

M. Peterson, *The Dimensions of Consequentialism: Ethics, Equality and Risk*, Cambridge University Press, 2013, 217 pages. ISBN: 9781107033030. Hardback: £55.00/\$90.00.

Reviewed by Ralf M. Bader, Merton College, University of Oxford

Martin Peterson's book *The Dimensions of Consequentialism* attempts to develop a multi-dimensional version of consequentialism, according to which there are a number of irreducible aspects belonging to different dimensions that determine an act's rightness. These dimensions are meant to include well-being, equality and risk. The irreducibility of the different aspects is taken to imply that conflicts between them are irresolvable. Rightness and wrongness accordingly turn out not to be binary but instead to be a matter of degree, such that an act can be somewhat right and somewhat wrong due to being optimal with respect to some aspect but sub-optimal with respect to another aspect. Peterson considers this theory to have a number of advantages when it comes to accounting for various intuitions that have been taken to conflict with traditional versions of consequentialism, as well as when it comes to making sense of moral dilemmas.

INDIVIDUATING ASPECTS

In setting out the basic notions of the theory, Peterson informs us that "something counts as a moral aspect if and only if it directly influences an act's deontic status, irrespective of how other aspects are altered. That something directly influences the deontic status of an act should be understood as a claim about functional relationships: an aspect, a , directly influences the deontic status, d , of an act if and only if d is a function of a " (p. 3). A multi-dimensional theory then is one according to which the deontic status of an act can only be understood as a function of multiple aspects, whereas a one-dimensional theory holds that the deontic status can be understood in terms of a one-place function.

This way of construing aspects and characterising the distinction between one-dimensional and multi-dimensional theories is problematic for a number of reasons:

1. Understanding what it is for something to directly influence something else in terms of the latter being a function of the former is not particularly plausible. This is because functional relationships can be chained since functions can be composed, whereas (any intuitive understanding of) direct influence does not allow for chaining. If x is a function of y , and y is a function of z , then x will be a function of z , i.e. if $x = f(y)$ and if $y = g(z)$, then $x = (g \circ f)(z)$. This, however, does not mean that z directly influences x . Instead, the influence may very well be only indirect insofar as it is mediated by y .

2. In order to apply this distinction, one needs to be given a way of individuating aspects. For any set of aspects, it is always possible to define a single complex aspect that has all of them as its components and that yields the same outputs as them.¹ Unless one is told how one is to rule out such gimmicky aspects, which requires one to privilege certain individuations, every theory will trivially be classified as a one-dimensional theory. Yet, identifying a privileged individuation is not an easy task. Is average utilitarianism a multi-dimensional (or rather multi-aspect)² theory, given that it considers the deontic status of an act to be a function of the well-being of all the individuals in question, making it a many-place function that takes utility vectors as its arguments, or is it a one-dimensional theory given that it can be characterised as a one-place function taking average utility quantities as its arguments? If average quantities are admissible as argument places, then why is not the case that other disjunctively characterised aspects are equally admissible?
3. Classifying something as a moral aspect if the deontic status of an act is a function thereof irrespective of how other aspects are altered is incompatible with any remotely plausible multi-dimensional theory. If there are multiple aspects, then this requires that d be a function of each aspect a , irrespective of how the other aspects are altered. This, however, is only possible if each aspect yields the same output as the other aspects for all inputs. Accordingly, this condition requires that any multi-dimensional theory only allows there to be a single deontic status that attaches to all acts, no matter which features of the different aspects they have. That is, the deontic status d cannot be a function of a as well as a function of d' , where these aspects are such that they can be varied independently, unless both functions are constant functions having the very same output. Otherwise, d would take on contradictory values, which would be problematic in its own right, and would also contradict the claim that d was a function of a since a would then only partially account for the act's deontic status,

¹“If it were possible to aggregate two or more conflicting aspects into a new – possibly very complex – aspect, then those aspects could obviously be reduced into a single aspect, which would entail that the theory in question was actually one-dimensional. By definition, multi-dimensional consequentialists believe that an act's deontic status depends on two or more *irreducible* aspects” (p. 3). Peterson does not provide a satisfactory explanation as to how this notion of irreducibility is to be understood (he briefly touches on the issue on p. 4 footnote 6). It seems that it would be most plausible for it to be some form of normative irreducibility. However, in that case one cannot characterise the contrast between one-dimensional and multi-dimensional theories in terms of that between one-place and many-place functions, given that functions are not sensitive to considerations of normative (ir)reducibility.

