

Should Philosophers Begin to Employ New Methods If They Want to Become More Societally Relevant?

Van de Poel, Ibo

Publication date

2020

Document Version

Final published version

Published in

Philosophy in the Age Of Science?

Citation (APA)

Van de Poel, I. (2020). Should Philosophers Begin to Employ New Methods If They Want to Become More Societally Relevant? In J. Hermann, J. Hopster, W. Kalf, & M. Klenk (Eds.), *Philosophy in the Age Of Science?: Inquiries into Philosophical Progress, Method, and Societal Relevance* (pp. 231-244). Rowman & Littlefield.

Important note

To cite this publication, please use the final published version (if applicable). Please check the document version above.

Copyright

Other than for strictly personal use, it is not permitted to download, forward or distribute the text or part of it, without the consent of the author(s) and/or copyright holder(s), unless the work is under an open content license such as Creative Commons.

Takedown policy

Please contact us and provide details if you believe this document breaches copyrights. We will remove access to the work immediately and investigate your claim.

Green Open Access added to TU Delft Institutional Repository

'You share, we take care!' - Taverne project

<https://www.openaccess.nl/en/you-share-we-take-care>

Otherwise as indicated in the copyright section: the publisher is the copyright holder of this work and the author uses the Dutch legislation to make this work public.

Chapter 13

Should Philosophers Begin to Employ New Methods If They Want to Become More Societally Relevant?

Ibo van de Poel

1 Introduction¹

In his essay ‘Philosophy Inside Out’, Philip Kitcher pleads for a transformation in philosophy in order for the discipline to become more societally relevant (again) (Kitcher, 2011). Among others, he argues that topics and questions that are now in the periphery should become more central, and vice versa. He also discusses a number of examples of philosophical inquiries that have in his view made a start in making philosophy more societally relevant and he says about these ‘What binds the endeavors together is a concern for philosophical questions that matter, rather than a shared method’ (Kitcher, 2011, p. 259) .

The question I want to pose in this contribution is whether making philosophy more societally relevant is indeed just a matter of asking other questions, or also has methodological implications. One obvious complication in answering this question is that it is not so clear what the (traditional) philosophical method is and even whether there is a method, or a series of methods, that is characteristic of philosophy. Moreover, this may well be different for different branches of philosophy.

In order to give a meaningful answer to the mentioned question, I will focus on more analytically inclined branches of philosophy, and I will assume that conceptual and analytic methods are central in such branches of philosophy. Now, terms like ‘analysis’ and ‘conceptual’ can still stand for a number of things, but they seem to typically refer to attempts to better understand certain philosophical notions and their relations, and they typically also involve breaking down questions and concepts into smaller parts. In general, analysis may perhaps best be understood as a strategy that decomposes a problem in its component parts; often the decomposition will imply a regression or reduction of the problem to a more fundamental level on which it might be solved (Beaney, 2011).

Analysis may also be understood by what it is not. It is typically not empirical, and it is typically not synthetic. Rather than understanding a problem in context and its complexity, it is often aimed at creating a more simple or clearer version of the problem that lends itself to decomposition and analysis. As a consequence, many analytic philosophers have a tendency to focus on what I would like to call *toy problems*. Toy problems are simplified versions of real-world problems that lend themselves more easily for philosophical analysis.

While a focus on toy problems might facilitate philosophical progress, it may also come at the cost of social relevance. At least, it requires a translation back into what philosophical theories and concepts that were developed on basis of toy problems teach us about real-world problems. I will argue that this translation should not be taken lightly, and requires more than just applying theoretical philosophical insights to real-world problems. Instead, it also has distinct methodological consequences, I will suggest.

In this contribution, I will particularly focus on two such methodological consequences. One is that addressing real-world problems requires an increased attention for what might be called ‘problem formulation’. Rather than assuming that real-world problems are somehow similar to the toy problems they know, philosophers should start with a thorough philosophical analysis of real-world problems in context, and in doing so they can take inspiration from design

and design thinking, I will argue. Second, if philosophy is to contribute to addressing, or even solving, real-world problems, it should also involve design and experimentation, I will suggest.

It must be admitted that these ideas are not entirely new. In recent years, we have, for example, seen an increased attention for empirical ethics (cf. Musschenga, 2005) and for experimental moral philosophy (e.g. Alfano & Loeb, 2016). Also, the role of design in moral philosophy has been the point of attention (van den Hoven, Miller, & Pogge, 2017). These are, however, still developments on the fringes of philosophy rather than turning philosophy inside out. I consider it therefore worthwhile to spell out in some detail why such methodological renewal is needed if philosophy is to become not only accidentally but more structurally societally relevant.

