

Balancing small against large burdens

ALEX VOORHOEVE

Department of Philosophy, Logic and Scientific Method, London School of Economics, London, UK
Department of Bioethics, National Institutes of Health, Bethesda, MD, USA

Abstract: Common principles for resource allocation in health care can prioritise the alleviation of small health burdens over lifesaving treatment. I argue that there is some evidence that these principles are at odds with a sizable share of public opinion, which holds that saving a life should take priority over any number of cures for minor ailments. I propose two possible explanations for this opinion, one debunking and one vindicatory. I also outline how well-designed surveys and moral enquiry could help decide between them. Finally, I consider how priority-setting principles could be adjusted if the view that saving a life always trumps alleviating small burdens were vindicated.

Submitted 14 January 2017; accepted 23 January 2017

Introduction

In allocating public resources for health, how do people balance individually small benefits for a multitude against individually large benefits for a few? And are their judgements well-founded? These questions garnered widespread attention because of a well-known priority-setting exercise for Medicaid in the US state of Oregon in 1990.¹ A state commission proposed to prioritise health interventions according to cost-effectiveness, defined as cost per quality-adjusted life-year (QALY).² This resulted in tooth capping receiving higher priority than treatment for terminal appendicitis (Morell, 1990). To see the

Department of Philosophy, Logic and Scientific Method, London School of Economics, Houghton Street, London, UK. Email: a.e.voorhoeve@lse.ac.uk

1 Medicaid provides health services for the poor and is funded through a mix of state and federal contributions.

2 One QALY is an amount of health-related well-being that is just as valuable to a person as one year in full health. For example, a person gains one QALY by living one extra year in perfect health or by being cured of a health problem that would have imposed on them for the next ten years an annual health burden of 0.10. See Bognar and Hirose (2014) for an overview of the measures that have been developed for determining the impact of health on quality of life.

implications of this proposal, suppose that one held constant the state's Medicaid budget and spent it on interventions ordered from lowest to highest cost per QALY until the budget was exhausted. It would then be possible that tooth capping would be funded, but life-saving treatment for acute appendicitis would not be. This implication was a key contributor to public outrage at the proposal, which led Oregon to abandon the ranking of interventions by cost-effectiveness (Morell, 1990).

On a common view, a reason for this outrage was that the commission's ranking method failed to reflect people's view that a given increment in health-related well-being should get greater, but finite, extra weight when it would go to someone facing a large individual disease burden than when it would instead go to another person facing a small individual disease burden. This analysis is backed by surveys in which the median respondent gives such finite extra weight to the interests of those facing larger burdens (see Ubel *et al.*, 1996; Nord & Johansen, 2014). From this diagnosis, it follows that priority-setting would better align with the public's views if it were based on a severity-weighted cost-effectiveness function of the kind proposed by the Norwegian Committee on Priority Setting in the Health Sector, on which a QALY gained by someone facing a large health burden is three times more valuable (and can therefore permissibly cost three times more) than a QALY gained by someone facing a small health burden (Ottersen *et al.*, 2016).

Those who advocate this approach do so not merely in order to accommodate the public. They also believe that special concern for those facing larger health burdens is justified, for, they hold, an improvement in the fate of the worse off is especially valuable both because it improves total well-being and because it reduces unfair inequality (Ottersen *et al.*, 2016).

It is noteworthy that, on this view, the mere fact that tooth cappings were prioritised over a life-saving operation was not what generated justified outrage. Rather, the commission's mistake was to apply the wrong weights when ranking deaths from appendicitis against the loss in well-being due to uncapped teeth. But with the purportedly right weights, tooth cappings could still trump saving someone from acute appendicitis. For example, on the Norwegian Committee's proposal, if curing someone of acute appendicitis were more than three times more costly per QALY than tooth cappings, the latter would still be prioritised over the former.

In this article, I shall put forward an alternative analysis of people's objections to the Oregon Commission's recommendations. Just like the common diagnosis, my alternative interpretation involves both an empirical and a normative claim. The empirical claim is that a sizable share of people reject the idea that many very small burdens can jointly outweigh a few very large ones. The normative claim is that these people may well be right. This

alternative diagnosis has policy implications that contrast with the aforementioned Norwegian policy. For, if correct, it would lend support to a plan announced (but not yet implemented) by the Dutch government to eliminate public spending on the prevention and alleviation of ailments that impose only very small health burdens, independently of their cost-effectiveness,³ in order to focus public resources on averting more substantial burdens (CVZ, 2012a, 2012b; Voorhoeve, 2017a).

My argument proceeds as follows. In the ‘Survey evidence’ section, I consider empirical data on people’s attitudes. I argue that there is evidence that a substantial share of subjects believe that saving a young adult’s life trumps saving a multitude from very minor harms, no matter how numerous this multitude. In the remainder of the paper, I offer two competing explanations of this judgement. In the ‘A debunking explanation’ section, I consider the hypothesis that this judgement stems from people’s faulty intuitive appreciation of the moral significance of large numbers. I argue that while it has some initial plausibility, this debunking explanation, as yet, lacks empirical support. In the ‘A possibly vindicatory explanation’ section, I put forward an alternative hypothesis: that those who are unwilling to prioritise alleviating small harms over saving a life are motivated by a form of sympathetic identification with, and respect for, a person who has much more at stake than others. In the ‘How to test our competing hypotheses’ section, I explain how these two competing explanations are empirically distinguishable. I summarise my findings in the ‘Conclusion’ section.