²Peterson allows that different aspects can belong to the same dimension. Cf. “Strictly speaking, multi-dimensional consequentialism need thus not be a multi-dimensional theory. In extreme cases, all aspects that determine an act's deontic status could belong to the same dimension” (p. 5).

which would make it the case that d would instead be a function of the different aspects taken together.³

Whilst the official definition is a non-starter, a more promising characterisation is given later on: “something counts as an aspect if and only if the deontic status of an act varies if we hold constant everything but the putative aspect in question” (p. 15).⁴ This implies that for something to be an aspect, it must be possible to vary it whilst holding everything else constant. This, however, is incompatible with Peterson’s commitments, threatening to render his version of multi-dimensional consequentialism incoherent. In particular, Peterson wants to hold that both equality and individual well-being are aspects. The problem now is that it is not possible to hold well-being facts fixed whilst varying equality, given that equality is itself a function of well-being and as such not separable therefrom.

IRRATIONALITY

When it comes to decision-making Peterson argues that we should randomise, using a ‘weighted-force rule’ that considers both the strength and degree to which an act is right (cf. p. 119). This choice procedure, however, would seem to be subject to money-pump arguments and hence classify as irrational. By randomising one can end up in a sub-optimal situation.⁵ Peterson is aware of this objection, but dismisses it on the basis that it would show too much. His argument is that, in the same way that the randomisation strategy ensures that an infinite number of choices will lead with certainty to a situation in which the agent has lost everything, standard expected utility theory has the consequence that an agent will suffer a total loss with complete certainty given an infinite number of gambles, and would thus be equally ruled out (cf. p. 138).

The infinite gambles case, however, differs in important ways from the ran-

³It might be suggested that, rather than a moral aspect being something that determines the deontic status of an act, it is something that determines an aspect thereof. This characterisation has the advantage of not being incoherent. However, it does not help us all that much since we are now understanding a moral aspect in terms of an aspect of a deontic status.

⁴It is not clear whether this condition is to be understood such that the function has to be the same no matter the level at which the other aspects are held constant, in which case this would imply that aspects have to be separable, requiring that the manner in which the deontic status of an act is influenced by one aspect has to be independent of how the other aspects influence its deontic status. If so, the condition would be rather substantive, ruling out many plausible theories by stipulation. If not, it would amount to the rather minimal condition that aspects not have inessential features.

⁵Peterson employs money-pump argument against various theories that allow for incomplete orderings. Given that advocates of incomplete orderings hold that choices between non-comparable options (or options that are on a par) are in an important sense arbitrary, the only difference between Peterson’s account and the theories that he rejects consists in the fact that his account includes a weighting function. The weighting function, however, does not affect whether one is susceptible to money-pumping (at best, it affects the degree to which one can be money-pumped).

domisation case. In order to establish the irrationality of randomisation, it is not necessary to appeal to what happens when applying this decision rule an infinite number of times. It is enough to show that there is a chance of a loss without any corresponding or compensating chance of a gain. By randomising, the agent is certain to be made not better off but possibly worse off – there is a chance that the agent will lose and a chance that the agent will end up in a non-comparable situation, but there is no chance of the agent gaining. As the number of choices increases, the chance of loosing increases. Yet, in order to show that randomisation is problematic, it suffices to show that there is some chance (no matter how small) of losing. Simply staying with the status quo whenever facing non-comparable alternatives will dominate randomisation.⁶ It is this that makes randomising irrational and that distinguishes it from acting in accordance with standard expected utility theory.⁷

COLLAPSE

Multi-dimensional consequentialism is at risk of collapsing into traditional consequentialism. In particular, one might worry that saying that rightness and wrongness are not binary features but come in degrees is not of any great significance, insofar as this simply implies that instead of performing an act that is right, we now have to perform an act that is right to the greatest degree. Even though rightness might not be a binary notion, we can define a binary notion rightness*, whereby it is this latter notion that determines what is to be done:

$$\text{right}^*(\phi) = \begin{cases} 1 & \text{iff } \text{right}(\phi) \geq \text{right}(\psi) \text{ for all } \psi \\ 0 & \text{iff } \text{right}(\psi) > \text{right}(\phi) \text{ for some } \psi \end{cases}$$

