

4. Weighing the Costs and Benefits of Public Policy: On the Dangers of Single Metric Accounting

Johanna Thoma

This chapter presents two related challenges to the idea that, to ensure policy evaluation is comprehensive, all costs and benefits should be aggregated into a single, equity-weighted wellbeing metric. The first is to point out how, even allowing for equity-weighting, the use of a single metric limits the extent to which we can take distributional concerns into account. The second challenge starts from the observation that in this and many other ways, aggregating diverse effects into a single metric of evaluation necessarily involves settling many moral questions that reasonable people disagree about. This raises serious questions as to what role such a method of policy evaluation can and should play in informing policy-making in liberal democracies. Ultimately, to ensure comprehensiveness of policy evaluation in a wider sense, namely, that all the diverse effects that reasonable people might think matter are kept score of, we need multiple metrics as inputs to public deliberation.

On Comprehensive Policy Evaluation and Single Metric Accounting

When policy interventions have far-reaching consequences for years to come it is important to make sure to take proper

account of all relevant effects in evaluating one's options. This is the case for decisions made at times of crisis, like those in response to the COVID-19 pandemic, as well as for decisions made in ordinary times. As Paul Dolan puts it, there is the initial splash of the policy, but there are also the ripple effects. We must pay attention to both [1, Chapter 3 of this book]. He laments that, in the UK at least, there has been too much of a focus on one type of initial effect of pandemic responses, measured using a single metric, namely, number of lives lost in the short-term. Too narrow a focus on this single measure has made us lose sight of costly ripple effects on long-term health (including mental health) and the economic prospects of young generations in particular.

It is hard to disagree with a call for public policy evaluation to be comprehensive, in particular when the stakes are high. Yet clearly, this is also a point that often needs (re-)emphasising. The pandemic appears to be such a moment in time. Beyond a call for comprehensiveness, however, Dolan also defends a particular form that policy evaluation should take, drawing on his long record of ground-breaking research on happiness and public policy. Interestingly, a single metric also features as part of his proposed strategy for making sure everything is taken into account: all costs and benefits, short- and long-term, should be aggregated using a single metric. And then policies should be evaluated in terms of the value and distribution of that metric. More specifically, the metric Dolan proposes to be used is subjective wellbeing adjusted life years, or WELLBYs [1 p4].¹ That is, what we should be aggregating is how many life years are gained and lost as a consequence of the policy intervention, where

life years are weighted by a measure of the self-reported quality of people's experiences. This WELLBY metric is then used in an equity-weighted cost-benefit or cost-effectiveness analysis, allowing for gains and losses in WELLBYs to count for more when they occur to the worst off.²

WELLBYs will remind the reader of the more well-known QALY (quality-adjusted life year) metric, widely used in policy evaluation in public health (although, interestingly, as Dolan points out, little-used so far during the pandemic). But QALYs are restricted to the evaluation of health-related outcomes, whereas WELLBYs are intended to be, in line with the guiding thought of the proposal at hand, more comprehensive. And while the measurement of QALYs typically relies on hypothetical comparative judgements people make about different health states, the measurement of WELLBYs is supposed to be based, as much as possible, on direct reports of how different outcomes are actually experienced when lived through. Here Dolan puts his cards on the table [1 p5]: What WELLBYs aim to capture is a hedonistic notion of wellbeing; a good life is one of good (pleasurable, purposeful) experiences.

Relative to a hedonistic standard of what a good life is, people's own hypothetical comparative judgements often come out biased. Using a hedonistic metric as a measure of costs and benefits in a standard cost-benefit analysis is a way of implementing a hedonistic form of utilitarianism, according to which the sum-total of subjective wellbeing should be maximised. Equity-weighting one's cost-benefit analysis, on the other hand, implements a hedonistic form of prioritarianism, according to which the experiences of the worst-off count for more.