Survey evidence

In this section, I review four studies that provide what is, to my knowledge, the best (although, as we shall see, imperfect) evidence available of people’s views on cases in which one must balance saving some from minor health burdens against saving others from great burdens.

The aforementioned study by Ubel *et al.* (1996) offered university students ($n = 42$) a choice between saving the lives of ten people who would otherwise die of appendicitis and saving a number of others from a lesser health burden. The latter ranged from a very small burden (a cyst on a tendon in one hand that was described as yielding a discernible bulge and infrequent, minor pain, but no functional limitations) to a sizable burden (a benign

³ Extant policy proposals make an exception for the treatments of minor ailments that save money in the long run, that would be administratively too burdensome to eliminate or that also prevent larger harms (CVZ, 2012a, 2012b).

meningioma – a tumour in the layers of tissue that surround the brain – described as causing very frequent, often debilitating headaches, but not affecting life expectancy). In each choice, subjects were asked which number of people cured of the illness causing the smaller burden would “equal the benefit” brought about by saving ten lives (Ubel *et al.*, 1996, p. 111). For our purposes, it is of interest what share of subjects responded to these questions with ‘no number’ or ‘infinity’. This information, which is represented in Table 1, can be inferred from the data in Ubel *et al.* (1996) if one supposes that every subject responding with ‘no number/infinity’ to the meningioma question also responded in this way to the cyst question.

These findings suggest that a substantial minority (35.7%) believed both: (a) that no number of cured cysts can outweigh saving several lives; and (b) that curing a large number of meningioma cases should take priority over saving lives. This combination of responses is interesting because it is not explicable in terms of familiar theories of distributive justice. In holding that curing a number of sizable ailments (less severe than death) can take priority over saving lives, these subjects reject the leximin principle of distributive justice (this principle holds that one ought to first maximally improve the least well-off position and then proceed on to maximally improve the second least well-off position, etc.). But they also reject well-known alternatives to leximin. A common interpretation of utilitarianism in the context of priority setting in health is that we ought to maximise total health-related well-being in a given population. Obviously, on this view, curing a sufficiently large number of cysts outweighs saving ten lives. As mentioned in the introduction, the same is true of a standard version of pluralist egalitarianism, on which benefits to the badly off receive finite extra weight; it is just that, on this egalitarian view, the number of cysts required to outweigh a life is greater than under utilitarianism (Ottersen *et al.*, 2016; see also Fleurbaey & Tungodden, 2010).

Table 1. Share of responses to paired questions from Ubel *et al.* (1996) ($n = 42$)

		Meningioma case	
		Answer is a natural number	Answer is ‘no number/infinity’
Which number of people cured of the condition would yield a benefit equal to saving ten lives?			
Cyst case	Answer is a natural number	59.5%	0 (by assumption)
	Answer is ‘no number/infinity’	35.7%	4.8%

Note: The bold percentage represents the share of respondents who were willing to let a multitude of sizable harms, but no number of minor harms, outweigh ten lives.

We should note, however, that the study has two shortcomings: first, it involved a small, non-representative group of subjects; and second, respondents were asked what number of lesser ailments averted would generate the same benefit as averting ten deaths. But one can coherently believe that it is impermissible to save a great many people from the harm of a cyst rather than ten from death even when the former would generate a greater total benefit. After all, on non-consequentialist moral views, it is not always permissible to do what will produce the greatest total benefit. One should therefore ask not about *total benefits* but about *permissible action*: for which number, if any, of cured cysts it would be permissible to cure the cysts rather than save ten people's lives. The study may therefore underestimate the share of people who would hold that saving lives always takes priority over curing cysts.⁴

Pinto-Prades and Lopez-Nicolas (1998) get somewhat closer to the correct framing by asking what types of interventions people would prefer (although it would have been even better had the study asked what it would be *permissible to choose*). Subjects (Spanish students, $n = 83$) were told they could either establish a neonatal care unit that would save the lives of ten newborns or implement a policy that would treat a "very large number (e.g. 100,000)" of others for an ailment that causes a health burden less bad than death. The latter burden varied depending on the question; the smallest burden was living with moderate pain and discomfort while still being able to perform all of one's ordinary activities.

Unfortunately, because the paper reports only the median response for each choice, one can draw conclusions only about the majority answer in each case. We learn that a majority chose to save the ten neonates' lives rather than treat the moderate pain of a "very large number" of people and that a majority was willing to alleviate the large burdens (still short of death) of a multitude rather than save the lives of ten neonates (Pinto-Prades & Lopez-Nicolas, 1998, p. 290). All we can conclude, therefore, is that at least some subjects must have displayed the combination of attitudes in question: that averting a small number of deaths trumps alleviating a multitude of small burdens, but that alleviating numerous quite large burdens takes priority over averting a small number of deaths.

