outside the scope of this discussion, but we include the point here to acknowledge that the rule of rescue is consistent with highly rationalistic forms of moral argument, in which reason, rather than sympathy is paramount.

² Strictly speaking this example does not help illustrate rescue reasoning as, after all, the Archbishop also needs to be rescued. We thank an anonymous referee for this observation. However the point of the example is to contrast impersonal moral reasoning, which arguably points in the direction of saving the Archbishop, and a more direct and personal form of moral reasoning which strongly supports saving one’s mother and arguably silences other moral concerns. It is suggested that on one account, the basis of the rule of rescue is the immediate sympathetic pull to those in special circumstances that is constitutive of sound ethical thought. The special circumstances can include close personal relations and extreme and urgent need or peril.

Finally, it is possible to argue that the rule of rescue has a type of symbolic, or perhaps expressive, value. This account is not, in fact, in competition with the earlier accounts, but can be used to clarify or perhaps augment their appeal. Consider what it would be to neglect someone in immediate peril on the grounds that more good elsewhere can be done with the money. This will strike many people as impossibly callous or inhumane. One of the things that is most important about us as human beings is our fellow feeling for other human beings in difficulties. Life would be much diminished if we did not treat urgent threats to other people as exerting an enormous pull on us. We express our common humanity and give meaning and value to our lives and those of others by responding to each other when we are in severe need. We would undermine our nature as human beings if we calculate that on occasion extreme need can be outweighed by other considerations, such as the greater good we could do by alternative actions. The rule of rescue, then, is a way of expressing or symbolizing our nature as creatures who respond to each others' severe needs.

5 Criticisms of the rule of rescue

Although the arguments given in the last section may seem strong, they are not always treated as compelling. There are at least two ways of attempting to counter them. First, it could be argued that the rule of rescue never applies. It is simply wrong to rescue a little girl down the well if, for the same money, two little girls will be saved at some point in the future using the same money to build fences around wells. Second, it could be suggested that while the reasoning does apply to the paradigm cases, the medical cases are very different and the reasoning does not typically carry forward (although when the conditions of the paradigm cases are exactly replicated in a medical situation it will) (Cookson et al. 2008).

The first reply above is based on the obvious point that allocating resources according to the rule of rescue is likely to be highly inefficient in a broader context. As just seen, we would need to give priority to save life here and now, possibly at the cost of saving life elsewhere or in the future. Furthermore, we would be required to give priority to cases that might involve death over those where threats to health but not to life are at stake. In QALY terms this could lead to very significant QALY loss. If we follow the rule of rescue, then we must save the little girl, the miners or the sailors even if with the same resources we could conduct a public health project that would improve the health of tens of thousands. Thus, we would be forsaking potential large gains in health in order to save a small number of lives.

It is unlikely, however, that society could be persuaded to give up rescue attempts in cases that meet the conditions set out in Sect. 3 above. However, as we noted, medical cases rarely, if ever, meet all these conditions to the full degree. Medical 'rescue-style' cases, as we understand them, involve a threat to the life of a particular, known, individual that could be reduced by the application of a particular, known, treatment. Hence they are a subset of cases of medical need; essentially severe and urgent. Seen this way, the above charge of inefficiency really shades into another; that it is *unfair* to those who are not in immediate peril. Why should a health service give such a high priority to the immediate saving of, or extension to, life? The invention of the QALY

was tacit recognition that extension to one patient's life has to be balanced against improvements to another patient's life or health. What is the justification for ignoring one type of health problem for the sake of another? Why should sympathy, duty, or symbolic value only attach to the immediate saving of life and not the reduction of other suffering? It is very hard to see why. Although cost-effectiveness analysis is often associated with a utilitarian moral theory because of its practice of maximization of a particular currency, it is in fact a concern with fairness that all individuals should have an equal claim to some shared resource, that can be seen as providing the normative force behind cost-effectiveness analysis.

Finally, incorporating the rule of rescue in medical practice would be chaotic, constantly pulling us away from routine, but important tasks. It will be very disruptive and make rational planning of services difficult. It is true that we have evolved systems to accommodate emergency treatments without disrupting everything else, but applying a version of the rule of rescue appears to expand the scope of such services dramatically. How can we manage a system where we need constantly to be alert to cases—common cases—that take priority over all other types of case?