Before I discuss in more detail the two mentioned methodological consequences of turning philosophy inside out, I start with a brief discussion of the use of toy problems in philosophy and the difference between such toy problems and real-world problems.

2 Toy Problems

Toy problems are simplified or stylized examples of real-world problems. They can be found in many disciplines, also outside philosophy; think of prisoner dilemmas in game theory, or the travelling salesman problem in computer science. Toy problems are usually not easy problems; in fact, they are often quite hard problems. The reason to introduce toy problems, rather than directly addressing real-world problems, is often not that they are easy to solve but that they allow to think sharper and more precisely about a problem. Because they are stylized, they often allow for more analytic approaches to come to potential solutions. The term ‘toy problem’ is thus used here in a technical sense, without any pejorative connotations.

Particularly in more analytically inclined philosophy, many examples of toy problems may be found. As an illustration, I will here focus on one particular example: the trolley problem (see, e.g. Thomson, 1985). The trolley problem was originally developed by Philippa Foot

(1967). Particularly in the last decade, many new variations of the trolley problem have been developed, and it has also been taken up outside philosophy, for example, in psychology and neuroscience. In the original trolley problem, there is a runaway trolley that is approaching five people tied to the track and it is about to kill them. However, there is a switch in the track that allows you to send the trolley to another track, where it will kill only one person. You are standing some distance away from the track next to a lever with which you can turn the switch (or not). The question is what you should do. If you pull the switch, you deliberately kill one person (and save five), but if you do nothing, you let five people die. Variations to the trolley problem involve, for example, a scenario where one can throw a fat man (from a bridge) on the track, so that the five people are saved (and the fat man is killed).

Philosophers may have good reasons to introduce toy problems, like the trolley problem. Often a problem becomes clearer if one strips it from its irrelevant characteristics and decontextualizes it. In this way, one can focus on the ‘real’ problem at issue rather than a contingent aspect of it. Another advantage of toy problems might be that they more easily lend themselves for thought experiments, a method popular among some philosophers to develop better philosophical concepts and theories. Anyway, toy problems may often be a useful way to start thinking about complex problems by starting with a simpler and clearer version of the problem.

The focus on toy problems, however, also brings risks and potential disadvantages. One risk is that philosophers may wrongly think that to become socially relevant it is enough to approach real-world (moral) problems as if they are similar to existing toy problems in philosophy. An example is the modelling of the real-world moral problem of how to design self-driving cars as a trolley problem (e.g. Lin, 2016; Bonnefon, Shariff, & Rahwan, 2016). The idea here is that when self-driving cars cannot avoid a collision, they have to make a choice who to kill (or at least to put at risk). It is not difficult to imagine that this choice involves scenarios that are similar to those in trolley problems.

However, modelling the real-world problem of how to morally design self-driving cars as a trolley problem is problematic or even misleading as has been pointed out by several authors (Goodall, 2016; Himmelreich, 2018; Nyholm & Smids, 2016). One issue is that the current state of technology does not allow self-driving cars, or the underlying algorithms, to make the required choices; sensor technologies are, for example, not good enough to reliably detect different types of road users, while the trolley problem framing assumes such possibilities. One may counter, that this may become technologically feasible in the future and is therefore still morally relevant. Even if this be true, what the trolley problem framing *de facto* does is to deflect attention away from more short-term and more urgent moral problems that self-driving cars already raise.

The straightforward application of the trolley problem to self-driving cars then results in what Nordmann (2007) has called speculative ethics. Such an ethics follows an if-then reasoning pattern: It states that if technology X becomes feasible, then moral problem Y will arise. Typically, speculative ethics then tends to take the if clause as given and squarely focuses on moral problem Y. The problem of such a reasoning pattern is not just that it makes speculative assumptions about the future, but also that it tends to forget that what is currently technologically possible may well raise ethical issues other than problem Y. So while speculative ethics may be somewhat useful in anticipating possible future ethical issues, it fails as a method to address current real-world ethical problems in technology.

