RESEARCH PAPER
Penultimate version. Final version forthcoming in: *Prometheus: Critical Studies in Innovation.* Special issue: Designing for Value Change (2022).

Future value change: identifying realistic possibilities and risks

Jeroen Hopster[1]

1 Faculty of Behavioural, Management and Social Sciences (BMS), Philosophy (WIJSB), University of Twente, Twente, The Netherlands

ABSTRACT

The co-shaping of technology and values is a topic of increasing interest among philosophers of technology. Part of this interest pertains to anticipating future value change, or what Danaher (2021) calls the investigation of "axiological futurism". However, this investigation faces a challenge: "axiological possibility space" is vast, and we currently lack a clear account of how this space should be demarcated. It stands to reason that speculations about how values might change over time should exclude farfetched possibilities and be restricted to possibilities that can be dubbed to be realistic instead. But what does this realism criterion entail? This paper introduces the notion of *realistic possibilities* as a key conceptual advancement to the study of axiological futurism and offers suggestions as to how realistic possibilities of future value change might be identified. Additionally, I propose two slight modifications to the approach of axiological futurism. First, I argue that axiological futurism can benefit from a thoroughly historicised understanding of moral change. Secondly, I argue that when employed in normative contexts, the axiological futurist should seek to identify realistic possibilities that come along with substantial normative risks.

Introduction

While philosophers have long sought to unearth morality's unshakable foundations, in recent years, there has been a growing scholarly interest in how morality evolves over time. Examples of recent and current moral shifts in Western societies include the moralisation of CO_2-emissions, the demoralisation of homosexuality, and the elevated moral status of several non-human species. Taking a broader look at history, moral change abounds and manifests itself at various timescales. As a result, researchers from multiple fields have taken an interest in the topic, scrutinising, for instance, the evolutionary roots of moral cooperation (Boehm 2012; Sterelny 2021; Tomasello 2016), the sociobiological and material pressures that have shaped moral systems throughout human history (Buchanan & Powell 2017; Kitcher 2011; Morris 2015), as well as the moral

CONTACT: j.k.g.hopster@utwente.nl

revolutions the have occurred over the last few centuries (Appiah 2010; Baker 2019; Pleasants 2018).

Philosophers and ethicists of technology, too, have broached the topic of moral change, specifically by studying the dynamic relation between values and technology (Kudina 2019). While some work on technomoral change (Swierstra 2013) is historically oriented (Nickel et al. 2021; Hopster et al. MS), recent lines of investigation also include conceptual, normative, and anticipatory analyses of value change. For instance, van de Poel (2021) has articulated a typology of value change, while the technomoral scenario approach advanced by Boenink et al. (2010) is geared explicitly towards anticipating future trajectories of technomoral change. Such anticipatory work, in turn, is often connected with normative aims. By comparison, policymakers regularly invoke the normative planning method of *backcasting*: start with outlining a desirable future and subsequently work backwards to disclose pathways to realise it. Something similar might be done in assessing emerging technologies (Brey 2012), or for purposes of value sensitive design (Friedman and Hendry 2019): identify technomoral futures which differ in terms of their moral desirability and subsequently adjust technology design to make the actualisation of the more desirable futures more probable.

Anticipatory models and scenarios do not aim to predict *what will occur* in the future. Attempts at predicting the future of complex societal dynamics have a record of being hubristic (Jasanoff 2003; van de Poel 2016). Instead, the current generation of approaches in foresight ethics (Floridi & Strait 2020) seeks to anticipate *what might occur* in the future. Given the limitations of our knowledge, as well as the indeterminacy of how causal pathways towards the future unfold, this implies that analysts should anticipate a range of possible moral futures (cf. Boenink et al. 2010). On the one hand, this moderates futurists' predictive ambitions, but it makes the undertaking more rigorous in turn. Rather than succumbing to the tunnel-vision of envisioning only one future, the aim of anticipatory studies should be to generate a sketch of the broader landscape of future possibilities.

A proviso needs to be added, however: *mere possibilities* that do not stand a serious chance of becoming actual provide are a distraction to futurists and should be excluded from this broader vista. In other words, some sort of plausibility requirement should play a role in scenarios of future value change. As of yet, little reflection has been given regarding the precise nature of this requirement. Arguably, the most rigorous methodological proposal to date for studying the future dynamics of moral change has been outlined by Danaher (2021). Under the header of "axiological futurism", Danaher proposes systematically exploring "axiological possibility space" and outlines tools for navigating it. Danaher is aware that the vastness of axiological possibility space poses a challenge for this undertaking. He concurs that speculations about how values might change over time should exclude farfetched possibilities and be restricted to possibilities that can be dubbed *realistic*, in some relevant sense. But how, exactly, should this realism criterion be understood? Neither Danaher's work, nor existing research on technomoral change (e.g. Swierstra 2013), provides a worked-out view to answer this question. Possibilities come in different sizes and shapes. Philosophers commonly distinguish between logical, metaphysical, and physical possibilities, and various further modal categories can be distinguished: epistemic possibilities, conceptual possibilities, and so on. Which of these, if any, are the kinds of possibilities that anticipatory scenarios should focus on?