The philosophical debate on hedonism as a theory of well-being, and on utilitarianism and prioritarianism is vast,³ and, it is safe to say, will never reach a consensus on the central issues. Instead of offering a critical discussion of these particular moral views and their application in public policy evaluation, I would like to take a step back and consider the more general idea that a single metric should be used in order to aggregate the many effects of public policies and to articulate the trade-offs involved in making a decision. I will offer two related challenges to this idea.

The first is to point out how the use of a single metric limits the extent to which we can take 'distributional concerns' into account, making the moral commitments embodied in such approaches narrower than it might initially seem. The possibility of equity-weighting is often presented as a way to flexibly accommodate those with more egalitarian leanings, giving the appearance of a 'broad church'. But equity-weighting can only accommodate concerns about the distribution of the single metric that is being equity-weighted. And those with egalitarian leanings often also care about the distribution of other things (or indeed equality of a non-distributional kind).

The second challenge starts from the observation that in this, and in many other ways, the choice of a single metric and its implementation in the aggregation of diverse effects necessarily involves settling many moral questions that reasonable people (including even, sometimes, philosophers) disagree on. This raises serious questions as to what role such a method of policy evaluation can and should play in informing policy-making in liberal democracies. Ultimately, to ensure comprehensiveness of

policy evaluation in a wider sense, namely, that all the diverse effects that reasonable people might think matter are kept score of, we need multiple metrics as an input into public deliberation.

What Else Might Matter

The hedonistic theory of individual wellbeing is only one of several major kinds of theories of what is good for individuals. Some have argued that what is good for you is to get the things you desire, or would desire under some ideal conditions, whether you desire pleasurable experiences or not. The pandemic and the policy responses to it around the world have massively frustrated people's life plans. According to desire-fulfilment theories, this would matter a great deal whether or not going through with those plans would have made people happy. Others have argued that there are some things that are good for people whether they enjoy or desire them or not. If education, engagement with the arts, regular gatherings with family, or direct contact with different ways of life are such things, again the pandemic and policy responses have undermined these in ways that a hedonistic wellbeing measure might not capture.

Using a metric for policy evaluation based on any one of the available theories of wellbeing means using a metric that risks leaving out or not doing full justice to some of the things at least some reasonable people take to be morally relevant. Moreover, many people believe that there are things other than wellbeing that matter, such as freedom (as is again very relevant during the pandemic), or the preservation of the environment for its own sake. In this respect, too, equity-weighted cost-benefit analysis using a broadly hedonistic wellbeing measure cannot

capture everything deemed to be morally relevant by at least some parts of the population. Perhaps less obviously, I want to focus here on how even those who agree that the hedonistic theory is the correct theory of individual wellbeing, and that all else that matters are ‘distributional concerns’ may not be satisfied that equity-weighted subjective wellbeing-based cost-benefit analysis captures everything that is morally relevant. This is because the framework can accommodate some, but not all distributional concerns.

Just as there is lively debate on what the correct theory of individual wellbeing is, there is debate on the metric(s) of equality or priority: Equality *of what* or priority in access *to what* matters? Having a single metric for the purposes of policy evaluation, and accommodating distributional concerns through equity-weighting commits us to using that same single metric as our metric of equality or priority. In the proposal under discussion, what matters is an equitable distribution of subjective wellbeing, or giving greater weight to the subjective wellbeing of those who are worse-off in terms of subjective wellbeing. As Dolan puts the idea of equity-weighting, the claims individuals have to resources depend both on the gain in WELLBYs they can expect as a result of those resources, and on their current and expected lifetime suffering or wellbeing compared to others (as well as the WELLBY effects they have on others) [1].

But those who defend some form of distributional equality or priority often care about the distributions of things other than wellbeing (be it hedonic or not). For instance, resources, capabilities, or opportunities for wellbeing are alternative potential metrics of equality or priority.⁴ Again, it is easy to think of ways

in which distributions of these alternative metrics may have been affected in ways not perfectly correlated with wellbeing itself during the ongoing pandemic, which would not be captured by the proposed method of policy evaluation under discussion. For instance, the pandemic and policy response have affected – and have most likely diminished – the ways in which resources can be translated into wellbeing, but despite that, you might think it still matters that some simply have more than others. And the pandemic response has taken away many opportunities for welfare people may not actually have made use of, but which you might think were nevertheless important for them to have as much of as those who did use them.

