

1 The Priority Idea

1.1 Dolores and Felicity

Dolores has had a rough life. Very rough. She has been homeless, on and off, for much of her adult life. She has been mentally ill since adolescence, and her illness makes her charmless and grasping, so she does not attract much sympathy from those who observe her plight. Family members help her occasionally but find helping her unrewarding, so they don't do anywhere nearly as much for her as they know they should. She has bounced in and out of halfway houses and government-supplied apartments for the disabled, always violating whatever rules are in force and chafing at restrictions on her freedom. She prizes her freedom but lacks the capacity to use it to her advantage in any way. She has lived on a meager income grant for disabled individuals. She has few resources, but even worse, very little capacity to transform whatever resources she has into well-being.¹

In contrast, consider Felicity. She has been extremely fortunate in her life and is looking forward to a happy old age. She has been enormously successful in business, amassing not only stupendous wealth but also great creative achievements of the sort any of us, in sensible moods, would want for ourselves. She has been blessed in family life and friendship, and when not working or engaged productively with family and friends, plays hard and with zest.

Felicity has already reached a very high level of lifetime well-being and is headed to further peaks beyond that ridgeline. As already mentioned, she has secure and stable command of immense resources that she can use for her own benefit. Moreover, part of her charm and good fortune is that she has immense capacity to turn resources into personal fulfillment (well-being or welfare). Give her a new yacht, she beams and will gain great benefit from this change in her life, notwithstanding the fact that she already owns around 100 magnificent yachts.

Now imagine you happen to face a rare opportunity, at a small cost to yourself, to do a significant good turn for one of two people ready at hand – as it happens, to either Dolores or Felicity. You can give one or the other a large bowl of delicious gelato that will otherwise go to waste. You have no other viable alternatives. There are no further complexities or nuances in the situation. Nobody will notice what you choose or what follows. There will be no significant long-term or indirect effects of giving the ice cream to Felicity or alternatively to Dolores. Although Felicity has recently had some fine ice cream, she would, being Felicity, gain a significant well-being boost

¹ The description of Dolores applies to a real person. (“Dolores” is not her real name.)

from getting this ice-cream gift from you. Dolores would get some benefit, but smaller. She has bad teeth, which tingle with something sweet in her mouth, diminishing the enjoyment she gets from the ice cream. She will for sure drop the bowl or get distracted and only consume a fraction of what you give her. Your action in giving the ice cream to one or the other will give rise to no effects except to bring about a one-unit enjoyment and well-being gain for Dolores, if you give it to her, or to bring about a three-unit enjoyment and well-being gain for Felicity, if you give it to her. To put the numbers in perspective, assume Felicity's lifetime well-being without this benefit will be 10,000 and Dolores's 800. What should you do?

1.2 Some Clarifications and a Statement of the Priority View

I admit this is a far-fetched and trivial scenario, but bear with me. I'm just making a point. Notice first of all that the influential and prominent doctrine called *utilitarianism*, as a criterion of morally right action, will definitely yield the judgment that you ought to give the ice cream to Felicity, not Dolores. Utilitarianism says, one ought always to do whatever would bring about the best reachable outcome, and in this small decision problem, the difference your action will make is, all things considered, either to bring about a one-unit well-being gain or alternatively a three-unit gain. Utilitarianism says you should do what will bring about the greater gain and thereby the better outcome. Utilitarianism tells us that in deciding what to do, from a moral perspective, one ought to pay no heed to the distribution of well-being gains across persons, just the sum total.

Some take this distribution-free character to be a black mark against utilitarianism. However, according to many more who have considered utilitarianism, the fact that it turns a blind eye to issues of distribution is among the least of its defects. These further supposed defects are beyond the topic of this Element, which focuses on fair distribution issues. Two fair distribution issues loom large. One comes out in the Dolores and Felicity parable: we should be maximizing, not the sum of well-being, but a transformation of this well-being sum that gives priority to gains for those who are worse off. Utilitarianism has built into it no special solicitude for the wretched of the Earth, and this is a defect. The utilitarian response to this demand to favor the badly off is discussed in Section 4.

The second big distributional concern starts with the thought that striving hard, maybe against the grain of your given personality proclivities, to be decent and nice to other people increases the moral value of gaining a boost in well-being for you or preventing a loss. What has been described as the unblinking

accountant's eye of the utilitarian counts a same-sized welfare benefit going to either a deserving or an undeserving person as having in itself exactly the same moral value. The opposed view is that we should give priority to the deserving. More on this second issue in Section 3.