This article proposes an answer to this question, and thereby seeks to further the project of axiological futurism. Its answer, to wit, will be to take the plausibility requirement at face value: axiological futurists should seek to approximate the historically indexed notion of "real

possibilities".[1] In the spirit of constructive criticism, I will argue for this claim while engaging with Danaher's (2021) framework, which provides the most rigorous methodological outline for studying future value change to date. But I take the lessons to be equally applicable to the technomoral change literature and to anticipatory approaches in the ethics of technology more broadly: these should be restricted to identifying realistic possibilities, and a historically oriented approach can provide important insights as to what such possibilities may amount to.

The second proposal of this article pertains to the normative aims of axiological futurism. Axiological futurism may be regarded as a value-neutral tool, whose aims are purely anticipatory. But as noted, the tool can be incorporated in frameworks of anticipatory ethics and value sensitive design, whose aims are explicitly normative. I will argue that when employing the tool with normative aims, axiological futurists can benefit from adhering to a further criterion. Not only should they focus on identifying realistic possibilities; additionally, they should seek to identify realistic possibilities that are distinctly risky. I make this argument by drawing an analogy with anticipatory efforts in climate discourse. Recent debates in climate science, reflected in the latest report of the Intergovernmental Panel on Climate Change (IPCC 2021), suggest that for purposes of decision-making in the face of deep uncertainty, a concerted effort to pin down the respective likelihoods of different future scenarios may not be very helpful. Instead, the bulk of decision-relevant information pertains to the question of which impacts that are regarded as particularly harmful, constitute realistic risks. Since similar – if not greater – uncertainty besets anticipations of future value change, axiological futurists are well advised to adopt a similar focus on identifying risky outcomes that are particularly harmful, when employing axiological futurism as a normative tool.

Axiological Futurism and the "Mere Possibilities" Challenge

Let me begin by recapitulating how Danaher conceives of the project of anticipating future value change, an inquiry he calls "axiological futurism". A key aspect of the methodology he proposes to conduct this inquiry centres on the notion of an "axiological possibility space". An axiological possibility space represents a possible constellation of future values. The representation need not be comprehensive: it might single out a few future possibilities rather than sketching the range of possibilities in its entirety. For present purposes, however, let us consider the entire set of possible future value trajectories. Thus understood, studying future value change involves the systematic exploration of axiological possibility space.

What possibilities are encapsulated in this space? For one, this can be gauged by studying the extant diversity of moral theories and codes adopted in current societies. As Danaher (pp. 4–5) observes, moral frameworks typically specify the following four components:

(i) *An axiology*: a theory specifying what counts as good and what counts as bad ("values" and "disvalues").

(ii) *Agents and patients*: a specification of what makes someone or something a moral subject and/or an object of moral concern.

(iii) *Internal relations*: a specification of the relationships between the different elements within an axiology, such as their respective priority.

[1] Real possibilities have previously been discussed in the philosophical literature, albeit in a rather different context, namely the metaphysical debate over indeterminism and free will (Müller et al. 2018). Since this metaphysical context is tangential to the current discussion, I set it aside in this paper.

(iv) *Pattern of outcome*: a specification of the appropriate pattern according to which moral goods ought to be procured (e.g. maximisation of goods; sufficiency for each agent; *etc*.).

These parameters do not only vary amongst extant moral frameworks, but also provide an indication of how morality might change over time: the identification of moral values and disvalues may shift, the circle of moral concern may expand or contract, the priority of values and subjects may change, and the favoured pattern of outcome may alter (cf. van de Poel 2021).