Moreover, even this evidence is weak because of shortcomings in the study. First, it would have been better had there been no suggested upper bound on

⁴ Indeed, Damschroder *et al.* (2007, p. 270) report that subjects are far more likely to respond that 'no number' of cures for minor ailments can permissibly be prioritised over a cure for ten severe impairments ('choice framing') than they are to say that curing 'no number' of minor ailments yields a benefit equivalent to curing ten cases of a very serious ailment ('benefit framing').

the size of the affected population. Second, there was a lack of clarity about the effects of the possible interventions on lifetime health. No mention was made of the quality of life the neonates would have if saved from death. In addition, the effect of the treatment for the lesser conditions was left unstated (Pinto-Prades & Lopez-Nicolas, 1998, p. 294). Third, the use of neonates is not ideal because some believe that, due to their as-yet underdeveloped cognitive capacities, their claim to life is less strong than that of a teenager or young adult (McMahan, 2002).

In contrast to the preceding studies, Damschroder *et al.* (2007, Study 1) draws on a relatively large set of subjects chosen to represent the US population ($n = 827$). Participants were asked to trade off saving a small number from a greater individual burden against saving a larger number from lesser burdens. The largest burden considered was quadriplegia and the smallest was paralysis in one foot. In a choice between curing ten quadriplegics and up to one million people of foot paralysis, 40% of respondents said that they would always cure the quadriplegics. By contrast, in a choice between curing ten people of quadriplegia and up to one million paraplegics, only 7% said that they would always cure the quadriplegics.⁵ As a follow-up, respondents who chose to aid the quadriplegics rather than up to a million people with a lesser health burden were asked whether there was a number of people in the world such that they would save that number from the lesser burden rather than the ten quadriplegics. A large majority (59%) replied that there was no such number (Damschroder *et al.*, 2007, p. 270). If we assume this share applies uniformly across all such answers and (as seems sensible) that no one who chose to cure ten quadriplegics rather than a million paraplegics also chose to cure a million from foot paralysis rather than ten quadriplegics, then we can compute the proportions listed in Table 2. These suggest that around a fifth of subjects (19.5%) endorse the idea that even billions of cases of foot paralysis cannot permissibly take priority over ten cases of quadriplegia, but that a multitude of cases of paraplegia can do so. I submit that this should be treated as a lower bound on the share of respondents who hold that no number of small harms can outweigh a very large harm, but that some number of considerable harms can do so. For the gap in severity (from foot paralysis to quadriplegia) is less great than it might be; even the smaller harm (foot paralysis) is quite substantial – far more severe, for example, than the cyst in Ubel *et al.* (1996).

⁵ These numbers are deduced from the odds ratios in Damschroder *et al.* (2007, Table 4) in combination with the report that 17% of all answers given were “off scale refusals” (Damschroder *et al.*, 2007, p. 270).

Table 2. Share of responses to paired questions inferred from Damschroder *et al.* (2007) ($n = 827$)

For which number of people cured of the condition would you cure that condition rather than ten quadriplegics?		Paraplegia case	
		Answer is less than 'the Earth's population'	Answer is 'no number in the Earth's population'
Foot paralysis case	Answer is less than 'the Earth's population'	76.4%	0 (by assumption)
	Answer is 'no number in the Earth's population'	19.5%	4.1%

Note: The bold percentage represents those willing to let a multitude of cases of paraplegia, but no number (short of the number of people in the world) of cases of foot paralysis, outweigh ten cases of quadriplegia.

The final study I shall review concerns income rather than health.⁶ The authors asked subjects (British and Norwegian students, $n = 642$) the following question (Cowell *et al.*, 2015, p. 771):

“Suppose one poor person benefits from an income increase of £ G while all the rich, no matter how many there are, suffer an income reduction of £1. If G were large enough, would this be a good idea?”

Subjects were also presented with a follow-up question about their adherence to the leximin principle applied to income. Overall, 33.0% of subjects both: (a) endorsed imposing a small monetary loss on the better off, no matter how numerous, to greatly improve the income of a single badly off person; and (b) rejected the leximin principle.

This may be taken to indicate the support of around a third of subjects for the view at issue. This indicator is imperfect, however, for two reasons: first, a loss of £1 could be regarded as yielding no appreciable loss in well-being to a rich person. Consequently, the responses might provide an overestimate of support for the idea that no number of small but appreciable benefits to

⁶ I shall not discuss in detail a number of further studies that, while suggestive, do not address our central issue. Choudhry *et al.* (1997), Olsen (2000), Rodriguez-Miguez and Pinto-Prades (2002) and Hukin and Tsuchiya (2005) all find that, when the gap in the size of individual benefits is very large, people tend to ‘concentrate’ a given sum of benefits among a few rather than ‘disperse’ them among many. Schumacher *et al.* (2017) find that when the gain to a recipient is large relative to the cost to those who pay for this gain, many subjects are insensitive to the number of payers. However, their maximum number of payers is 32. None of these studies test whether subjects endorse the idea that a large benefit to one otherwise badly off person would outweigh *any* number of small benefits to better-off people.

the better off can outweigh one very large benefit to a badly off person. Second, the survey questions in Cowell *et al.* (2015) suggest the imposition of a loss on some for the sake of a gain to others. Some subjects may therefore have conceived of this as a question about forcible transfers from some for the sake of another. An aversion to such transfers (e.g. on grounds that they constitute an interference with people's liberty) might therefore lead to an underestimation of support for giving priority to large benefits for a badly off person over small benefits to a multitude of better-off people.