A preliminary point to notice is that although the Dolores and Felicity example posed is stylized and simple, in broad outline, this type of decision problem often recurs and, described in some ways, is ubiquitous. People who are poor by wealth and income measures tend to be worse off than others in significant quality of life. It may be true, as they say, that the best things in life are free, but those best things tend to have prerequisites that are not free, rather costly. Also, another feature of the Felicity–Dolores example shows up often: people who are imprudent, or in some other way are poor transformers of resources into well-being gains for themselves, will tend to be worse off than others even if their access to resources is similar to the access of those who are better transformers. Of course, huge numbers of people suffer cruel oppression. But the thing about oppression is, even if you are very stalwart, it grinds you down and tends to render your quality of life Dolores-like.

Another complication in framing the Felicity–Dolores beneficence problem has already been noted two paragraphs back: you might wonder if Felicity has made herself into a person who is an efficient transformer of resources into well-being and made herself into a reasonable, prudent person and should get moral credit for those character improvements. You might incline to believe Felicity is especially deserving, and on this ground is a more apt recipient of beneficence than Dolores. You might also wonder how Dolores came to be in steady peril. Maybe she is partly to blame for the fact that she has so evidently traipsed down the wrong path. Thinking along these lines starts to look like an excuse for not helping and for not worrying much about how the background causes of social misery might be fixed.

To simplify things, let us set this issue to the side, for now. As already mentioned, we return to issues of deservingness, moral worth, blameworthiness, and personal responsibility in Section 3. To proceed, just assume that Felicity is far better off than Dolores through no merit of Felicity and through no fault of Dolores.

A final preliminary clarification, before venturing a verdict on the decision problem, whether to help Felicity or Dolores, is that to focus on the issue at hand, we should explicitly state that you do not have some prior duty to do something else with the resources you are thinking of giving to one of these two individuals. It's yours to hand out. To put it cautiously, if there are any deontic moral duties or moral rights that register at the fundamental level of moral principle, they do not have any application to this example. We put it cautiously

because not everyone accepts such deontic constraints, and perhaps, at the fundamental moral level, no one should.

With this stage-setting in place, I myself have a strong intuition to the effect that you ought to offer the bowl of gelato to Dolores, not Felicity. Here I can only appeal to the reader's intuition about the issue, once it is clearly posed. Sure, Felicity will get more out of it. The good at stake here is enjoyment, and we stipulate that Felicity would get more enjoyment from eating the ice cream than would Dolores. However, the thought is that given how badly off Dolores is in lifetime terms, it is morally more important, morally better to provide the benefit to her rather than to Felicity. It is good for either of them to get the ice cream, given either would enjoy it, and the fact that Felicity would enjoy it more is a reason to give it to her, but this reason is outweighed by a countervailing reason: a well-being increase for a person matters more, the worse off the person would otherwise be over the course of her life. A slightly misleading shorthand slogan expressing this claim is that benefits matter more; the worse off the person is who gets the benefits.

In other words, in considering who should be the recipient of your beneficence, you should give priority to the worse off. This is the basic claim of the moral doctrine that has come to be known as prioritarianism or the priority view.

1.3 Equality or Priority?

In a famous essay on "Equality," Thomas Nagel describes an example similar to our Felicity–Dolores example (Nagel 1979). Nagel imagines a parent who is contemplating a family move that would have differential effects on her two children. One child has a severe physical disability, so his life prospects are significantly worse than those of his well-functioning sibling. We can suppose that by any plausible measure of individual well-being, the able child is heading for greater lifetime well-being than the disabled child will reach. The family faces a stark choice: either move to a city, enabling the handicapped child to get special medical treatment, or move to a suburb, lacking access to ready special care for this child but clearly providing better opportunities for the flourishing of the already well-functioning child. Nagel adds this twist: the able child will gain more benefit from the move to the suburb than the disabled child would gain from the move to the city. Assume we must choose one or the other of these two options. Assume also that there are no other reasons, except for the effects on the children's welfare, that favor either move.

In both the Dolores–Felicity example and Nagel's example, the numbers matter. If the difference between the gain the able child would get from the suburb choice and the gain the disabled child would obtain from the city choice

were very slight, the reason to favor the disabled child would increase. If the difference between how well-off the able child will be apart from this choice and how well-off the disabled child will be is very small, the reason to favor the disabled child would diminish.