To give an example of the former – changes at the axiological level – consider the value of equality. As Danaher observes:

> "Philosophers have identified dimensions or parameters along which different conceptions of [this value] can vary. A theory of equality, for example might vary along two dimensions: equality of opportunity and equality of outcome. Given these two parameters, a researcher can construct a simple 2 × 2 logical space for the value of equality, classifying different possible axiologies depending on whether they score high or low on those two dimensions." (p. 6)

Put in modal terms, philosophers have previously identified different conceptual possibilities associated with the value of equality. In turn, these can be rendered as parameters in axiological possibility space. Which possibilities, however, are admissible? Here axiological futurism faces a challenge, as Danaher acknowledges:

> "Presumably, axiological possibility space is vast — much larger than anyone can really imagine. But equally, many of the 'possible' axiologies within this space are not that plausible or interesting: e.g. a world in which the subjective pleasure we experience while scratching our knees is the only recognised good may be *possible* (in some thin sense of the word 'possible') but is not very plausible and should not concern us greatly. [Therefore] we need some constraints on the boundaries of axiological possibility space to make the project feasible." (p. 5)

Danaher rightly stresses that axiological futurists should not rest content with outlining *mere possibilities*. Instead, they should take on the more challenging task of outlining possibilities that are realistic, in some relevant sense. But how to judge whether this is the case?

Danaher provides two suggestions to this effect. First, a wide range of evidential sources provide insight into what is – and is not – possible in terms of moral change, much of which goes beyond strictly philosophical work. Values differ in current societies, they have varied throughout the course of history, and they are associated with distinct psychological traits. These variations and associations shed light on the diversity of possible value systems, as well as some of the constraints thereupon. To give one example, Danaher refers to moral foundations theory, which suggests that there are five or six robust dimensions of value in human moral psychology (Graham et al., 2013). If this is correct, and these five or six dimensions are recurring pillars for any moral system, then this sets constraints on the kinds of variations that moral systems can take. Similar kinds of insight about the breadth and constraints of moral possibility space might arise from fields of inquiry such as evolutionary anthropology (e.g. Henrich 2020), comparative human ecology (e.g. Flanagan 2017), as well as the study of human history (e.g. Morris 2015).

Secondly, constraints can be discerned from existing work in axiological theory. This work suggests, for instance, that there is a fixed set of items that can be included in any possible list of

goods/bads, such as "subjective pleasure, desire satisfaction, knowledge, friendship, beauty, education, health, money, family [and] food" (Danaher, p. 5). The same holds for the kinds of entities that are ascribed a certain moral status (e.g. all humans, all sentient beings, all living entities), or the relationships that exist between them (e.g. equal treatment vs hierarchical treatment). While there certainly exists a substantial amount of moral diversity, we already seem to possess a good idea of what this diversity amounts to.

I submit that the general outline of Danaher's twofold response to the "mere possibilities" challenge is along the right lines, but its details need refining. Danaher is certainly right to draw lessons from empirical sciences and history in ascertaining what is – and is not – possible in terms of future value change. In fact, I maintain that these lessons might be taken even further and should outweigh the conceptual considerations that Danaher additionally appeals to. Consider his appeal to existing work in axiology, which suggests that the sets of goods that can be contained in any value system are fairly static. This appeal seems difficult to reconcile with the apparently open-ended character of the value changes that have occurred throughout history. The value of sustainability, for instance, does not seem to be reducible to any of the items in the abovementioned set of goods. Historically, this value was only clearly articulated – and has become widely endorsed – since the last decades of the 20th century (van de Poel 2021). Or consider the value of privacy, which – like many other values – has been interpreted differently in different historical epochs (Holvast 2009), transforming from a physical notion to an informational notion. Given that the conceptualisation of values evolves over time, an inventory of the conceptual diversity currently recognised in philosophical work is likely to provide a skewed and overly narrow picture.

There is reason to think, then, that axiological possibility space should be rendered more dynamic than Danaher's account allows for. My proposal, to this effect, is that the axiological futurist's framework be more thoroughly historicised. This proposal goes hand in hand with a suggestion to further explicate the conceptual aim of axiological futurism. What, exactly, are the kinds of possibilities that axiological futurists should seek to identify? I answer this question in the next section, arguing that axiological possibility space should be understood in terms of the temporally indexed notion of realistic possibilities.

Realistic Possibilities: A Historicist Approach

Real possibilities are those possibilities that might actualise, conditional on the state of the world at a specific moment in time. Hence, at some point in time t, the set of real possibilities consists in all states-of-affairs whose realisation is compatible with the state-of-the-world at time t (Betz 2016). Real possibilities differ from other kinds of possibility that regularly feature in philosophical discussions, such as logical, conceptual and physical possibilities, in virtue of being a temporal notion. They are anchored in concrete situations and oriented towards the future: what is really possible at any given moment is what can temporally evolve from a concrete situation against the background of what the world is like (Müller et al., 2018). As time passes, real possibilities can become more proximate or remote: what was once a far-fetched possibility may turn into an adjacent possibility over time, or vice versa. The remoteness of a possibility depends on its ease of realisation: the more difficult this realisation is, the less realistic it becomes (Hopster 2018).