Beyond the issues already raised, the studies reviewed all face the following challenge. These surveys pose questions with which respondents are unacquainted. Subjects have little opportunity for deliberation. Unsurprisingly, therefore, inconsistent responses are common (Ubel *et al.*, 1996; Damschroder *et al.*, 2007). Consequently, the answers elicited are, at best, imperfect indicators of people's considered commitments.

There is, in sum, a strong case for better-framed questionnaire research conducted jointly by behavioural scientists and moral philosophers. Still, despite their shortcomings, extant studies suggest that at least a substantial minority of subjects endorse both: (a) that one ought to save one life rather than spare any number of the better off a minor harm; and (b) that one ought to prioritise saving a very large number of people from sizable harm over saving one life.

A debunking explanation

What might explain the judgements of this group? In their work on our attitudes to genocide, Paul Slovic and Daniel Västfjäll provide a first step towards a hypothesis, proposing that “we may not ‘feel’ much difference, nor [much] value the difference, between saving 87 [people from a given harm] and saving 88” (Slovic & Västfjäll, 2013, p. 98). This remark suggests that, intuitively, the perceived marginal moral value of the number of averted harms of a given size is positive, but declining. However, to explain the judgements in question, we would need to assume that the perceived total value of averting a number of harms has an upper limit and that this limit is a positive function of the size of the harms averted. Figure 1 illustrates the requisite perceived value function for the cases from Ubel *et al.* (1996). In the figure, the perceived value of curing a number of cysts asymptotically approaches V_C as the number tends towards infinity. The perceived value of curing a number of meningiomas approaches V_M as the number tends towards infinity. If the felt value of curing ten cases of terminal appendicitis lies in between these two limits, then we arrive at the to-be-explained intuitive verdict that no number of cured cysts, but some number of cured meningiomas, outweighs the value of saving ten lives.




Figure 1. A possible explanation of the belief that a number of meningiomas, but no number of cysts, can permissibly outweigh ten lives

Now, Slovic and Västfjäll do not offer direct support for their “declining marginal perceived significance of numbers” hypothesis.⁷ They do, however, draw attention to the common finding that we display diminished sensitivity to changes in a perceptual stimulus (such as sound, brightness or temperature) as the magnitude of the stimulus increases. They also suggest that our sense of the significance of numbers may display a similar pattern.

Interestingly, some work on our intuitive numbers sense appears to support their conjecture. For example, we perceive a pair of numbers that differ by a given amount as lying ‘closer together’ on an intuitive ‘mental number line’ when these numbers are larger – so that, say, 10 is perceived as further from 20 than 90 is from 100 (Longo & Lourenco, 2007). Another way of putting this is that, at least in our untrained perception, a given increase in number is more significant when it takes place from a smaller base. For example, when asked to place numbers on a line with 0 and 10 as endpoints, kindergartners and members of Amazonian tribes without a formal education put 3 in the middle, suggesting that an increase from 0 to 3 is perceived as just as significant as an increase from 3 to 10 (Varshney & Sun, 2013, p. 28).

There are, however, some reasons to be sceptical of the hypothesised explanation of responses to our trade-off cases, for the way we intuitively map numbers seems to depend on education. Fourth-graders, for example, place 3 at three-tenths of the way between 0 and 10 (Varshney & Sun, 2013,

⁷ Slovic and Västfjäll (2013, pp. 97–98) cite data that they claim support the “diminishing marginal perceived value of averting harms” hypothesis, but in fact the cited studies support a *different* hypothesis, viz. that the perceived value of saving a number of people from harm depends on the *share* of the at-risk population of which they are a part, so that, say, saving 1000 people out of a population of 2000 is perceived as more valuable than saving 1000 out of a population of 10,000.

p. 28). The ‘mental number line’ of the educated subjects in the surveys reviewed above may not, therefore, have exhibited the form of diminished marginal sensitivity in question. But even if it did, the literature on our numbers sense does not appear to offer evidence in favour of the idea that the total perceived value of saving a number of people from a given harm has an upper limit, for it suggests that our untrained numbers sense is approximately logarithmic, and logarithmic functions do not have an upper bound (Dehaene, 2011, pp. 64–66, 265–266). Finally, there is a need to motivate a step from how we *represent* numbers on a mental line to the way we *value* quantities of people saved. I conclude that while the hypothesis has some plausibility and is therefore worth investigating, there are, as yet, important lacunae in the case for it.

Suppose, however, that it were to explain the judgements in question. Slovic and Västfjäll mention grounds for concluding that these judgements are unreliable, for, they argue, the perceived diminishing marginal importance of saving an additional person from harm is a *misperception*. The equal importance of every person’s life, they claim, implies that we should value every additional life saved as much as the preceding one – saving 101 rather than 100 people generates just as much additional moral value as saving three rather than two (Slovic & Västfjäll, 2013, p. 97). If we extend their reasoning from lives saved to other benefits, then it follows that providing n people with a benefit of a given size is n times as important as providing one person with this benefit, so that (assuming that all benefits have finite value) some number of small individual benefits can take priority over saving one life.