Nagel asserts that there would be an egalitarian reason to make the choice that favors the disabled child. This assertion might be read as another way of saying there is a prioritarian reason to favor the disabled child. That is to say, the fact that the one child is very badly off, in absolute terms, gives an extra moral reason to aid him if we can.

But a more straightforward way to interpret Nagel's claim is to recognize it makes an ineliminably comparative claim: the fact that the one child is far worse off than the other is what generates a moral reason to opt for the city move. We should at least, to some extent, be equalizing the outcomes or expectable outcomes of people.

Are these two different notions or really the same idea in different guise? In a remarkable essay, delivered as the Lindley Lecture at the University of Kansas, Derek Parfit (1995) clarified for philosophers the distinction between equality and priority. Priority says, a benefit obtained for a person is in itself morally more valuable, the worse off in absolute terms the person would otherwise be. Equality says, it is intrinsically morally valuable to bring it about that different people's condition is the same or closer to the same. So stated, the two views are views about the moral assessment of outcomes. They also become two views about what we morally ought to do if we add that, for the prioritarian, we ought to bring about outcomes in which the sum of people's benefits adjusted by each one's prioritarian weight, and for the egalitarian, we ought to bring about outcomes in which people's condition is the same or closer to the same.

Parfit explains the distinction between equality and priority in these words:

It may help to use this analogy. People at higher altitudes find it harder to breathe. Is this because they are higher up than other people? In one sense, yes. But they would find it just as hard to breathe even if there were no other people who were lower down. In the same way, on the Priority View, benefits to the worse-off matter more, but that is only because these people are at a lower *absolute* level. It is irrelevant that these people are worse off *than others*. Benefits to them would matter just as much even if there *were* no others who were better off (Parfit 1995, 23).

For now, we simply acknowledge that having an inclination to believe that one ought to give the gelato to Dolores, not Felicity, and an inclination to believe that the parent ought to move to the city, not the country does not necessarily indicate an inclination to believe the priority view is correct. The same judgments have alternative rationales. One is egalitarian.

Another possible rationale is sufficientarian (Frankfurt 1988; Crisp 2003; Benbaji 2005, 2006). Sufficiency says that what matters is that everyone has enough. In a slogan, what matters morally in itself is not that some people have more or less than others but rather that some do not have enough – they face excessively low well-being prospects. The prospects in question could be deemed to be prospects for the moment just ahead, or a stretch of future time, or the whole future course of the person's life, or the person's entire lifetime from birth to death. These same options for ranking are open for priority and equality as well.

Sufficiency, when advanced as a view about what is morally right and wrong, says that what each of us owes others is that they have enough and as a view about social justice says that the imperative of justice is to institute social arrangements that provide enough to each and every person. Applied to the Dolores and Felicity decision problem, the sufficiency advocate will say that on its natural interpretation, Felicity is above sufficiency and Dolores is below, and this is the crucial feature that determines what would be right to do in this case. One ought to help Dolores, not Felicity, and one ought to do so because she alone is below the sufficiency threshold, not to bring the distribution of well-being closer to equality, and not on the basis that Dolores is worse off. If we revise the example so that the well-being gap between the two potential beneficiaries of one's choice is the same, but both individuals are above sufficiency, this moral imperative to favor Dolores disappears.

Yet another alternative explanation of at least some moral judgments resembling priority is negatively weighted utilitarianism (NWU). This is a version of utilitarianism that gives extra moral weight to reducing bads that people undergo as compared to providing them goods. For example, if several people are experiencing pain, there is more reason to help the person with the worst pain at that time, independently of how well-off or badly off that person is in terms of the lifetime well-being she is headed toward. Negatively weighted utilitarianism implies that if we had to choose between preventing well-off Felicity from breaking her leg or instead bringing it about that badly off Dolores gets a great vacation in the Alps, we should favor, to some degree, preventing Felicity from undergoing the looming bad thing. Priority firmly disagrees. Siding with priority, one could acknowledge that we humans may have psychological dispositions to respond more strongly to bads than goods, but deny this has any normative reason-giving weight.

1.4 Axiology and Morally Right Conduct

So far the discussion has been focused on a question of right conduct: Is it morally right in some circumstances to bring about a smaller benefit for

a worse-off person than instead to bring about a larger benefit for an already better-off person? To this question, the proposed reply is Yes.