While real possibilities are a metaphysical notion, they have their epistemic counterpart in what I call *realistic possibilities*: possibilities that we judge to be realistic to the best of our knowledge. When it comes to sketching scenarios about what might realistically occur in the future, our aim is to approximate the real possibilities as close as we can. The totality of our background

knowledge is relevant for this purpose: at minimum, our identification of realistic possibilities should be compatible with this background knowledge (Betz 2010). But arguably, this minimal constraint does not suffice: many propositions that are not obviously excluded by our background knowledge seem highly improbable, nonetheless. To reiterate Danaher's (p. 5) example, in a thin sense of being "possible", there might be a possible future world in which the subjective pleasure we experience while scratching our knees is the only recognised good. Arguably, our background knowledge does not strictly *exclude* this possibility. However, neither is there any positive support for this proposition. In the absence of such support, it should be disqualified as a realistic possibility.

What matters, then, for a possible state-of-affairs to be included in the set of realistic possibilities is not only that our background knowledge does not exclude it, but also that we can give it some positive epistemic support. Such support might consist, for instance, of outlining the mechanism which could give rise to this state-of-affairs: if this mechanism is well understood and deemed scientifically plausible, then the state-of-affairs is realistically possible. Historical precedents are another source of insight into realistic possibilities: if some state-of-affairs obtained in the past, then *ceteris paribus* it might also obtain in the future. Considerations like these are quite relevant, for instance, in the context of assessing long-term risks of climate change (Parker and Risbey 2015). Our historical and geological record of mass extinctions, rapid temperature shifts and massive sea-level rise provide an important evidential source to take the prospect of radical long-term change due to anthropogenic greenhouse-gas emissions very seriously (Hopster 2020a).

Now, let us return to the project of axiological futurism. My conceptual claim is that the possibilities that axiological futurists should be after are *realistic possibilities*, understood in the above sense. Danaher's (2021) approach to axiological futurism is largely consonant with this proposal: it is thoroughly multidisciplinary, takes input from evidence from various sources and pays specific attention to historical examples to calibrate axiological possibility space. But as noted, Danaher's reliance on conceptual claims about the goods contained in all axiologies might be unduly restrictive and insufficiently anchored in an historical account of what state-of-affairs might evolve starting from our current position. His approach, I submit, can be historicised more fully. Armed with the notion of realistic possibilities, let me suggest four avenues to re-adjust the approach of axiological futurism along these lines.

A first suggestion is to give explicit attention to processes moralisation and demoralisation (Buchanan and Powell 2017) in the anticipatory framework and to seriously entertain the notion, drawing on historical evidence, that such processes are surprisingly open-ended. Moral values can be reinterpreted over time, as the previously given example of privacy suggests. But human practices can also be moralised anew, leading to new values such as sustainability. Furthermore, practices may lose their moral significance. For examples of the latter, consider the value of chastity, which has lost much of its moral significance in the Western world since the sexual revolution of the 1960s and onwards (Hopster et al. MS). Or consider bastardy, which was heavily moralised until the early 20th century in England (Baker 2019), but has arguably become entirely devoid of moral significance in English society today. A further dimension of moralisation is the objectification of moral issues (Hopster and Klenk 2020; Wright 2021), which can similarly shift as time passes. Rather than setting out with a fixed list of goods and bads, axiological futurists might be better served by incorporating the sliding scales of moralisation and objectification as key dimensions in axiological possibility pace.

How can we anticipate, however, which objects and entities may be (de)moralised in the future? A second suggestion for the axiological futurist is to take seriously the idea that morality

evolves through societal pressures (Anderson 2016; Hopster 2020b; Smyth 2020) and that moral norms and values can often be understood as a functional response to the problems of communal life that societies have historically faced (Kitcher 2011). If this view is along the right lines, then we should expect that the major challenges societies will face in the future will similarly solicit a moral response. What will be the main challenge that societies are likely to face over the next few decades? And what pressures might these challenges exert on moral norms and values? To view moral evolution as a process of historical, societal learning (Hopster 2020a), which typically occurs reactively in the face of the challenges that sociotechnical and environmental predicaments engender, provides a helpful angle to get some hold of the realistic possibilities of the future.