6 Could there be a 'rescue-adjusted QALY'?

Perhaps these criticisms could be answered by modifying the QALY so that the calculation also includes the additional value of rescue. Rescue could be given extra weight in our calculations but not the apparent absolute weight called for by adherents of the rule of rescue. On this account, the saving of a life in rescue cases would be given an extra value. The QALY formula should be adjusted in such a way that it yields the judgements that we would intuitively make in rescue cases, allowing more spending in a rescue case than a standard QALY calculation would permit.³ It is true that such an approach would still be subject to the 'one thought too many' criticism, but at the same time perhaps this is a small price to pay for an approach that has other advantages.

Each of the three criticisms of the previous section would be addressed. First, rather than being inefficient, the definition of efficiency will have changed. Policy will maximize rescue adjusted QALYS. Second, fairness is restored, or even improved. All claims are considered but are now given their 'true' weight. Finally rational calculation is possible.

The problem, however, is that the resulting solution is very hard to accept. Suppose there is a rescue situation which needs to use resources that could have saved five lives over a longer period. Suppose the formula we have derived permits the rescue. This would imply that we value a rescue more than we do five lives. But is this really how we should capture the intuition? Do we want to endorse an approach that could have the consequence that sacrificing five people is acceptable but we must draw the line at six? Furthermore, do we really want to say that our common humanity is of a level of importance that we should allow five people to die so that we can express it? The point is not so much that it is hard to get the formula right, but that using some sort of formula to capture the value of responding to an emergency seems to falsify the

³ We first heard this suggestion from Donald Franklin.

nature of the value of humanity. Kant suggested that human life has a ‘dignity’ and not a ‘price’ (Kant 1975/1998, p. 43). QALY calculation has allowed us, controversially, to put a price on life in some cases, but it seems much harder to put a price on the ‘expression of our common humanity’. In sum, while the notion of a ‘rescue-adjusted QALY’ may have some merit (and we will return to this), it does not seem to engage properly with the motivations of rescue intuitions.

7 Reconciliation through division of labour

If it is not possible to derive a formula that does justice to the notion of humanity underlying rescue intuitions it may seem that we are condemned to choose between heartless cost-effectiveness analysis, some sort of unconvincing weighting, or finally unsystematic intuitive judgements, leading also to arbitrary unfairness. However, here we wish to argue that another approach is already implicit in social decision making in the medical sphere.

To see this, consider a familiar objection to QALY calculation, and an equally familiar reply. The objection is that QALY calculation discriminates against the elderly. If two patients present with the same condition and a drug will offer a cure, then under conditions of scarcity the drug will be given to a younger patient who will be expected to live longer and thereby gain more QALYs from the treatment. And this, it is often claimed, is a form of unfair discrimination. In response it is said that the QALY system as used in practice does not have this effect. Therapies are evaluated on the basis of the average cost-per-QALY they offer, which will inevitably differ between patients. Once the drug is approved, it is available for prescription to all patients for whom it is clinically indicated, and not just those for whom the QALY calculation individually is favourable (Brock 2004).⁴

The lesson to learn from this discussion is that one set of considerations is used to decide whether a treatment is to be made generally available, and a different set of considerations is used to decide who should receive it. This division of labour is key to understanding how to reconcile cost-effectiveness analysis with the rule of rescue.

Consider again the paradigm cases of rescue: the girl down the well, the trapped miners or stranded sailors. In all of these cases, we do not raise questions of cost; rather we do what we can with all available resources. But this notion of ‘all available resources’ is very revealing. Why is one set of resources available, rather than another? Presumably it would have been possible to invest in technology to make rescue easier, as many societies have done in the medical case with defibrillators, and, in earlier times, life-belts near harbours. Yet at some point, in the clear light of day, and with no pressing immediate emergency, and no identified victim in front of us, a decision was made to equip ourselves to one level and not more or less (or more likely a number of decisions were taken that had the unintended consequence that we are equipped to a particular degree).

⁴ It is often pointed out that this reply is not sufficient to silence the critics, for there are conditions that predominantly affect the elderly and drugs for those conditions will typically have a higher average cost per QALY than drugs that benefit younger groups. This may well be so, but it does not affect the point made in the text above.