But a different question can also be raised, to which a version of priority might be proposed as the right answer. We can distinguish ways the world might go or states of affairs and outcomes, the latter being states of affairs that our acts or omissions might bring about. A for-sure outcome of an action is the difference your action would make to how the world goes – the difference in terms of what matters, what we should care about. These states of affairs can be assessed from many perspectives, in many different ways. A prudential ranking of states of affairs from a particular agent's perspective ranks the states of affairs as better or worse according to the degree to which she is well-off or badly off in one state of affairs compared to others. Let us consider impartial rankings. These assess states of affairs according to standards that do not rank one state of affairs above another on the basis that in one, some particular individuals or some particular group of individuals or members of some social group, identical in relevant respects to others, are better off, and in the alternate scenario, worse off. Let us further consider individualistic impartial rankings. These are impartial rankings that are determinable, given only information about individual persons (and other individual animals) and their characteristics. For simplicity, we set aside the important issues involving how to balance outcomes for human persons and outcomes for other animals.

Priority is an axiological (outcome evaluation) standard that ranks outcomes as morally better or worse according to the aggregate sum of priority-weighted well-being (welfare) they contain. The greater the aggregate sum of priority-weighted well-being, the better the outcome. This is in a way a broad-church, ecumenical view. Its doctrinal commitment is limited to the assertion that a certain curve is concave. This curved line is formed by measuring, along the horizontal axis, the well-being level of an individual, and along the vertical axis, the moral value of obtaining a further slight increment of well-being for an individual at any level of well-being.

In a way ecumenical: but embedded in priority are two controversial commitments. One is that the evaluation of the outcomes our actions might bring about depends only on their impact on the well-being of individuals. Well-being here names whatever in itself makes someone's life go better rather than worse for her. This is controversial in ruling out claimed impersonal goods as positively affecting the value of outcomes. Suppose the Grand Canyon will survive or not, depending on what you do, but its survival will never affect anyone's well-being (maybe all animal life on Earth is extinguished, and no extraterrestrials will ever land here and appreciate the Canyon's majestic beauty). A further implication is that the welfare of collectives and groups as such,

over and above the welfare of the individual members of the groups, makes no difference at all to the evaluation of outcomes.

The second controversial commitment embedded in prioritarian outcome assessment is that well-being is not a moralized concept. This needs explaining. The thought is that your making choices or having attitudes that are morally right or wrong, morally creditable or discreditable, does not in itself make your own well-being increase or decrease. Your well-being depends on whether your self-interest is advanced, not your interests in being moral or giving due consideration and respect to others as morality demands. Being a moral hero does not in itself add an iota of advantage to you that registers as well-being increase, and being a moral scoundrel does not in itself subtract an iota of advantage that registers as well-being decrease. Being moral or immoral is one thing and achieving high or low well-being in your life is something else entirely.

It is plain, I hope, why this commitment seems plausible and compelling to many of us. But some find this commitment not just incorrect and implausible but preposterous. Victor Tadros (2016,1) imagines a dialogue between Genghis Khan, the great Mongol conqueror of the twelfth century, and his mother, while he is still a child. Young Genghis opines that he does not want to follow a recognized career path in his society and his mother says, “Your father and I will support you and approve of whatever you choose to do in life that makes you happy.” Take “happiness” here as signifying well-being/welfare. Tadros observes that the mother’s view is silly. Becoming a world conqueror by way of wrongfully slaughtering hundreds of thousands of mostly innocent persons in itself blights the life of the conqueror, in itself makes his life go horribly bad for him. Even if Genghis Khan lives a supremely happy life, or one full of satisfied desires, or one that achieves great goods of achievement, friendship, and love, being horribly immoral in itself makes your life go badly for you. On Tadros’s view, well-being has to be understood as a moralized concept.

This may be one of those philosophical disputes in which adherents of opposed positions believe their opponents are not just wrong but obviously wrong. Against Tadros, I would say that avoiding a career path that includes the bloody slaughter of innocents with no remotely sufficient justification would definitely be choiceworthy but need not be prudent. It all depends. I would not urge my militarily precocious child to become a Genghis Khan for our times, but that is because my aims for my child extend beyond trying to promote her well-being and also include her becoming firmly disposed to interact with others in ways that are morally right, fair, and just.

A spare revision of priority, relaxing the first commitment, will be introduced in Section 3. This says, people’s being more rather than less morally deserving

enhances the impersonal moral value of welfare gains that accrue to them, just as people's being worse off enhances the impersonal moral value of such welfare gains.