So, equity-weighted WELLBY-based cost-benefit analysis is narrower in the ways in which it can accommodate distributional concerns than it might initially seem. And in large part, this is down to the ambition of using a single metric to both capture expected harms and benefits, and to capture distributional concerns. Whatever your single measure of costs and benefits, an equity-weighted cost-benefit analysis can only take into account the distribution of that single metric, and will miss other distributional concerns that have been defended by prioritarrians and egalitarians. A prominent branch of egalitarianism, social egalitarianism, is not even concerned directly with distribution at all, but rather with making sure that people can engage with each other as equals.⁵

Finally, even granting a single metric for the purposes of capturing harms and benefits and for distributional concerns, there remains an important ambiguity about what kind of distributional concern equity-weighting involves in the context of risk;

that is, when policies impose probabilities of harms and benefits on people, rather than certainties. In such contexts, is what matters the *ex ante* distribution of risks of harm and chances of benefit, or is it the *ex post* distribution of harms and benefits? Given how central risk is to the evaluation of policy responses to the COVID-19 pandemic, the next section will explain this ambiguity in some more detail, and argue that talk of a single metric of policy evaluation risks obfuscating the issue.

Distributing Harms, Distributing Risks

Consider the following choice problem loosely based on an example by Peter Diamond [16]. You are in charge of making sure one of your equally well-off flatmates – Amal or Bella – moves out, and this outcome would be equally bad for each. Do you:

- A. Choose Amal,
- B. Choose Bella, or
- C. Throw a fair coin, giving each a 50% chance of staying?

There seems to be an intuitive equity case for C, even though *ex post*, the outcomes of all three choices have the same wellbeing distribution. If you agree, this is likely because you think the distribution of chances of harms and benefits matters. To show how such an intuition might extend to a stylised policy case, suppose that as a policy-maker, you have to choose between the following two prospects for a population of ten million:

- D. Everybody faces an additional 0.002% risk each of a loss of 30 WELLBYs.

E. One thousand people who are currently at welfare levels that are representative of the population at large face an additional 10% chance each of a loss of 30 WELLBYs.

Here, there seems to be an intuitive equity case for D, even though it is virtually certain that the loss of WELLBYs is larger in D and no more equally distributed ex post. If you agree, then again this seems to express concern for the distribution of risks of harm, which are much more concentrated on a few individuals in E.

Cost-benefit analysis in the social welfare function tradition can implement equity-weighting in two main different ways in the context of risk: It can either introduce equity weights on the ex ante expectations of harms and benefits a proposed policy imposes on individuals; or it can equity-weight the ex post distributions of harms and benefits in the population for each potential policy outcome, and recommend the policy option with the best expectation of equity-weighted outcomes.⁶ The recommendations of the two approaches can come apart, as they (likely) would in the two examples just described. The first strategy is sensitive to the distribution of risks, would recommend option C in the first case, and, with the right parameter choices, could recommend option D in the second case. The second strategy is insensitive to ex ante distributions of risks, and would be indifferent between the options in the first case, and recommend E in the second case.

There is a lively debate about which of these two strategies is better. There are also ways to combine them.⁷ The point I want to raise here is not only that defence of an equity-weighted cost-benefit

analysis is ambiguous on this morally important question. It is also that talk of a single metric may obfuscate the issue. Both the *ex ante* and the *ex post* approaches (and any combination between them) use a single wellbeing metric, and then merely proceed to combine this with probabilities and equity-weights in different ways. But in so doing, the *ex ante* approach implements a distributive concern for a currency different from wellbeing, namely chances of wellbeing, or conversely, risks of harm. If we think that the distribution of such chances and risks matters, so if we favour the *ex ante* or a mixed approach, then there are really two things we need to keep track of: The expected wellbeing distribution in the population as a consequence of policies, and how the risks of harms and chances of benefits are distributed in the population. If we don't, there is something else that our approach to policy evaluation does not keep track of that at least some people find morally relevant.