In the medical case, the main sites of rescue are the emergency room, the operating theatre and the intensive care unit. In each case, care is given apparently without regard to cost. Yet in each case clear-headed decisions have been made, no doubt involving cost-effectiveness analysis, to provide a certain level of rescue resources. Once provided, they can be used without regard to cost, but it would be very misleading to say that considerations of cost have not entered into the equation at any point.

Note, however that what is true of the emergency room is equally true of the general practitioner's surgery, at least some of the time. Decisions have been made to allow the doctor to prescribe certain drugs when faced with patients in particular circumstances. The doctor is not required to carry out a cost-effectiveness analysis before treating each patient. In fact it is likely that we would wish to discourage the doctor from having such thoughts, or at least from making them central to deliberation. Rather, a doctor should be concerned, primarily, with the health of the patient and what can be done to improve it. The constraints of the system may make it the case that some possible therapies are not available, but the doctor's primary duty is to do the best for the patient with the materials at hand.

What we see, then, is an institutional division of labour, which is common to the most exotic rescue case and the most mundane routine medical treatment. A set of decisions is made in the committee room about the resources that are to be made available in front-line services, such as the doctor's office or the emergency room.⁵ These decisions will largely be made on grounds of cost-effectiveness. Emergencies will be predicted, though not faced. All lives will be statistical, or perhaps anonymous, rather than identified.⁶ Budget implications of any decision are likely to be significant. Rescues are not the immediate business. The agenda is to stock pharmacies, or to order a particular number and type of ambulances, or plan the training of emergency and non-emergency teams. The motivation of the committee members will be that of planners: to achieve the goals of population health to the greatest degree possible, subject to possible constraints, as well as the available budget.

Once the equipment is in place, then it is in the hands of staff who have a different motivation: saving the lives or improving the health of the people in front of them.

⁵ To denote the metaphorical and real places where cost-effectiveness analysis takes place, we will generally refer to the 'committee room', and to denote those locales where front-line services are delivered we will generally refer to the 'emergency room'.

⁶ Perhaps the briefest way to explain anonymous, identified, and statistical lives is to use the decision theorist's familiar example of balls in an urn. Suppose there are five agents, A, B...E, and that there are 5 balls in an urn, 4 black and one white. And let us assume that drawing the one white ball is to be assigned some negative consequence (in this case, death). In the case of 'identified' life, it is assumed some external agent or exogenous event has simply decided that the negative consequence will befall a particular agent. That is, in identified life cases, it is simply the case that 'X gets it'. In statistical life cases, assume that each agent draws a ball from the urn and then, regardless of the outcome, returns the ball to the urn for the next agent to choose. This entails that the white ball might never be drawn, or be drawn by all five, or, "everyone might get it, or no one might get it". In the anonymous life case, the agents each choose a ball from the urn, but do not return the ball. Whoever draws the white ball suffers the negative consequence, or, 'someone's gonna get it'. It is interesting that each agent faces a 1 in 5 chance of the negative consequence in the statistical and anonymous lives case, despite their being very different situations from an ethical perspective.

Something much closer to the humanitarian impulse of the rule of rescue is in play at the level of front-line services.

It is not at all difficult to understand the benefits of this division of labour. Unless budgets are unlimited it does not seem be possible to run any sort of large medical service without reference to some notion of cost-effectiveness and rational procurement. Under such circumstances, there must be such a thing as a possible treatment or procedure that is simply too expensive. Yet at the same time it seems highly undesirable for such calculations to be made in respect of any particular patient. We seem prepared to accept the idea that a statistical human life has a price, but far less prepared to put a price (in the form of a limit in expenditure) on the life of a person standing in front of us. It is much better for a doctor to tell a patient that decisions taken by others at another stage of the process have made certain treatments unavailable. The doctor/patient relationship will probably survive the doctor's complaint that his or her hands are tied. It will be put under intense pressure if the decision not to prescribe the treatment has been made by use of cost-effectiveness analysis in the emergency room.

There is some similarity here between our discussion of the division of moral labour through institutions and Nagel's (Nagel 1991: ch. 6) discussion of the division of moral labour. Our view and Nagel's are both attempts to reconcile two conflicting standpoints: in Nagel's case, the two standpoints are the personal and the impersonal, where in ours the conflicts are between the humanitarian motivation and the motivation to do as much moral good as possible per pound spent. And like Nagel, we argue that this is a conflict that exists within individuals (Nagel 1986; and see Wolff 2011: 90, 102), and because of the very fact of the difficulty of these intra-personal conflicts, we allocate some decisions to be made at an institutional level. (Our suggestion also bears comparison with Goodin 1995, who argues for different principles of morality to operate at a 'public' and a 'personal' level.)