Notice, the impartial standard for outcome ranking that we seek is a moral standard. We aim to find a standard that ranks outcomes of possible actions by the moral value they contain. It's morally a good thing when people fare better in welfare, and more morally valuable, according to priority, the worse off in absolute terms the beneficiaries would otherwise be.

Outcome evaluation is ranking outcomes of possible actions, and the ranking has to give rise to reasons to choose one action or another. The simplest connection between the moral ranking of outcomes and the issue of what we morally ought to choose and do is consequentialist. Consequentialism says, one morally ought always to do whatever would bring about the best reachable consequences (outcome), the one ranked highest by whatever is the correct outcome assessment standard. But there are less simple connections possible. For example, one might hold a deontological morality, according to which, one may be forbidden to do what would bring about the best reachable outcome, on the ground that doing so would violate a moral constraint we should obey, and one may not be required to bring about the best reachable outcome, because doing so would require more self-sacrifice than one is duty-bound to make.

Consider this last idea, that outcome rankings do not necessarily give rise to duties to bring about better outcomes. In this view, the outcome ranking still provides reasons for choice, just not necessarily decisive ones. Let's say the outcome ranking might at the limit just generate purely optional, take it or leave it reasons for choice, but more plausibly gives rise to a moral duty of beneficence, a duty to improve the world by bringing about impartially better rather than worse states of affairs, and the beneficence duty in turn might be deemed more or less weighty in the determination of what all things considered one morally ought to do.

To characterize the prioritarian outcome evaluation standard, it helps to contrast it with competitors (here I follow Adler 2012). First, consider utilitarianism. This standard is elegantly simple. It says, for each outcome, find the aggregate of well-being summed across persons it contains. The best outcome is the one with the highest sum, and any other outcome is worse – the greater the gap between it and the best. The associated doctrine of morally right action says: do whatever would bring about the best outcome, that is, the greatest reachable sum of well-being. Any other act is worse, the greater the gap between it and the best.

Strictly speaking, utilitarianism does not require summing across persons. All the information one needs, to determine the right course of action is what

would be the impact on total welfare of each of the acts one could now choose and execute. Whether this welfare is gained by collectives or individuals, or smeared across the two types of entity, does not signify. In contrast, to apply priority one must identify individual persons and track each person across time, to know what lifetime welfare level each is heading toward. In this way priority respects the “separateness of persons.” But in fact the classical utilitarians, certainly J. S. Mill (1979), did interpret utility as accruing to individuals (and other animals). The individual is the unique container of utility and in this way morally special.

Priority makes a small but significant adjustment to utilitarian outcome ranking (and to the associated doctrine of morally right action). Like utilitarianism, priority is *welfarist*: all that matters for the ranking of states of affairs is how the individual persons existing in those states of affairs are faring in terms of welfare or well-being. Consider outcomes in which the same persons exist. Pareto says that if in one state of affairs compared to a second, one person is better off in terms of welfare and no one is worse off, the first state of affairs is better than the second. Pareto indifference says that if no one is better off, and no one worse off, in terms of welfare, in one state of affairs compared to a second, the two states of affairs are equally good. Welfare anonymity says that two states of affairs are equally good if the welfare levels of persons in one is a permutation of the welfare levels of persons in the other. Welfarism consists of the three conditions just stated. Both priority and utilitarianism satisfy the Pareto norms and welfare anonymity.

Priority breaks away from utilitarianism by commitment to a broadly egalitarian norm, Pigou–Dalton. This says that a transfer of welfare without loss from a person with greater welfare to one with less, provided the transfer does not leave the person who gets the transfer at a higher welfare level than the other, and provided no one else’s welfare is thereby changed, makes the resulting state of affairs an improvement. This is the condition that bends the line that graphs a person’s present well-being level against the moral value of a small increase in her well-being. A continuity condition smooths this curve. This says that if one state of affairs is better than another, states of affairs close to the first are also better; the same holds if one state of affairs is worse than another. By contrast, while agreeing on continuity, the utilitarian is committed to what Krister Byqvist (2014) has called *transactional equity*. A shift in an outcome that takes well-being from one person and redistributes that same amount of well-being to others with no loss and no other well-being impact on anyone else brings about an outcome that is no worse and no better than the first. This is the case, according to transactional equity, even if the well-being is distributed from a worse-off person to a better-off person.