This issue is especially relevant in the pandemic given how central risk is in the management of the pandemic response. The virus itself poses a risk of death or serious adverse health outcomes that is much higher for some parts of the population rather than others. And different potential policy responses differ in the extent to which they concentrate or spread risks of harmful outcomes, such as unemployment, within the population. If avoiding an unequal spread of risks of harm is acknowledged as a distinct policy goal, this may in some cases lead us to accept lower and no more equitably distributed expected aggregate wellbeing in the population *ex post*, as in the stylised policy example above.

The Problem of Reasonable Disagreement

The last two sections canvassed a number of things that matter morally, according to at least some of the people who have thought about them, but that the framework of policy evaluation proposed by Dolan will not account for – and that indeed any framework insisting on a single metric to capture both costs and benefits, as well as distributional concerns will fail to account for. There are many more ways in which implementing an equity-weighted WELLBY-based cost-benefit analysis involves settling on specific answers to a number of contentious moral questions: What is the right theory of wellbeing? If it is hedonistic in general, which experiences count as bad, which as good? Is wellbeing all that matters? Whose wellbeing matters? What should the equity weights be? Should they be applied *ex ante* or *ex post* (or both)? And so on. These are all questions that reasonable, thoughtful and well-informed people disagree about. So, policy evaluation using this framework can only hope to be comprehensive in the sense that all effects of a policy are accounted for against the background of specific answers to these questions. It can't be comprehensive in the sense that it accounts for everything that at least some reasonable people take to be morally relevant.

This problem is not specific to Dolan's proposal, of course. It arises especially starkly for proposals that aspire to aggregate all potential effects of policies into a single metric, and output specific policy recommendations (rather than, say, simply present an array of potentially relevant considerations to policy-makers). Doing so must involve making judgements about what matters, how it matters, and how trade-offs between the things that matter

are to be made. Of course, public decision-makers cannot get around making such judgements eventually. They have to choose, after all, and they must do so against a backdrop of reasonable disagreement in the population they aim to serve. The question is what role comprehensive frameworks of policy evaluation, presented and advocated for by social scientists – such as the one under discussion – can and should play in this eventual decision-making process.

The danger, as I see it, is illustrated by this caricature: If social scientists were to simply present policy recommendations based on evaluations in terms of a single metric capturing many different effects to policy-makers and to the public without further context and qualification, this would not only mask all of the contentious moral decisions that went into the construction of that metric, but it would also endow the recommendation with the authority of scientific expertise, making it hard for public decision-makers to diverge from the recommendation.⁸ And that would be a threat to the liberal democratic ideal of how public decision-making in the face of reasonable disagreement should be done: Value conflicts should be resolved by democratically elected officials in a way that is open to public scrutiny.⁹ Of course, policy-making needs social scientific input, and recent philosophy of science is also rich in demonstrations that social science, just like any science, can't help but be value-laden.¹⁰ But there are clearly ways in which social scientists can make sure to help, rather than undermine, democratic decision-making. When it comes to comprehensive frameworks for policy evaluation like the one advocated by Dolan, there seem to be two main strategies for doing so.

One strategy is to work closely with the public and democratically elected officials to devolve as much as possible all important value judgements, so that the resulting recommendations would have democratic legitimacy.¹¹ There are some suggestions in Dolan's text pointing towards this kind of strategy. For instance, he writes '[e]ven if we never end using a single metric as the final arbiter on what to do, the processes by which we discuss the data required to generate one, and debates about how to make the diverse array of human experiences commensurable with one another, will lead to policy decisions that better account for the myriad of ripple effects they generate' [1 p3]. Moreover, he suggests drawing on, and generating more evidence on, what the public thinks about various kinds of trade-off to inform, for instance, how equity-weights are set, while at the same time conceding that '[e]mpirical investigation of these issues can only get us so far. We need to ensure that the policymaking processes better reflect the myriad of concerns and impacts' [1 p11].