Building on this approach, a third suggestion to the axiological futurist is to come up with a more detailed framework of what is involved in stabilising and disrupting values (Hopster 2021b). For instance, to what extent do socioeconomic and technological background conditions correlate with the adoption of specific value regimes? To what extent do such conditions enable, or even determine, the respective moral frameworks? As Danaher acknowledges, historical work (esp. Morris 2015) can provide a rich source for answering these questions, and so can various insights from human biology, sociology and anthropology (e.g. Flanagan 2017). But the merits of these sources notwithstanding, a fully-fledged account of how values are stabilised in the interplay with broader societal and technological dynamics is still forthcoming. In this regard, I propose that it might be specifically helpful to scrutinise the promise of *(techno)moral niche construction* (Severini 2016; Hopster et al. MS) as a conceptual framework to describe and anticipate processes of moral change.

My fourth suggestion to the axiological futurist is to study indicators of value change. Consider moral disagreement: is this a reliable indicator of impending moral change? Or does this depend on the nature of the disagreement in question (e.g. Hansson 2018)? Conversely, is the historical inertness of a value a reliable indicator of its future stability? Or does this depend, for instance, on whether the inertness has resulted from a process of diverse and critical interrogation (Hopster 2017; cf. Longino 1990)? As it stands, these questions are under-theorised. But as noted, the project of investigating moral change, as outlined in the introduction of this article, has only recently taken off in earnest. A more rigorous theoretical framework of how morality can change, and of what is predictive of moral changes, will give a better hold on assessing the realistic possibilities of value change that lie ahead.

Lessons from Climate Scholarship: Scenarios and Risk

There are some notable parallels between anticipatory projects in the philosophy of technology and in climate scholarship. Global warming is a slow-moving and long-term process; some of its impacts are likely to be felt over the course of decades, centuries and beyond (Gardiner 2011). As a result, much of climate science – as well as climate policy and climate ethics – is decidedly future-oriented. The instruments that climate scientists have developed to make projections about climate futures are very sophisticated and have emerged through decades of concerted scientific effort (Winsberg 2018). This holds, in particular, for the computer simulations used in the Climate Model Intercomparison Project (CMIP), which constitute the basis for the climate projections outlined by the IPCC.

The advanced state of the art in climate modelling makes this field an interesting example for anticipatory endeavours in the philosophy of technology, which are by and large still in an explorative stage. A further commonality between anticipatory efforts in these domains is that both

are couched in substantial uncertainty. Their sophistication notwithstanding, projections derived from climate models come along with various uncertainties (Hopster 2022), for instance due to potential measurement errors, the natural variability of the climate system, the unknown external forcings on the climate system, the idealisations of simulation models, as well as the potential aggrandisement of model biases through their merger in ensembles like CMIP (e.g. Baumberger et al., 2017).

Building on CMIP models, key pieces of information provided to decision-makers are the projected pathways of the Earth's mean surface temperature during the 21st century, conditional on different anthropogenic greenhouse gas emission scenarios, as well as the likelihood that certain tipping points in the climate system will be crossed. While these projections and likelihoods are carefully argued for and meticulously justified, this is done in a framework primarily oriented towards scientific understanding rather than a framework oriented towards decision-making about climate risks (cf. Sutton 2019). By way of illustration, the statement that because of climate change, the Atlantic gulf stream might be overturned during the 21st century has been dubbed "highly unlikely" by the IPCC, based on modelling efforts (IPCC 2014). Judged by the criteria outlined in the previous section, however, it should probably be included in the set of realistic possibilities. Furthermore, it is a possibility with a serious risk attached to it, as the overturning of the gulf stream is likely to have many corollary impacts that will engender major harm.

For these reasons, in recent years several climate scholars have grown critical about the dominance of CMIP models, combined with the IPCC's predictive statements couched in terms of likelihood, in informing climate policy-makers (e.g. Shepherd et al. 2018). Arguably, given the uncertainty that is inherent in modelling efforts, CMIP projections should not be the main focal point to anchor policy-decisions in the face of the imminent dangers of global warming. Instead, as a policy instrument, these projections should be complemented with scenario approaches, which are better suited for communicating realistic possibilities in the face of deep uncertainty. A prominent scenario approach that is currently being developed along these lines is the storyline approach (Sillmann et al. 2021). A storyline is defined as "a physically self-consistent unfolding of past events, or of plausible future events or pathways" (Shepherd et al. 2018, p. 555). Storylines do not have any specific probability attached to them. Instead, typical storylines single out compound risks that arise because of climate change (Zseischler et al. 2020). In doing so, they focus on identifying realistic possibilities that matter from a normative point of view. This approach fits well with core principles of disaster risk management (King et al. 2015). In its latest report, the IPCC (2021) similarly endorses the principles of a disaster risk management approach, and has increased its emphasis on communicating findings regarding low-likelihood, high-impact events.