This aggregative principle has some plausibility. It is, of course, endorsed by familiar views of distributive justice in health, including the standard form of cost-effectiveness analysis initially applied in Oregon, as well as the form of equity-weighted cost-effectiveness analysis now proposed in Norway. However, given the lack of evidence for the hypothesis that the judgements in question are based on the perception that the value of saving a number of people from a given harm has an upper limit, we should also consider an alternative explanation of these judgements. I shall now do so.

A possibly vindicatory explanation

The alternative I propose is that at least some individuals employ the following two-step procedure when making trade-offs. First, they compare how much is at stake for each individual on either side (e.g. being cured of a cyst versus having one’s life saved). If what is at stake for a person on one side pales in comparison with what is at stake for a person on the other side, then the claims of those who have much more at stake win out, irrespective of the

numbers involved. However, if the gap between what is at stake for individuals on either side is not too great, then the numbers on either side (as well as the size of the individual benefits at issue) count; indeed, every additional person benefited matters just as much as the preceding person.

More precisely, for the relatively straightforward type of cases we have been considering, the view I have just sketched can be put as follows (Voorhoeve, 2014, p. 66)⁸:

- (1) Every person whose well-being is at issue has a claim.
- (2) The strength of a person's claim increases:
 - (2.a) the larger the magnitude of their feasible gain in well-being; and
 - (2.b) the less well-off they are compared to others.⁹
- (3) People's claims compete just in case their claims cannot be jointly satisfied.
- (4) A claim is relevant just in case it is sufficiently large when compared one-to-one to the strongest claim with which it is in competition.
- (5) The decision-maker should satisfy the greatest sum of strength-weighted, relevant claims.

I shall refer to this view as Aggregate Relevant Claims (ARC). This view incorporates an aggregative moral principle of the kind championed by Slovic and Västfjäll (2013), but it constrains it by a non-aggregative principle that compares competing claims one-to-one. What grounds the latter principle and where does it draw the line between relevant and irrelevant claims?

I submit that the non-aggregative approach originates in a form of sympathetic identification (Voorhoeve, 2014, 2017b). While in everyday English the term 'sympathy' often signifies feeling along with another's sorrow, I shall here use the term in the sense famously articulated by Adam Smith in *The Theory of Moral Sentiments*, in which sympathy involves placing oneself in another's position with a set of standardised, or proper, attitudes (Smith, 1982). This form of sympathy is, at its heart, anti-aggregative. One can imaginatively place oneself in another person's shoes, but one cannot do so for a group, except by placing oneself in the shoes of a representative individual or by taking up the position of each person in turn, one at a time. (This explains why research suggests that moral reasoning that relies on sympathetic identification alone does not take account of numbers; see Kogut & Ritov, 2005.)

Drawing on this form of sympathy, a non-aggregative approach to distributive justice involves imaginatively taking up each person's legitimate perspective on the situation, one person at a time. Call a person whose perspective

⁸ While in what follows I offer my own version of this view, there are others (e.g. see Kamm, 1993, Chapters 8–10; Kamm, 2007, pp. 484–486; Temkin, 2012, Chapters 2 and 3).

⁹ I owe this way of determining the strength of a person's claim to Lange (2017).

we are taking up *P*. Taking up *P*'s legitimate perspective means taking on their natural extra concern for themselves up to the degree that this is morally permissible. It also involves giving the minimally required weight to the interests of others.¹⁰ (In line with what I take to be common-sense morality, I shall assume that it is permissible to have a substantially greater degree of concern for oneself than for a stranger.) From this permissible personal perspective, we then rank the importance of the competing claims one-to-one; that is to say, we compare the maximum value *P* can permissibly assign to the satisfaction of *P*'s own claim to the minimum value *P* can permissibly assign to the satisfaction of the stranger's competing claim.

Assuming that this ranking is complete, there are two possibilities: one is that, from *P*'s legitimate perspective, it is at least as important to satisfy *P*'s claim as the stranger's claim. Given the fact that *P* may give additional weight to their own interests, this will occur not merely when *P* has more at stake than another with whom their claims compete, but also when they have somewhat less at stake. In this case, we might say that while comparing claims one-to-one, we can sympathise with *P* pressing their claim, because *P* can permissibly value the benefit they might receive at least as much as the benefit that the other person might receive.

The other possibility is that so much more is at stake for the other person that even from *P*'s permissible perspective the satisfaction of the other's competing claim is valued more highly. In this case, if we confine ourselves to one-to-one comparisons, we cannot sympathise with *P* pressing their claim, because even *P* must in good conscience value the stranger's claim over their own.

To illustrate these contrasting possibilities, consider again the cases from Ubel *et al.* (1996). Suppose first that *P* has a debilitating, painful meningioma and the strongest competing claim on behalf of a stranger is to life-saving treatment for appendicitis. It seems within the bounds of permissible self-concern for *P* to value their own cure of the meningioma over a stranger's life. (This manifests itself in a number of ways. For example, it would, I submit, be permissible to use one's own money to cure oneself of a meningioma rather than to purchase life-saving treatment for a stranger who would die without one's assistance.) By contrast, it would be indecently selfish to value a cure of the aforementioned cyst on one's hand over a stranger's life.