To assuage worries about a lack of democratic legitimacy of methods of evaluation that take sides on morally contentious questions, however, there would need to be democratic input on all the contentious value assumptions. And as the foregoing aimed to illustrate, these go far beyond the ways in which equity-weights are set, but also concern what the appropriate currency of distributive justice is to begin with, and much more. Ensuring the democratic legitimacy of every element of the analysis would be a large undertaking.¹² And it is, moreover, not clear to me that such an undertaking would result in anything like an equity-weighted WELLBY-based cost-benefit analysis. For instance, there is some evidence that many people

are reluctant to trade-off especially large burdens against any number of smaller burdens, which is antithetical to this framework.¹³

The alternative strategy involves presenting one's preferred framework of policy evaluation (or its specific applications) to policy-makers as only one of several reasonable ways of evaluating policy options. Rather than a way of settling what policy-makers should do, the analysis would be an input into public decision-making. To serve as a good basis for public discussion and eventual policy choice, there has to be transparency about all the value judgements that went into the assessment and that are inherent in the general framework.¹⁴ But if we care about comprehensiveness in the wider sense – that there is proper accounting of all the things that reasonable people might find morally relevant – we also have to make sure that alternative frameworks and metrics are presented to the public, to enable there to be an informed public debate amongst people with different values, to reveal whether there are options that can be endorsed from any or most moral perspectives, and ultimately to give policy-makers informed options as to which values to pursue.

Of course, ensuring such wider comprehensiveness is not the responsibility of any one scientist or research team or even subfield, but rather of the scientific community at large and the science policy that sets its parameters. It is from this wider perspective that I think the call for a single metric of policy evaluation is problematic. From within some particular value frameworks (for instance, a hedonistic *ex post* prioritarian one), the call for a single metric of policy evaluation makes sense (assuming there is transparency about what goes into the metric) and

can be a means of comprehensively aggregating everything morally relevant within that moral framework.¹⁵ But from the wider political perspective, where the goal should be to ensure that the outputs of policy-relevant social science enable and inform public discourse in the context of reasonable disagreement, what we need are multiple metrics and frameworks.

Conclusion

I have presented two related challenges to the idea that, to ensure policy evaluation is comprehensive, all costs and benefits should be aggregated into a single, equity-weighted metric. Firstly, the only distributional concerns such equity weighting can accommodate concern the distribution of that single metric. But those with prioritarian or egalitarian leanings often care about the distributions of things other than our chosen metric of costs and benefits (subjective wellbeing in Dolan's case). Moreover, this is just one of many ways in which aggregating diverse effects into a single metric of policy evaluation involves settling on specific answers to controversial moral questions that reasonable people disagree on. The second challenge is that this raises serious questions as to what role such a method of policy evaluation, and advocacy for it by social scientists, can and should play in informing policy-making in liberal democracies. There is a wider sense of comprehensiveness of policy analysis, where the ideal is that everything that reasonable people might think is morally relevant is kept proper score of as an input into public deliberation and choice. Given reasonable people disagree on many important questions of value, achieving such wider comprehensiveness requires the use of multiple metrics.

Acknowledgements

I thank the participants of the Beveridge 2.0 Symposium on Wellbeing at the LSE for helpful discussion, my colleagues Alex Voorhoeve and Liam Kofi Bright for valuable feedback on an earlier draft, and Paul Dolan for the opportunity to comment on his work.

Notes

¹ See [2] for a fuller defence of the use of WELLBYs as a currency for the evaluation of public health interventions, and [3] for an overview of the subjective wellbeing approach to cost-benefit analysis and its potential advantages and disadvantages more generally.

² See [4] on equity-weighting using the more standard QALY metric.

³ [5] and [6] are useful overviews of the philosophical debate on theories of wellbeing, including hedonism. [7] provides a good overview of defences and critiques of utilitarianism. See [8] for a recent, and [9] for the classic defence of the prioritarian view, and [10] for an application of prioritarianism to the public health response to the pandemic. For a defence of a utilitarian pandemic response, see [11], and for an exploration of the difference that the use of utilitarian or prioritarian social welfare functions in cost-benefit analysis is likely to make in evaluating pandemic responses, see [12].