There are two lessons of this recent debate in climate scholarship, I submit, that should be taken to heart by axiological futurists. The first lesson speaks to Danaher's suggestion to extent axiological futurism with the help of computer-assisted models (Danaher 2021, *passim*). While such efforts are likely to be valuable, they should not be pursued in isolation. Model-based approaches should be developed alongside more qualitatively oriented scenarios, as each of these pursuits comes with its own advantages (Challinor 2018 et al.). When it comes to anticipating value change, scenario approaches may be even more important than in climate science. This is because the dynamics of value change are arguably more historically contingent, and therefore more difficult to predict, than the dynamics of climate change. Under conditions of greater uncertainty, the usefulness of scenario-approaches, relative to modelling-approaches, increases. Concretely, one might take from this that current theorising on technomoral change (e.g. Boenink et al. 2010), which has specifically adopted the scenario-approach as its preferred method, should not simply be

discarded to give way to Danaher's more formal framework. Instead, axiological possibility spaces may be regarded as an instrument to contribute to the rigour of technomoral scenarios.

The second lesson we can take from climate scholarship concerns the normative aims of anticipatory studies. As discussed in the previous section, axiological possibility space should be tailored to identifying realistic possibilities. An important aspect of the criticism of current approaches to climate modelling is that modellers seek to outline the full range of scientifically credible possibilities. In doing so, they insufficiently focus on highlighting what – from a risk perspective – are the most salient possibilities among these: the outliers and extremes, the tipping points, and the catastrophes that may ensue. Hence, from a normative point of view, climate modellers should arguably have a distinct focus on higlighting possibilities of substantial harm – provided, of course, that these possibilities satisfy the epistemic standard of being realistic (Hopster 2021a).

Axiological futurists should not fall into the same trap. When operationalized in the context of normative frameworks, such as anticipatory ethics or value sensitive design, the aim of anticipatory endeavours should not be to outline axiological possibility space in a fully comprehensive manner. There may simply be too many realistic possibilities of value change to entertain, not all of which are equally significant from a normative point of view. Furthermore, given the deep uncertainty at issue, efforts to predict which future value changes are particularly likely to occur may be futile. Instead, a more promising route is to first identify which future value changes, from a normative point of view, would be particularly significant – for instance because they directly compromise the normative aims of a given technological design, or because they would render the design harmful. Hence, the first step is to pinpoint the most salient normative risks posed by future value change. Subsequently, building on the epistemic toolkit of the axiological futurist, engineers and ethicists should ascertain whether these prospects of "risky future value change" can be regarded as realistic. If they pass the realism threshold, then it is crucial that engineers make adaptable designs, such that the potentially adverse consequences of risky value change can be mitigated.

Conclusion

I have argued for a further crystallisation of the conceptual aims of axiological futurism and proposed avenues to advance it, both as an anticipatory framework and as a normative tool. First, anticipatory accounts of value change can benefit from being firmly anchored in our historical understanding of moral change. Accordingly, the notion of an "axiological possibility space" can be usefully spelt out in terms of realistic possibilities. Realistic possibilities are historically conditioned. One important strategy to identify realistic possibilities of future value change is by considering which processes of moralisation and demoralisation might occur, in the wake of future challenges that societies are likely to face. Secondly, I have argued that we should take seriously the lessons from recent climate change scholarship, which serve to underline that where morally relevant – and potentially harmful – yet uncertain changes are at play, purely anticipatory scenarios can be toothless. Such scenarios should be combined with a disaster risk mitigation approach, aimed at avoiding particularly bad outcomes.

While these considerations are theoretical, they can find a more practical application, for instance, in processes of value sensitive design. What engineers can take from the former historical lesson, is to think seriously about processes moralisation and demoralisation, and the historical pressures thereto, as being crucial to the dynamics of value change. To make technological designs

that are resilient to value change, then, a key task for engineers is to assess whether the values embedded in their designs are likely to shift in terms of their moral significance and which features of their designs might plausibly be moralised in the future. What engineers can take from the latter, risk-based approach, is that assessing their designs to identify which evaluative components constitute the most likely candidates to be subject to future value change may not be the best way to proceed. Instead, efforts should be made to identify realistic value changes that make a design specifically vulnerable from a moral point of view, and to mitigate the potentially harmful implications that might ensue.

References

Anderson E. 2016. The Social Epistemology of Morality: Learning from the Forgotten History of the Abolition of Slavery. In Brady M, and Fricker M, eds. *The Epistemic Life of Groups: Essays in the Epistemology of Collectives*: 75–94. Oxford: Oxford University Press.

Appiah KA. 2010. *The Honor Code: How Moral Revolutions Happen*. New York: W.W. Norton & Company.