The framing of the moral problem of self-driving cars as a trolley problem is also problematic in other respects. It, for example, tends to neglect the uncertainty that surrounds real-world problems. Trolley problems have certain outcomes (five people will die if it stays on the current track, for example); in as far as uncertainty is introduced in versions of the trolley problem, it is usually a very stylized form of uncertainty. In the real world, uncertainty is not only much more endemic but we are often unaware of uncertainties, or do not know what type of uncertainty we are confronted with; e.g. whether uncertainty is merely stochastic, or that there are events that may occur we cannot even yet think of.

Another respect in which trolley problems are different from real-world problems is that they present moral problems as clear-cut dilemmas. Although the choice in such dilemmas is very hard (if not impossible), it is assumed that the set of options is given and limited. Moral dilemmas frame moral problems as multiple-choice problems with pre-given options, but in the real world options are almost never given; rather, they need to be constructed or thought out, and we may never be able to say that we have listed all possible options.

It is of course not really amazing that the trolley problem is not a good way to frame real-world moral problems like the self-driving car design problem. After all, the original trolley problem was never intended as a good model for such real-world problems. Rather, it was a deliberate simplification and stylization with the aim to put more light on a specific moral issue. Once this is recognized, there are more sophisticated ways to make trolley problems relevant for real-world problems.

Trolley problems may, for example, help to settle the dispute between utilitarian and deontological moral theories; or they may help to formulate moral principles that address people's moral intuitions about such cases (cf. Kamm, 2016). Such principles may then be applied, for example, in the design of collision algorithms for self-driving cars (Keeling, 2019). This route to make the trolley problem relevant to real-world problems is more indirect, but also more sophisticated, as it does not, at least not necessarily, assume a direct analogy between the real-world problem of self-driving cars and the trolley problem.

Still, I believe that also this route is not entirely unproblematic. One risk is that it carries in it the idea of a dichotomy between what is called 'theoretical philosophy' and so-called 'applied philosophy'. The idea is roughly this: Theoretical philosophy does the real philosophizing based on toy problems, and the insights can then be applied in applied philosophy where context is added. The addition of context is then just seen as 'application', rather than as something requiring philosophical inquiry. This idea is wrong, I believe.

It is wrong because even if we can somehow 'apply' more general (and abstract) philosophical or moral principles to real-world cases, this application is by no means

straightforward. It in fact requires a *specification* of the principles for the case and context at hand, and this specification is certainly not deductive, but instead requires context-specific moral judgements and philosophical investigations (Beauchamp & Sugden, 1984; Richardson, 1997; van de Poel, 2013).

But, perhaps even more important, before we can start thinking about applying more general or theoretical insights to real-world cases, we first need to understand what the problem is in such cases, rather than just assuming that the problem is somehow similar to an existing philosophical toy problem. This requires first of all attention for problem formulation, which is the topic of the next section.

3 Problem Formulation

Toy problems fit analytic methods and approaches in two ways. First, their formulation itself is an example of the application of an analytic approach, as they imply the reduction of a real-world problem to a problem that is considered more fundamental. They also fit an analytic approach in the sense that toy problems lend themselves to types of analysis that cannot be applied directly to real-world problems due to their complexity and messiness.

Below, I will argue that if we want to address real-world problems such analytic approaches are not sufficient, and that we need to employ more synthetic approaches and methods as well. The main reason is that in the case of real-world problems, unlike that of toy problems, it is often initially unclear what the problem is. Real-world problems are typically ill-structured problems surrounded by many uncertainties and ambiguities. They first need to be translated in to well, or at least better, formulated problems before they can become subject to methods like conceptual analysis that analytical philosophers like to apply.

The first step in addressing real-world problems is thus what may be called *problem formulation*. Problem formulation itself does not just require analysis and analytic methods, but also synthesis and synthetic methods, as they are employed, for example, in design. Here, I

understand analysis as breaking down a problem in, more fundamental, elements, while synthesis is aimed at creating something new out of existing elements or components.

Why does problem formulation require not only analytic but also synthetic activities? The answer may be found in the literature on ill-structured problems (Cross, 1989; Rittel & Webber, 1973; Simon, 1973).² Ill-structured problems are usually described as problems that do not have one best solution (Simon, 1973). Instead, they have a range of possible solutions of which some may be better than others, but which may not all be ordered on one scale from better to worse. The main reason for this is that usually different criteria are important in judging solutions to ill-structured problems; options that score well on one criterion may not score well on another. This is a characteristic that ill-structured problems have in common with moral dilemmas, even if the latter are presented as multiple-choice problems. However, while philosophers highlight the ill-structured character of the choice problem that moral dilemmas pose, they tend to neglect the fact that real-world moral problems are also ill-structured in many other respects.