¹⁰ Following Smith (1982, especially Sections I.I.III–IV and II.II.II), I focus on a moralised form of sympathy that involves placing oneself in another's shoes whilst assuming morally permissible desires and emotions. Because people's actual sentiments may not be permissible ones, sympathy, so conceived, need not involve feeling along with people's actual sentiments. The idea of taking up a person's permissible personal perspective in the course of balancing people's interests is also prominent in Kamm (1993, Chapters 8–10).

Now, of course, this process of imaginative identification always yields fellow-feeling with a desire of a person with the strongest claim to see their claim satisfied. After all, their legitimate partiality towards themselves merely reinforces a claim that is already stronger from an impartial point of view. It follows that this process of identification with each person's perspective may lead us to sympathise with the claims of opposing sides – as in the case of the meningioma versus appendicitis trade-off. In such a scenario, the process of placing oneself in each person's position yields a vivid sense of conflicting interests. However, we may also find that we sympathise with the claims of only one party – as in the cyst versus appendicitis trade-off. In the latter circumstances, the non-aggregative approach yields not discord, but concord: from every person's legitimate perspective, saving one from life-threatening appendicitis should take priority over curing a cyst.

I submit that this distinction may explain the views of those subjects in Ubel *et al.* (1996) who held that no number of cysts, but some number of meningiomas, could outweigh a life. In the former case, the non-aggregative approach yields a clear, univocal verdict: save the life. By contrast, in the latter case, the non-aggregative approach yields disagreement; it is impossible to satisfy everyone's legitimate claims, so some other approach is necessary for deciding whom to help. An approach that suggests itself is to employ an aggregative moral principle of the kind proposed in Slovic and Västfjäll (2013, p. 97).

The criterion for relevant claims that emerges from this discussion, then, is this: a claim is relevant if and only if it is permissible for its possessor to value the satisfaction of their claim over the satisfaction of the strongest competing claim when these claims are compared one-to-one. By employing this criterion, ARC respects a form of unanimity that arises when one takes up each person's point of view. But when such imaginary identification does not yield unanimity, ARC appeals to a form of weighted majority voting in which the weight of each person's vote is the strength of their claim.

ARC, then, is based on the idea that both the outlined non-aggregative approach and the proposed aggregative approach have some validity and that one therefore ought to attempt to mediate between them. Thomas Nagel (1979, pp. 118 and 123) expresses this idea as follows:

“Both the separate [non-aggregative] and the conglomerate [aggregative] methods count everyone fully and equally. The difference between them is that the second moves beyond individual points of view to something more comprehensive than any of them, though based on them. The first stays closer to the points of view of the individuals considered. ... It is obvious that [these] conceptions of moral equality ... are extremely different. They seem to be radically opposed to one another, and it is very difficult to see

how one might decide among them. My own view is that we do not have to. A plausible social morality will show the influence of [both of] them.”

It seems to me that if people’s objection to prioritising minor over major harms were grounded in ARC, then it would be hard to dismiss it as a product of unreflective, biased judgement. Nonetheless, as limned here, ARC only deals with a very limited range of decision problems.¹¹ It therefore needs further development and critical examination before one can confidently endorse it or something like it.¹²

How to test our competing hypotheses

Having put forward competing explanations of our target judgements, I shall now outline two ways in which these proposed explanations are empirically distinguishable.

First, only the hypothesis that people adhere to ARC predicts that there will be a link between answers to the following questions (Voorhoeve, 2014, Section VIII): (i) For a benefit of size S , is there a number of people one can provide with that benefit such that one ought to aid them rather than save a life? (ii) Does one have a personal prerogative¹³ to secure a benefit of size S for oneself rather than save a stranger’s life? Now, it seems to me that, in everyday morality, the boundaries of the latter prerogative are not precisely delineated; indeed, some have argued that they are irremediably imprecise (Parfit, 2011, Chapter 6). Nonetheless, there are circumstances in which it is clear that one may favour oneself; as mentioned, the meningioma versus appendicitis case in Ubel *et al.* (1996) seems to be such a case. There are also circumstances in which the balance of interests is such that one may not favour one’s dear self; the cyst versus appendicitis case seems to be an example. If the proposed version of ARC is what motivates people, then these circumstances will correspond, respectively, to cases in which they think it is permissible to prioritise many lesser harms over a single greater harm and to cases in which they think this is impermissible.

11 For example, as formulated, ARC says nothing about cases involving risk or cases in which people on ‘one side’ have claims of differing strength (e.g. when one can either save ten lives or, instead, cure both some of their meningiomas and others of their cysts).

12 For some such work, see Temkin (2012, Chapters 2 and 3), Kelleher (2014), Voorhoeve (2014, 2017b), Gustafsson (2015), Halstead (2016), Badano (2016), Tomlin (forthcoming) and Privitera (2017).

13 Following Kamm (2007, Chapter 1), I assume this prerogative operates within the constraints posed by people’s rights. One is therefore to imagine that, in exercising it, one makes use only of one’s own resources and does not violate others’ rights.