8 Is division of labour really possible?

We have argued, then that a division of labour between cost-effectiveness analysis and 'rescue-reasoning' already exists in medical decision making. It exists in rescue situations and non-rescue situations. It exists in the emergency room and the local clinic. And we have also suggested that this is a beneficial situation from the point of view of doctor-patient relationships. And it should also be apparent that we regard this as an attractive way to reconcile what appeared to be two, irreconcilable, approaches to medical decision making. Yet there may be a lingering concern that the solution is too stark to be accurate. Is it really true that 'rescue' intuitions are absent in the committee room? Or that cost-effectiveness considerations never occur in the emergency room? If it is true that both types of considerations apply in both venues then it seems we are back to where we started: having to worry about how to combine two quite different styles of reasoning into a single decision procedure.

To start with the committee room, is it true that rescue intuitions have some place there too? Although the committee room will rarely have to deal with identified patients in known emergencies, nevertheless they can rationally predict that emergencies will occur and will need to be dealt with. Will considerations of basic humanity mean that

cost-effectiveness has to be abandoned, or at least modified, from time to time in the committee room? Here, we suggest that the notion of the rescue adjusted QALY may have some place after all, albeit in different terminology. The key elements of rescue that will arise in the allocation room are those of end of life and of severity of disease. In both cases, we have seen policy proposals that the calculation should be adjusted so that a higher cost per QALY is allowed (NICE 2009; Department of Health 2010). This, we suggest can be seen as moving in the direction of rescue, but doing so in the cold-light of day when such decisions have to be made as a matter of routine. The earlier objection—that humanitarian reasoning is debased by reducing it to a formula—still has some purchase, but the fact that the formula is used at a general level for resource allocation and not as a way of mediating a relationship between doctor and patient removes some of its sting. Hence at this level it seems satisfactory.

What, though, can we say about the emergency room. Is it really true that doctors should put considerations of cost completely to one side? This seems unrealistic and, indeed, unwise. There are regular campaigns, for example, to persuade doctors to prescribe generic medicines rather than more expensive branded equivalents. Doctors are often rightly reluctant to send their patients for expensive tests, and may spend a few weeks, or longer, in a state of ‘watchful waiting’ before doing so. The emergency room is rightly suffused with considerations of expense after all.

Yet it is interesting to reflect on how considerations of cost could enter the reasoning of the doctor. Robert Nozick distinguished two theories of rights: one in which a moral agent tries to minimize rights violations—the ‘rights as goals’ theory—and the other in which rights function as ‘side constraints’, setting absolute barriers to action (Nozick 1974). Although we would not want to follow Nozick’s model exactly in the current case, it reminds us that there is more than one way a consideration can function in one’s reasoning. Our suggestion is that the humane but responsible doctor will consider considerations of cost as a type of provisional side-constraint on action. By analogy consider, for example, how one chooses from a menu in a restaurant, when another person is paying. Some people will deliberately choose an expensive option, others a cheap one. But a common experience is simply to choose what you want, but keeping an eye out for cost, avoiding the most expensive dishes. Cost here functions as a background consideration. It is not a firm constraint, perhaps today you will order the lobster after all. Yet cost is something to be taken into account in background reasoning, albeit only when a certain threshold is crossed.

To apply the analogy to rescue cases, we may say that we will save lives at all costs, given the equipment available, yet at the back of our minds there may also be a thought that ‘enough is enough’. When the chances of a successful rescue are falling and the costs of the next step are increasing, the consideration of cost, reluctantly, comes into view. The rescue may be halted on cost-effectiveness grounds. Cost-effectiveness lurks in the background, but it is rarely central to the decision-making process, and it is not always over-riding even when it is relevant.