⁴ See Section 3 in [13].

⁵ See [14] for a classic statement of this view. [15] defends a pluralist view according to which both social equality and fair distribution matter.

⁶ On these approaches, see [17] and [18] who defends an ex post prioritarian approach.

⁷ [17] reviews some of the formal ways to combine them. [19] defends a hybrid egalitarian view that cares about both ex ante and ex post equality.

⁸ Or, if the recommendations diverge from those of other groups of scientists, e.g. medical ones, call into question the authority of scientific expertise.

⁹ Also see [20] on this worry in the context of the pandemic.

¹⁰ See, for instance, [21] and [22]. [23] raises the problem of value-ladenness specifically in the context of studying the effects of the pandemic.

¹¹ A similar proposal is made by Anna Alexandrova in the context of establishing ‘mixed claims’ (claims that are partly descriptive and partly evaluative) in wellbeing science, to establish a kind of objectivity for these claims [24]. Scientists should check value presuppositions for controversy, and then consult the relevant parties on their views on these controversies.

¹² Alexandrova is already concerned about the cost of such exercises in establishing mixed claims about wellbeing only, which would only be one input in a comprehensive policy evaluation [24]. Also note that by departing from classic cost-benefit analysis in abandoning preference in favour of subjective wellbeing as a measure of individual welfare, Dolan loses what some have taken to be the democratic appeal of cost-benefit analysis – that it is a neutral way of aggregating people’s preferences over policy options. See, e.g., [25].

¹³ As discussed in [26], which also defends the reasonableness of such reluctance.

¹⁴ This is a common minimal demand made by most philosophers of science writing on the value-ladenness of science, including, for instance, Alexandrova and, in the context of the pandemic, Stephen John [24, 23].

¹⁵ Other value frameworks, for instance ones that are pluralist about the objects of egalitarian concern, will themselves require more than a single metric of account, even before we consider the need to accommodate reasonable disagreement. See [15] for such a pluralist egalitarian critique of an exclusive focus on equity-weighted cost-effectiveness analysis in health care priority-setting.

References

1. Dolan P. Accounting for consequences and claims in policy. *LSE Public Policy Review*. 2021; 2(2): 1–8. DOI: <https://doi.org/10.31389/lseppr.43>. A version is republished as Chapter 3 of this book.
2. De Neve J, Clark A, Krekel C, Layard R, O'Donnell G. Taking a wellbeing years approach to policy choice. *BMJ*. 2020; 371(m3853). DOI: <https://doi.org/10.1136/bmj.m3853>
3. Dolan P, Fujiwara D. Happiness-based policy analysis. In: Adler M, Fleurbaey M (eds.), *The Oxford Handbook of Well-Being and Public Policy*. Oxford: Oxford University Press. 2016; 286–317. DOI: <https://doi.org/10.1093/oxfordhb/9780199325818.013.9>
4. Bleichrodt H, Diecidue E, Quiggin J. Equity weights in the allocation of health care: The rank-dependent QALY model. *Journal of Health Economics*. 2004; 23(1): 157–171. DOI: <https://doi.org/10.1016/j.jhealeco.2003.08.002>
5. Fletcher G, (ed.). *The Routledge Handbook of Philosophy of Well-being*. Abingdon: Routledge; 2016. DOI: <https://doi.org/10.4324/9781315682266>
6. Crisp R; Well-being. In: Zalta E, (ed.), *The Stanford Encyclopedia of Philosophy*. Metaphysics Research Lab, Stanford University. 2017: fall 2017 edition.
7. Sinnott-Armstrong W. Consequentialism. In: Zalta E, (ed.), *The Stanford Encyclopedia of Philosophy*. Metaphysics Research Lab, Stanford University. 2019: summer 2019 edition.
8. Adler M, Holtug N. Prioritarianism: A response to critics. *Philosophy, Politics and Economics*. 2019; 18(2): 101–144. DOI: <https://doi.org/10.1177/1470594X19828022>
9. Parfit D. *Equality or priority*. The Lindley Lecture, The University of Kansas. 1995.