The second empirical test is this: an adherent of ARC will hold that whether a particular claim is relevant depends on the strongest competing claim. They can therefore display a distinctive pattern of pairwise choices. To illustrate, when choosing only between saving one young person's life and curing a thousand moderate ailments, ARC can hold that one must do the latter, because the thousand's claims are relevant and together outweigh the single person's claim. Moreover, when choosing only between curing a thousand moderate ailments and curing a million minor ailments, ARC can require the latter, because the million minor ailments are relevant in this context and they together outweigh the thousand moderate ailments. Finally, when choosing only between saving one person's life and curing one million of the minor ailment, ARC can mandate saving the life, because the claims to be cured of the minor ailments are irrelevant in this context.¹⁴

By contrast, such a pattern of choice could not arise if people's choices were based on those outlined in the 'A debunking hypothesis' section above, for on the latter view, the value of an intervention does not depend on the set of feasible alternatives. If, for a subject, curing a thousand moderate ailments were to take priority over saving one life and curing a million minor ailments were to take priority over curing a thousand moderate ailments, then the subject would have to judge that curing a million minor ailments takes priority over saving a life, if this hypothesis were true.

Conclusion

We have examined two questions: first, how do people believe a public body should trade off alleviating small health burdens for many against alleviating very large burdens for a few? Second, are their opinions well-founded?

In response to the first question, I have argued that the available evidence is of surprisingly low quality, in part because none of the studies reviewed poses questions in quite the right way. Nonetheless, these studies indicate that at least

¹⁴ While this pattern might appear to violate transitivity, such violation is merely apparent (Voorhoeve, 2014, Section VI), for, as argued in Broome (1991, Chapter 5), in judging whether a moral view violates formal principles of rational choice, two alternatives should be taken to be identical just in case, on the view, they share all pertinent features. But, according to ARC, saving a million from a minor harm when the only alternative is saving a thousand from moderate harm is importantly different from saving a million from a minor harm when the only alternative is curing one young person of a terminal illness, for in the former case it considers the minor harms to be relevant, whereas in the latter case it does not. To put it differently: in the former case, on ARC, sparing the million from a minor harm does not disrespect a person whose life is at stake, whereas in the latter case it does (Kamm, 2007, pp. 484–487). And, once alternatives are properly identified, the outlined pattern of choice does not violate transitivity.

a sizable minority believe that averting a very large burden should take priority over averting any number of minor burdens.

My answer to the second question has focused only on this group. I have put forward two competing hypotheses. On one hypothesis, such judgements are due to the perception that the total value of bestowing a given benefit on people approaches an upper bound as the number of people benefited tends towards infinity. If true, this explanation would undermine the validity of people's judgements because the hypothesised perception fails to capture the equal value of meeting each person's health needs.

By contrast, on another hypothesis, the judgements in question are motivated by sympathy with, and respect for, a person who has much more at stake than anyone with whom their interests compete. I have also argued that if this is indeed what motivates people, then their resistance to granting numerous small benefits priority over saving a life may well be reasonable.

While I hope to have made some progress, one conclusion is that we lack adequate empirical knowledge about what people's attitudes towards the relevant trade-offs are and what motivates them. I have indicated a number of ways in which collaborative work by behavioural scientists and moral philosophers could improve on extant studies. Such work is clearly relevant for policy-makers, for understanding the public's views is key to setting priorities in a way that engages with people's concerns. To return to our opening policy case study: suppose that a large share of the public reject Oregon-style rankings because they believe that tooth cappings must never be prioritised over curing an adult's terminal illness. Such people will not be inclined to accept priority-setting principles such as the ones proposed in Norway that merely give some limited extra weight to saving a life. Of course, these people's judgements may be based on a misperception of the importance of the number of people that can be benefited; in that case, the right response would be to attempt to persuade them of their vulnerability to this bias. However, if their objection is grounded in respect for the especially strong claims of a person who would, if aided, be cured of a terminal illness, then such an argument will probably not persuade them. Indeed, I have argued such a view appears to have considerable merit. If, on further examination, it proves to be a reasonable view, then it can motivate the contrasting policy under development in The Netherlands, which is to withdraw public funding from very minor ailments in order to focus it on larger burdens.

Acknowledgements

I am grateful to Ipek Gençsü, Adam Oliver and Trygve Ottersen for helpful discussion.

Financial support

This research was supported by the Intramural Research Program of the US National Institutes of Health, Department of Bioethics.

Disclaimer

The opinions expressed are the views of the author only. They do not represent any position or policy of the National Institutes of Health, the Public Health Service or the Department of Health and Human Services.