We suggest here that as an empirical account of moral psychology what is sometimes termed ‘threshold deontology’ is the kind of moral thinking that takes place in the emergency room. That is, one should act in a way consistent with one’s duty until cost reaches some threshold that suggests acting contrary to that particular duty; see Alexander (2000). And indeed, in every case in an emergency room where a life may

be lost, some decision is taken as to when to no longer apply resources to saving that life. For example, think of the familiar case in medical dramas where a doctor furiously trying to resuscitate a patient is told ‘Let it go. You’ve done all you can’, and then ‘time of death’ must be declared. Note that death is declared when the efforts to save the patient have stopped, not by some objective indicator of remaining brain or heart activity; this suggests that up to some limit there could always be a bit more done to save that life, but some consequential evaluation must take place when it is felt that the resources could be better used elsewhere than in trying to resuscitate a particular patient. This suggests that there is some small, implicit element of consequentialist concern with costs and benefits in almost every emergency scenario.

Similarly, as we have seen above, the committee room is not devoid of humanity. There will be norms implicit about how much humanitarian reasoning can enter into decisions made at the highest levels of allocation, and it will not necessarily be all ‘cold calculation’. There is an informal rule in environmental economics when eliciting preferences over environmental goods: ‘no fluffy bunnies’. This means that no images meant to prime the subject to over-value a particular good should be allowed in an objective valuation exercise. One can imagine as well that when committees make decisions there must be rules for how much affective priming will be allowed, but it will not be disallowed completely. Indeed, all NICE decisions will have input from patient groups where the issue being decided is intentionally given a human face.

A related idea to what we are suggesting draws on the role morality suggested by F.H. Bradley’s famous “My Station and its Duties” (1927).⁷ We generally want those in the committee room to follow the logic of promoting the best consequences, and we generally want those in the emergency room to act to promote health and life regardless of cost. Therefore, we create and design institutional roles which will allow individuals inhabiting those roles to ‘silence’, or ‘forswear calculation’ of, other types of consideration which might be given voice: in our case, the individual in the emergency room will (almost) silence considerations of cost, and the individual in the committee room will (almost) silence considerations of humanitarian ‘special pleading’.⁸ It is worth noting that the silencing considerations in the emergency room are specifically endorsed in the committee room; as Kamm puts it “...if a policy allocates some money to an institution like an emergency room, this might just be a way of saying that in some areas of life, however small, a different principle than is involved elsewhere governs distribution” (2007: 46, fn. 66).

Our reasoning here suggests that institutions can play an important role in resolving intra-personal value conflicts, as discussed here in the conflict between motives of pursuing the greatest good per pound spent and acting consistently with the rule of rescue. Where there are hard cases in conflicts of values we often create a division of labour between the choices made by individuals (in the emergency room) and the allocation of resources (in the committee room) for those individuals to make their choices. This issue has an obvious basis in normative ethics, but it may well be that

⁷ For a thorough discussion of this general line of thought which has influenced the views taken here see Ch. 7, “Roles and their Morality”, in Emmet (1966) *Rules, Roles and Relations*.

⁸ On motivational silencing, see McDowell (1978). On forswearing calculation, see Brennan and Pettit (1986).

different institutions separate moral labour between rationalist utilitarian calculation and non-calculating humanitarian motives, and these may well vary by the type of institution which has evolved to address a particular set of conflicting distributional aims (Elster 1992; Ostrom 2005), and hence this division of moral labour will need to draw on the study of social institutions in order to be properly addressed.

9 Conclusion

The rule of rescue has a strong intuitive pull. It seems to express our common humanity, and to refuse a rescue on grounds of cost appears morally horrendous, even in cases that do not share all the paradigm features of the rule of rescue. Yet at the same time in a complex, resource-constrained world, cost-effectiveness cannot be ignored. The two types of reasoning appear irreconcilable. We believe, however, that this appearance is misleading, and ordinary processes of medical decision making show how to reconcile the two. Resource allocation decision making broadly follows cost-effectiveness analysis, while emergency room and related 'bedside' decision making is much closer to rescue reasoning. There is good reason for this division of labour, although we have conceded that this simple picture does need to be modified to accommodate the different ways in which both styles of reasoning take place in both venues. Nevertheless, the key point remains: cost-effectiveness analysis is needed to decide which tools of rescue to provide. Rescue can then take place in a manner apparently unconstrained by cost.

Acknowledgments This work was supported by the Arts and Humanities Research Council, Ethics of Risk project held at University College London. It draws on a report undertaken for the Inter-Departmental Group on the Valuation of Life and Health (Wolff and Orr 2009). We are grateful to a seminar audience at the University of York's joint meeting of the Morrell Political Theory Workshop and School of Politics, Economics and Philosophy Seminar series for their comments on an earlier version of their paper. We would also like to thank Paul Anand, Albert Weale and two anonymous referees of this journal for their valuable suggestions

Open Access This article is distributed under the terms of the Creative Commons Attribution License which permits any use, distribution, and reproduction in any medium, provided the original author(s) and the source are credited.