Peterson objects to the idea of maximising rightness. “If you were to maximise rightness, this would not properly reflect the fact the other options were also right to a high degree” (p. 116). Instead, we are told that “the rational thing to do is to give the right-making features of each act their due” (p. 117). In response, we can note that what Peterson calls the non-weighted maximising approach (which is developed in the Appendix) fails to satisfy this desideratum. It specifies that $\mathcal{O}(\phi) = \max(\{O_i(\phi) : i \in I\})$ (cf. pp. 200-201). This means that if $O_a(\phi) = x$ and $O_i(\phi) = 0$ for all $i \neq a$, whereas $O_i(\psi) = x - \epsilon$ for all i , then $\mathcal{O}(\phi) = x$, whereas $\mathcal{O}(\psi) = x - \epsilon$. This is difficult to reconcile with the idea that all right-making features are given their due, insofar as the substantial differences between

⁶Dominance is to be construed in terms of the outcome not being worse in every state of nature and strictly better in some, rather than in terms of it being at least as good in every state and strictly better in some.

⁷Randomisation has further bizarre consequences. If one can either produce 100 units of utility or produce 90 units, then the former act would be entirely right whereas the latter would be almost right (cf. pp. 47-48 for an analogous case). Utilising the weighted-force rule and randomising amongst these options, however, seems completely irrational.

ϕ and ψ with respect to aspects $i \neq a$ do not have any impact on the status of these actions.⁸

In addition, the idea of defining a binary notion in terms of the non-binary notion does not rely on maximisation. One can specify some other function, such as the weighted-force rule. Peterson addresses this objection: “If agents should randomise among a set of pure acts, it could be claimed that it would be right to the highest degree to perform the corresponding randomised act. The best response to this objection is to note that the weighted force rule is a claim about rationality, not about morality” (p. 120). However, it is not clear how exactly the contrast between rationality and morality is to be construed. Moreover, it looks like the contrast between single- and multi-dimensional theories will no longer be particularly substantive. Advocates of the different theories will agree as to what should be done – they will simply differ in terms of whether this ‘should’ is one of morality or of rationality.

The theory is also at risk of collapsing if the supposed irreducibility of the different aspects should turn out to be only apparent. If different aspects are genuinely irreducible, then there is no way of trading them off against each other, no way of weighing them up. This means that one needs to render them somehow commensurable if degrees of rightness are to be compared across different aspects (something that is required if the weighted-force rule is to be applicable). In the Appendix, comparability is introduced by means of a 0-1 normalisation. “By letting each function be normalised to a scale from 0 to 1, the degree to which each moral aspect is fulfilled can then be directly compared across the set of all moral aspects” (p. 200). From normalisation it is only a small step to collapse. Peterson is aware of this problem. “It is true that *something* is aggregated . . . , but that entity is not the kind of entities (consequences) aggregated by one-dimensional consequentialists. . . . The aggregation-mechanisms . . . refer to a higher level of abstraction: the ultimate deontic level” (pp. 203-204).

This response, however, is not satisfactory. If we have commensurability of the different oughts, we should also have commensurability of the different aspects or values on which these oughts are based. If the deontic level can be normalised and aggregated, then it should likewise be possible to normalise and aggregate at the axiological level. If aspects are irreducible and not amenable to aggregation, then the corresponding oughts should likewise be irreducible and not amenable to aggregation. This structural correspondence is particularly plausible in the context of a consequentialist theory, given that the essence of consequentialism consist in a reduction of the deontic to the axiological, which means that the two are meant to go hand-in-hand.

⁸Understanding $\mathcal{O}(\phi)$ in terms of an average or aggregate of the different aspects (as happens in the case of what Peterson calls the weighted maximising approach) will run into difficulties once one allows that different acts can be determined with respect to different aspects.

OVERALL ASSESSMENT

Whilst the book contains some interesting discussions, it does not seem that multi-dimensional consequentialism constitutes a coherent alternative to traditional forms of consequentialism but instead risks either ending up being incoherent or collapsing into single-dimensional consequentialism.