Another characteristic of ill-structured problems is that they are open-ended (cf. Rittel & Webber, 1973). There is no point in time at which the problem is fully and definitively understood and all options are given. Of course, in practice, at some point, one needs to stop looking for new possible solutions and one has to choose what to do. However, this point is usually not reached because the problem has turned into a well-defined problem, although this may occasionally occur, but due to real-world constraints with respect to time and other resources.

The third characteristic of ill-structured problems that has particularly been stressed in the design literature is that typically problem formulation and the development of potential solutions run in parallel (e.g. Cross, 1989). In other words, in design one often does not first have a definitive formulation of the (design) problem and then develops potential solutions, but over time the problem becomes more definite, as one develops potential solutions.

What underlies this third characteristic is the fact that synthetic activities like devising options and tinkering with possible solutions help to better understand the (design) problem. Or

even stronger: A good problem formulation is usually impossible without such synthetic activities. Ill-structured problems are typically so underdetermined that the only way to come to grips with them is to start thinking about possible ways to deal with them (i.e. synthesis) in order to understand the problem better.

More precisely, my argument here is that problem formulation involves synthesis in two distinct ways. One is that in order to better understand ill-structured problems, we need to carry out synthetic activities: Without synthesis, a proper problem formulation is impossible (at least for ill-structured problems). A second is that problem formulation itself is synthetic in the sense that it is not just the discovery of a pre-given problem, but rather involves the *construction* of a problem in an initially indeterminate situation.

The latter aspect has also been highlighted by, for example, the pragmatist philosopher John Dewey. Dewey (1938, pp. 101–19) describes the process of inquiry as always starting with what he calls indeterminate situations, i.e. situations that create unease, in which something is ‘wrong’ or ‘missing’. What is particularly important is that Dewey distinguishes indeterminate situations from problematic situations. In an indeterminate situation, it is clear that something is wrong but what exactly the problem is, is unknown. In fact, for Dewey turning an indeterminate situation into a problematic one, or what I have called problem formulation, is the first step in inquiry. This step for Dewey is to be guided by what he elsewhere calls ‘intelligence’, and it will thus also involve more analytic activities, but it also involves the synthesis or construction of the problem out of, but not fully determined by, the indeterminate situation.

If we accept that in the case of ill-structured problems (or indeterminate situations as Dewey calls them), we need to *construct* rather than discover the appropriate problem formulation, it seems obvious that problem formulation is a synthetic activity. A *good* problem formulation, however, requires both synthesis and analysis. In this sense, the designing engineer who immediately perceives the indeterminate situation of self-driving cars as a technical problem is as mistaken as the philosopher who immediately perceives it as a trolley problem.

Ill-structured problems may not allow for a definitive or ‘best’ problem formulation. However, it does not follow that all possible problems formulations are equally good; in fact, I think we have good reasons to believe that some problem formulations are, in a given indeterminate situation, better than others. Do we then have any criteria for what counts as a good problem formulation?

I think that Dewey’s work suggests at least two possible criteria here. The first is that according to him, we should consider (philosophical) problems always in context, i.e. we should analyse what conditions led to the problem, and what the consequences are of possible courses of action (see, e.g. Dewey, 1939). When we (completely) detach a problem from its context, it may well stop to be a relevant problem, because what makes a situation indeterminate is its linkage to other events, objects and people. Completely abstracting away contextual linkages in problem formulation will either dissolve the problem (but without solving the initial indeterminate situation) or result in a problem that does not address the initial indeterminate situation.

The second criterion is that for Dewey the final step in inquiry is not just thinking out a (possible) solution, but actually exercising this solution and seeing whether this indeed solves the initial indeterminate situation. In his sense, for Dewey inquiry is always experimental. A good problem formulation then is helpful in formulating possible solutions that, when tried, help to overcome the unease or lack in the initial indeterminate situation.

4 The Need for Design and Experimentation

For Dewey, inquiry is not just an intellectual activity; at the core of his notion of inquiry is that it is *transformative*, i.e. inquiry transforms an indeterminate situation into a determinate one (e.g. Dewey, 1938, p. 104). For this reason, inquiry does not end with thinking out a desirable course of action, but by actually exercising this course of action, and seeing whether it works (experimentation). The transformative character of inquiry also explains why Dewey somewhat mysteriously says that practical judgements complete a situation (e.g. Dewey, 1915). For him, a

practical judgement is not just the intellectual conclusion of an inquiry, but rather it consists in actually carrying out the concluded action, and so making the initial indeterminate situation determinate. Of course, the action may fail to complete the situation, and in that case, the practical judgement has turned out to be wrong, and we need to continue our inquiry.