Baker R. 2019. *The Structure of Moral Revolution. Studies of Changes in the Morality of Abortion, Death, and the Bioethics Revolution*. Cambridge MA: MIT Press.

Baumberger C, Knutti R, Hirsch Hadorn G. 2017. Building Confidence in Climate Model Projections: An Analysis of Inferences from Fit. *WIREs Clim. Change 8*, e454.

Betz G. 2010. What's the Worst Case? The Methodology of Possibilistic Prediction. *Analyse & Kritik 1*: 87–106.

Betz G. 2016. Accounting for Possibilities in Decision Making. In: SO Hansson and GH Hadorn eds. *The Argumentative Turn in Policy Analysis: Reasoning under Uncertainty.* Springer: 135–170.

Boehm C. 2012. *Moral Origins: The Evolution of Virtue, Altruism, and Shame*. Basic Books.

Boenink, M, Swierstra T, Sternerding D. 2010. Anticipating the Interaction between Technology and Morality: A Scenario Study of Experimenting with Humans in Bionanotechnology. *Studies in Ethics, Law, and Technology 4*(2): 1–38.

Brey P. 2012. Anticipatory Ethics for Emerging Technologies. *Nanoethics 6*(1): 1–13.

Buchanan A, Powell R. 2017. *The Evolution of Moral Progress: A Biocultural Theory*. Oxford: Oxford University Press.

Challinor AJ, Adger WN, Benton TG, Conway D, Joshi M, Frame D. 2018 Transmission of Climate Risks Across Sectors and Borders. *Philosophical Transactions Royal Society A 376*: 1–23.

Danaher J. 2021. Axiological Futurism: The Systematic Study of the Future of Human Values. *Futures 132*: 102780.

Flanagan, O. 2017. *The Geography of Morals: Varieties of Moral Possibility*. Oxford: OUP.

Floridi L, Strait A. 2020. Ethical Foresight Analysis: What It Is and Why It Is Needed? *Minds & Machines 30*: 77–97.

Friedman B, Hendry DG. 2019. *Value Sensitive Design: Shaping Technology with Moral Imagination*. Cambridge (MA): MIT Press.

Gardiner SM. 2011. *A Perfect Moral Storm*. New York: Oxford University Press.

Graham J, Haidt J, Koleva S, Motyl M, Iyer R, Wojcik SP, et al. 2013. Moral Foundations Theory: The Pragmatic Validity of Moral Pluralism. *Advances in Experimental Social Psychology, 47*: 55–130.

Hansson, SO. 2018. Risk, Science and Policy. A Treacherous Triangle. *Ethical Perspectives 25*(3): 391–419.

Henrich J. 2020. *The WEIRDest People in the World: How the West became Psychologically Peculiar and Particularly Prosperous*. Penguin Ltd.

Hopster JKG. 2017. Two Accounts of Moral Objectivity: From Attitude-Independence to Standpoint Invariance. *Ethical Theory and Moral Practice 20*(4): 763–780.

Hopster JKG. 2018. *De andere afslag. Hoe had het leven anders kunnen lopen?* Amsterdam: Amsterdam University Press.

Hopster, JKG. 2020a. Shall we Adapt? Evolutionary Ethics and Climate Change. In: Hermann J, Hopster JKG, Kalf WF, Klenk MBOT eds. *Philosophy in the Age of Science? Inquiries into Philosophical Progress, Method, and Societal Relevance.* Rowman & Littlefield: 195–213.

Hopster JKG. 2020b. Explaining Historical Moral Convergence: The Empirical Case against Realist Intuitionism. *Philosophical Studies 177*: 1255–1273.

Hopster, JKG. 2021a. Climate Uncertainty, Real Possibilities, and the Precautionary Principle. *Erkenntnis 00*: 1–17.

Hopster, JKG. 2021b. What are Socially Disruptive Technologies? *Technology in Society 67*: 101750.

Hopster, JKG. 2022. Climate Change, Uncertainty and Policy. In: Pellegrino G, Di Paolo M eds. *Handbook of the Philosophy of Climate Change.* Springer.

Hopster JKG, Klenk MBOT. 2020. Why Metaethics needs Empirical Moral Psychology. *Crítica: Revista Hispanoamericana de Filosofía*. *52*(155): 27–54.

Hopster, JKG, Arora C, Blunden C, Eriksen C, Frank LE, Hermann JS, Klenk MBOT, O'Neill ERH, Steinert S. Pistols, Pills, Pork and Ploughs: The Structure of Technomoral Revolutions. Manuscript.