References

- Badano, G. (2016), 'On the moral importance of numbers, relevance, and the workings of non-aggregative reasoning', *Ethical Perspectives*, 23(3): 527–43.
- Bognar, G. and I. Hirose (2014), *The Ethics of Health Care Rationing: An Introduction*, Abingdon: Routledge.
- Broome, J. (1991), *Weighing Goods*, Oxford: Blackwell.
- Choudhry, N., P. Slaughter, K. Sykora, C. D. Naylor (1997), 'Distributional dilemmas in health policy: large benefits for a few or smaller benefits for many?', *Journal of Health Services Research and Policy*, 2(4): 212–16.
- Cowell, F., B. Tungodden, and M. Fleurbaey (2015), 'The tyranny puzzle in social preferences: An empirical investigation', *Social Choice and Welfare*, 45: 765–92.
- CVZ (2012a), *Uitvoeringstoets lage-ziektelastbenadering*, Diemen: CVZ.
- CVZ (2012b), *Technische rapportage ziekteelast*, Diemen: CVZ.
- Damschroder, L., T. R. Roberts, B. J. Zikmund-Fisher, and P. A. Ubel (2007), 'Why people refuse to make tradeoffs in person tradeoff elicitation: A matter of perspective?', *Medical Decision Making*, 27: 266–80.
- Dehaene, S. (2011), *The Number Sense: How the Mind Creates Mathematics, revised and updated edition*, New York: Oxford University Press.
- Fleurbaey, M. and B. Tungodden (2010), 'The tyranny of non-aggregation versus the tyranny of aggregation in social choices: a real dilemma', *Economic Theory*, 44: 399–414.
- Gustafsson, J. (2015), 'Sequential dominance and the anti-aggregation principle', *Philosophical Studies*, 172(6): 1593–601.
- Halstead, J. (2016), 'The numbers always count', *Ethics*, 126: 789–802.
- Hukin, A. and A. Tsuchiya (2005), Dispersion versus concentration of health benefits: preliminary report. HESG working paper. <http://hesg.org.uk/wp-content/plugins/papers-listing/public/> (accessed December 5, 2016).
- Kamm, F. M. (1993), *Morality, Mortality, Vol. I*, Oxford: Oxford University Press.
- Kamm, F. M. (2007), *Intricate Ethics*, Oxford: Oxford University Press.
- Kelleher, P. (2014), 'Relevance and non-consequentialist aggregation', *Utilitas*, 26: 385–408.
- Kogut, T. and I. Ritov (2005), 'The 'identified victim' effect: An identified group, or just a single individual?', *Journal of Behavioral Decision Making*, 18: 157–67.
- Lange, B. (2017), 'Restricted prioritarianism or competing claims?', *Utilitas*, 29: 137–52.
- Longo, M. R. and S. F. Lourenco (2007), 'Spatial attention and the mental number line: evidence for characteristic biases and compression', *Neuropsychologia*, 45: 1400–6.
- McMahan, J. (2002), *The Ethics of Killing: Problems at the Margins of Life*, Oxford: Oxford University Press.

- Morell, V. (1990), 'Oregon puts bold health plan on ice', *Science*, **249**: 468–71.
- Nagel, T. (1979), 'Equality', in his *Mortal Questions*, Cambridge: Cambridge University Press, 106–27.
- Nord, E. and R. Johansen (2014), 'Concerns for severity in priority setting in health care: A review of trade-off data in preference studies and implications for societal willingness to pay for a QALY', *Health Policy*, **116**: 281–8.
- Olsen, J.A. (2000), 'A note on eliciting distributive preferences for health', *Journal of Health Economics*, **19**: 541–50.
- Ottersen, T., R. Førde, M. Kakad, A. Kjellevold, H.O. Melberg, A. Moen, Å. Ringard, and O. F. Norheim (2016), 'A new proposal for priority setting in Norway: Open and fair', *Health Policy*, **120**: 246–51.
- Parfit, D. (2011), *On What Matters, Vol. 1*, Oxford: Oxford University Press.
- Pinto-Prades, J.-L. and A. Lopez-Nicolas (1998), 'More evidence of the plateau effect: A social perspective', *Medical Decision Making*, **18**: 287–94.
- Privitera, J. (2017), 'Aggregate relevant claims in rescue cases?', *Utilitas*, in press.
- Rodriguez-Miguez, E. and J.-L. Pinto-Prades (2002), 'Measuring the social importance of concentration or dispersion of individual health benefits', *Health Economics*, **11**: 43–53.
- Schumacher, H., I. Kesternich, M. Kosfeld, and J. Winter (2017), 'One, two, many – Insensitivity to group size in games with concentrated benefits and dispersed costs', *Review of Economic Studies*, **84**: 1346–77.
- Slovic, P. and D. Västfjäll (2013), 'The more who die, the less we care: psychic numbing and genocide', in A. Oliver (ed.), *Behavioural Public Policy*, Cambridge: Cambridge University Press, 94–109.
- Smith, A. (1982) [1790], *The Theory of Moral Sentiments*, D. D. Raphael and A. L. Macfie (eds), Indianapolis: Liberty Fund.
- Temkin, L. (2012), *Rethinking the Good*, Oxford: Oxford University Press.
- Tomlin, P. (forthcoming), 'On limited aggregation', *Philosophy and Public Affairs*.
- Ubel, P., G. Loewenstein, D. Scanlon, and M. Kamlet (1996), 'Individual utilities are inconsistent with rationing choices: A partial explanation of why Oregon's cost-effectiveness list failed', *Medical Decision Making*, **16**: 108–16.
- Varshney, L. R. and J. Z. Sun (2013), 'Why do we perceive logarithmically?', *Significance*, **10**(1): 28–31.
- Voorhoeve, A. (2014), 'How should we aggregate competing claims?', *Ethics*, **125**: 64–87.
- Voorhoeve, A. (2017a), 'Healthy nails versus long lives: an analysis of a Dutch priority-setting proposal. Working paper available on <http://personal.lse.ac.uk/voorhoev>
- Voorhoeve, A. (2017b), 'Why one should count only claims with which one can sympathize', *Public Health Ethics*, **10**: 148–56.