What makes Dewey's description of inquiry particularly relevant for the current purpose is that it is aimed at addressing real-world problems rather than toy problems. Even if one might disagree with Dewey on some points, what stands out is that philosophical inquiry that addresses real-world problems is not just intellectual and mental but is transformative. It involves and requires making changes in the real world, and for that reason, I believe that design and experimentation have a proper place in philosophical inquiries that are aimed at addressing real-world problems.

Let me first focus on design. I use the term design here to describe planned and deliberated activities aimed at changing the world in a particular way that is deemed desirable (cf. Simon, 1979). This notion of design encompasses technical (or engineering) design but also, for example, educational or policy design and the design of institutions. Design involves analysis and judgement in the sense that it is based on an analysis of how things actually are and a judgement about how things should be or can be improved. It also involves synthesis, in the sense of creating new artefacts or institutions that offer new options for action.

Since design is aimed at changing (and improving) the world, it has the same direction of fit as moral judgements. It is not aimed at descriptions that fit the world as it is, but rather at fitting the world to descriptions about what is desirable and good. Nevertheless, the way in which designers operate, and the methods they employ, is usually quite different from those of philosophers. The best way to explain the difference is perhaps to sketch the prototypical reaction of philosophers and designers to dilemmatic situations, like moral dilemmas (van de Poel, 2017a; van den Hoven, Lokhorst, and van de Poel, 2012).

The typical inclination of the philosopher in such situations is to delve deeper into the normative principles that underlie the dilemmatic choice; are there perhaps conflicting values or

moral principles at play that recommend (or even require) conflicting courses of action? And if so, does one normative principle or value perhaps trump the other or should we somehow balance the values? The philosopher's quest may also involve a consideration of whether values are incommensurable, and if so what that implies for the possibility of (rational) choice. As might be clear, the philosopher's approach is largely analytic in nature; it is aimed at finding the underlying problem and solving it at this deeper level, e.g. by arguing that one value trumps another so that the dilemma is solved.

The prototypical approach of a designer is completely different. The designer will typically start looking for new options and for solutions that take the sting out of the dilemma. So designers would, for example, propose a new option that accommodates both horns of the dilemma. So, when designers are confronted with trolley problems, they will propose to change the design of the tracks, or to add warning lights, or to provide the trolley with automatic detection systems that stop the trolley if there is something on the track. Philosophers will typically perceive such solutions as cheating, and indeed if we understand the trolley problem as a toy problem in moral philosophy such design responses are beside the point, because the trolley problem (as a toy problem) was deliberately devised to reveal a *moral dilemma*. However, if the trolley problem is approached as a real-world problem, the designer's response is perfectly appropriate, and it may reveal solutions to real-world moral problems that the philosopher may well overlook.

The take-away from the above is not that designers are better in addressing real-world problems than philosophers, or the other way around. Rather, my point is that addressing real-world problems requires both analytic methods and skills, in which philosophers often excel, as well as more synthetic methods and skills, which might be more the strength of designers.

Design is still largely a mental activity. It involves thinking out courses of actions and imagining ways in which a situation can be changed, but design is not yet the exercise of these courses of action; design itself does not yet rearrange the material and social world. If inquiry of real-world problems is to be transformative, we need therefore experimentation as well.

I will employ a broad notion of experimentation here, in which experimentation is understood as trying out something (e.g. a course of action), observing the consequences, and learning from it (cf. van de Poel, 2017b). This learning may be theoretical (as in a scientific experiment), but it may also be aimed at improving a certain course of action or intervention. Some have characterized the latter type of experiments as design experiments (e.g. Stoker & John, 2009). A design experiment is aimed at gradually improving an intervention (or course of action) by a repeated process of trying it out, observing its consequences, learning and then adapting the initial course of action, to try it out in a next experimental round. Such design experimentation typically would proceed in small steps (also to reduce the risks of experimentation), not unlike how Karl Popper (1945) has described the idea of piecemeal social experiments.

Good experimentation thus involves more than just trial-and-error. It is guided by proper problem formulation, analysis and design. It is not just trying out some action or intervention, but first requires thinking out an option of action that we have good reasons to expect to solve the initial indeterminate situation. Experimentation then comes at the final stages of inquiry, but it is an indispensable part of philosophical inquiries into real-world problems.