Holvast, J 2009. History of Privacy. In: Maty V, et al.: *The Future of Identity*. IFIP AICT 298: 13–42.

IPCC. 2021. Summary for Policymakers. In: Climate Change 2021: The Physical Science Basis. Contribution of Working Group I to the Sixth Assessment Report of the Intergovernmental Panel on Climate Change [Masson-Delmotte V, Zhai P, Pirani A, Connors SL, Péan C, Berger S, Caud N, Chen Y, Goldfarb L, Gomis MI, Huang M, Leitzell K, Lonnoy E, Matthews JBR, Maycock TK, Waterfield T, Yelekçi O, Yu R, Zhou B, eds.]. Cambridge: Cambridge University Press.

IPCC. 2014. Climate Change 2014: Mitigation of Climate Change. Contribution of Working Group III to the Fifth Assessment Report of the Intergovernmental Panel on Climate Change [Edenhofer O, Pichs-Madruga R, Sokona Y, Farahani E, Kadner S, Seyboth K, Adler A, Baum I, Brunner S, Eickemeier P, Kriemann B, Savolainen J, Schlömer S, von Stechow C, Zwickel T, Minx JC, eds.]. Cambridge: Cambridge University Press.

Jasanoff S. 2003. Technologies of Humility: Citizen Participation in Governing Science. *Minerva 41*(3): 223–244.

King D, Schrag D, Dadi Z, Qui Y, Ghosh A. 2015. *Climate Change: A Risk Assessment*. Cambridge University Centre for Science and Policy Rep.

Kitcher P. 2011. *The Ethical Project*. Cambridge MA: Harvard University Press.

Kudina O. 2019. *The Technological Mediation of Morality: Value Dynamism, and the Complex Interaction between Ethics and Technology*. Dissertation: University of Twente.

Longino HE. 1990. *Science as Social Knowledge: Values and Objectivity in Scientific Inquiry*. Princeton University Press.

Morris, I. 2015. *Foragers, Farmers, and Fossil Fuels: How Human Values Evolve.* Oxford: Princeton University Press.

Müller T, Rumberg A, Wagner V. 2018. An Introduction to Real Possibilities, Indeterminism, and Free Will: Three Contingencies of the Debate. *Synthese 196*(1): 1–10.

Nickel PJ, Kudina O, van de Poel I. 2021 (forthcoming). Moral Uncertainty in Technomoral Change: Bridging the Explanatory Gap. *Perspectives on Science.*

Parker WS, Risbey JS. 2015. False Precision, Surprise, and Improved Uncertainty Assessment. *Philosophical Transactions R. Soc. A 373:* 20140453.

Pleasants NJ. 2018. The Structure of Moral Revolutions. *Social Theory and Practice 44*(4): 567–592.

Poel, van de I. 2016. An Ethical Framework for Evaluating Experimental Technology. *Science and Engineering Ethics 22*(3): 667-86.

Poel, van de I. 2021. Design for Value Change*. Ethics and Information Technology 23*: 27–31.

Severini, E. 2016. Evolutionary Debunking Arguments and the Moral Niche. *Philosophia 44*: 865–875.

Shepherd TG, Boyd E, Calel RA et al. 2018. Storylines: An Alternative Approach to Representing Uncertainty in Physical Aspects of Climate Change. *Climatic Change 151*: 555–571.

Sillmann J., Shepherd TG, van den Hurk B, Hazeleger W, Martius O, Slingo J, Zscheischler J. 2021. Event-based Storylines to address Climate Risk. *Earth's Future 9*: e2020EF001783.

Smyth N. 2020. A Genealogy of Emancipatory Values. *Inquiry 00*: 1–31.

Sterelny K. 2021. *The Pleistocene Social Contract: Culture and Cooperation in Human Evolution*. Oxford: Oxford University Press.

Sutton RT. 2019. Climate Science Needs to take Risk Assessment Much More Seriously. *Bulletin of the American Meterological Society 100*(9): 1637–1642.

Swierstra T. 2013. Nanotechnology and Technomoral Change. *Ethics and Politics 15*(1): 200–219.

Tomasello M. 2016. *A Natural History of Human Morality*. Cambridge (MA): Harvard University Press.

Winsberg EB. 2018. *Philosophy and Climate Science*. Cambridge: Cambridge University Press.

Wright, JC. 2021. Morality as a Regulator of Divergence: Protecting against Deviance while Promoting Diversity. *Social Cognition 39*(1): pp. 81–98.

Zscheischler J, Martius O, Westra S et al. 2020. A Typology of Compound Weather and Climate Events. *Nat Rev Earth Environ 1:* 333–347.