References

- Alexander, L. (2000). Deontology at the threshold. *37 San Diego Law Review*, 893, 905–912.
- Bradley, F. H. (1927). *Ethical studies* (2nd ed.). Oxford: Oxford University Press.
- Brennan, G., & Pettit, P. (1986). Restrictive consequentialism. *Australasian Journal of Philosophy*, 64(4), 438–455.
- Brock, D. (2004). Ethical issues in the use of cost effectiveness analysis for the prioritization of health care resources. In S. Anand, F. Peter, & A. Sen (Eds.), *Public health, ethics, and equity*. Oxford: Oxford University Press.
- Cookson, R., & Culyer, A. J. (2010). Measuring overall population health - the use and abuse of QALYs. In A. Killoran & M. Kelly M (Eds.), *Evidence based public health: effectiveness and efficiency*. Oxford: Oxford University Press.
- Cookson, R., McCabe, C., & Tsuchiya, A. (2008). Public healthcare resource allocation and the rule of rescue. *Journal of Medical Ethics*, 34, 540–544.

- Department of Health. (2010). A new value-based approach to the pricing of branded medicines: a consultation. http://webarchive.nationalarchives.gov.uk/20130107105354/http://www.dh.gov.uk/en/Consultations/Liveconsultations/DH_122760. Accessed 15 April 2014.
- Elster, J. (1992). *Local justice: How institutions allocate scarce goods and necessary burdens*. New York: Russell Sage Foundation.
- Emmet, D. (1966). *Rules, roles and relations*. London: MacMillan.
- Godwin, William (1793/1976) *An Enquiry concerning Political Justice* (Harmondsworth: Penguin Press, 1976).
- Goodin, R. (1995). *Utilitarianism as a public philosophy*. Cambridge: Cambridge University Press.
- Harris, J. (1987). QALYfying the value of life. *Journal of Medical Ethics*, 13, 117–123.
- Jonsen, A. R. (1986). Bentham in a box: Technology assessment and health care allocation. *Journal of Law, Medicine and Ethics*, 14, 4–172.
- Kamm, F. M. (1998). *Morality, mortality*. Oxford: Oxford University Press.
- Kamm, F. M. (2007). *Intricate ethics: Rights, responsibilities, and permissible harm*. Oxford: Oxford University Press.
- Kant, I. (1785/1998) *Groundwork of the Metaphysics of Morals* (Cambridge: Cambridge University Press).
- McDowell, J. (1978). Virtue and reason. *The Monist*, lxii, 331–350.
- McKie, J., & Richardson, J. (2003). The rule of rescue. *Social Science and Medicine*, 56, 2407–2419.
- Nagel, T. (1986). *The view from nowhere*. Oxford: Oxford University Press.
- Nagel, T. (1991). *Equality and partiality*. Oxford: Oxford University Press.
- NICE. (2009). Appraising life-extending, end of life treatments. <http://www.nice.org.uk/media/E4A/79/SupplementaryAdviceTACEoL.pdf>. Accessed 15 April 2014.
- Nozick, R. (1974). *Anarchy, state, and utopia*. Oxford: Basil Blackwell.
- Ostrom, E. (2005). *Understanding institutional diversity*. Princeton, NJ: Princeton University Press.
- Scanlon, T. (1998). *What we owe to each other*. Cambridge, Mass.: Harvard University Press.
- Sheehan, M. (2007). Resources and the rule of rescue. *Journal of Applied Philosophy*, 24(4), 352–366.
- Williams, B. (1981). Persons, character and morality. In B. Williams (Ed.), *Moral luck* (pp. 1–19). Cambridge: Cambridge University Press.
- Wolff, J. (2011). *Ethics and public policy: A philosophical inquiry*. Abingdon, UK: Routledge Press.
- Wolff, J. & Orr, S. (2009). Cross-sector weighting and valuing of QALYs and VPFs: A report for the inter-departmental group for the valuation of life and health. <http://www.ucl.ac.uk/cpjh/docs/IGVLH.pdf>. Accessed 15 April 2014.