10. Nielsen L. Pandemic prioritarianism. *Journal of Medical Ethics*; 2021. DOI: <https://doi.org/10.1136/medethics-2020-106910>
11. Savulescu J, Persson I, Wilkinson D. Utilitarianism and the pandemic. *Bioethics*. 2020; 34(6): 620–632. DOI: <https://doi.org/10.1111/bioe.12771>
12. Ferrana M, Sevilla J, Bloom D. Addressing the COVID-19 pandemic: Comparing alternative value frameworks. *National Bureau of Economic Research*. 2021; w28601. DOI: <https://doi.org/10.3386/w28601>
13. Arneson R. Egalitarianism. In: Zalta E, (ed.), *The Stanford Encyclopedia of Philosophy. Metaphysics Research Lab, Stanford University*. 2013; summer 2013 edition.
14. Anderson E. What is the point of equality? *Ethics*. 1999; 109: 287–337. DOI: <https://doi.org/10.1086/233897>
15. Voorhoeve A. Why health-related inequalities matter and which ones do. In: Norheim O, Emanuel E, Millum J, (eds.), *Global Health Priority-Setting: Beyond Cost-Effectiveness*. Oxford: Oxford University Press. 2020; 145–162. DOI: <https://doi.org/10.1093/oso/9780190912765.003.0009>
16. Diamond P. Cardinal welfare, individualistic ethics, and interpersonal comparisons of utility: Comment. *Journal of Political Economy*. 1967; 75: 765–766. DOI: <https://doi.org/10.1086/259353>
17. Mongin P, Pivato M. Social evaluation under risk and uncertainty. In: Adler M, Fleurbaey M, editors. *The Oxford Handbook of Well-Being and Public Policy*. Oxford: Oxford University Press. 2016; 711–745. DOI: <https://doi.org/10.1093/oxfordhb/9780199325818.013.23>
18. Adler M. *Well-being and Fair Distribution*. Oxford: Oxford University Press; 2011.
19. Otsuka M, Voorhoeve A. Equality versus priority. In: Olsaretti S (ed.), *The Oxford Handbook of Distributive Justice*. Oxford: Oxford University Press. 2018; 65–85. DOI: <https://doi.org/10.1093/oxfordhb/9780199645121.013.1>

20. Stefansson O. *Three mistakes in the moral reasoning about the COVID-19 pandemic*. Institute for Future Studies Working Paper Series. 2020; 12.
21. Douglas H. *Science, Policy, and the Value-Free Ideal*. Pittsburgh: University of Pittsburgh Press; 2009. DOI: <https://doi.org/10.2307/j.ctt6wrc78>
22. Kincaid H, Dupré J, Wylie A. (eds.) *Value-Free Science?: Ideals and Illusions*. Oxford: Oxford University Press; 2007. DOI: <https://doi.org/10.1093/acprof:oso/9780195308969.001.0001>
23. John S. The ethics of lockdown: Communication, consequences, and the separateness of persons. *Kennedy Institute of Ethics Journal*. 2020; 30(3): 265–289. DOI: <https://doi.org/10.1353/ken.2020.0015>
24. Alexandrova A. *A Philosophy for the Science of Well-Being*. Oxford: Oxford University Press; 2017. DOI: <https://doi.org/10.1093/oso/9780199300518.001.0001>
25. Heath J. Cost-Benefit Analysis as an Expression of Liberal Neutrality. In: *The Machinery of Government*. Oxford: Oxford University Press; 2020. DOI: <https://doi.org/10.1093/oso/9780197509616.003.0005>
26. Voorhoeve A. Balancing small against large burdens. *Behavioural Public Policy*. 2018; 2(1): 125–142. DOI: <https://doi.org/10.1017/bpp.2017.4>