The reason why experimentation is indispensable is that real-world problems can, unlike toy problems, not be detached from their context. They occur in open systems that we do not, and perhaps cannot, fully understand and predict. Real-world problems are therefore always characterized by uncertainty, ambiguity and unknowns. For this reason, the final step of inquiry requires not just mental experimentation (e.g. through thought experiments or a computer simulation) but a real-world experiment as well. The experiment may fail, that is to say, it may not have the expected or desired outcome, but if it is used to learn and as a means for a new, and better, experiment, it has succeeded nevertheless despite its apparent failure.

5 Conclusion

There are many ways in which philosophy can be socially relevant and Kitcher may be right that it typically starts with asking the right questions. However, I think there are some methodological implications as well. I have argued that analytic methods alone are not enough to address real-world problems in philosophy. If philosophy wants to contribute to addressing real-world problems, it should pay more attention to problem formulation, and to design and experimentation.

Still, some caveats are in place with respect to this conclusion. One important caveat is that my argument is not only based on the assumption that it is desirable that philosophy becomes societally (more) relevant, but also on the premise that it should do so by better addressing real-world problems. Even if one accepts the first assumption, one may still reject the second, and argue that there are other ways in which philosophy can be – or become – societally (more) relevant. One might, for example, argue that societal relevance can be achieved not so much through contributing to the solution of real-world problems but rather through philosophy as a form of criticism. In fact, criticizing existing beliefs, or for that part potential solutions to real-world problems proposed by other disciplines, has always also been an important role for philosophy in society.

And even if one believes that philosophers should contribute to addressing real-world problems, one may point out that they need to do so within a broader division of intellectual labour. Social problems cannot be solved by philosophers alone, but rather require multidisciplinary cooperation. So the real question is probably not how philosophers should solve societal problems, but instead what their proper role is a larger division of labour aimed at addressing societal problems. My own experience (particularly in the field of design for values and responsible innovation) here is that one of the main added values of philosophy in such multidisciplinary cooperation is still in more traditional philosophical virtues like a certain rigor and sharpness in how concepts are being used. So my argument should certainly not be understood as one against the use of analytic methods or analytic rigor. Nevertheless, even within a broader division of labour, I think philosophers need some awareness of the importance

of problem formulation, design and experimentation because without such awareness, they will also not be able to appreciate the contribution of other disciplines, and to cooperate productively with others.

One final consideration is that if philosophers indeed become more societally relevant and more involved in addressing real-world problems, this might also lead to dirty hands problems. Their analyses will no longer be purely academic, but will make a difference in the real world. This increases their potential to do good, but also their potential to harm others. Moreover, if philosophy matters for the way real-world problems are addressed, philosophers will also become part of the power struggles that surround such problems. So one of the consequences is that philosophers would need to think more about the ethics of doing philosophy and ethics (cf. Hansson, 2017). This is of course not a reason to refrain from a role in addressing real-world problems, but nevertheless something one should be aware of.

Notes

References

Alfano, M., & Loeb, D. (2016). Experimental moral philosophy. In E. Zalta (Eds.), *Stanford encyclopedia of philosophy*. Retrieved from

<http://plato.stanford.edu/archives/spr2016/entries/experimental-moral/>

Beaney, M. (2011). Analysis. In E. Zalta (Eds.), *Stanford encyclopedia of philosophy*. Retrieved from <http://plato.stanford.edu/archives/sum2011/entries/analysis/>

Beauchamp, T. L., & Sugden, S. J. B. (1984). On eliminating the distinction between applied ethics and ethical theory. *Monist*, 67(4), 514–31.

<https://doi.org/10.5840/monist198467430>

- Bonnefon, J.-F., Shariff, A., & Rahwan, I. (2016). The social dilemma of autonomous vehicles. *Science*, 352(6293), 1573–76. <https://doi.org/10.1126/science.aaf2654>
- Cross, N. (1989). *Engineering design methods*. Chichester: Wiley.
- Dewey, J. (1915). The logic of judgments of practice. *The Journal of Philosophy, Psychology and Scientific Methods*, 12(19), 505–23.
- Dewey, J. (1938). *Logic, the theory of inquiry*. New York, NY: Rinehart and Winston.
- Dewey, J. (1939). *Theory of valuation*. Chicago, IL: Chicago University Press.
- Foot, P. (1967). The problem of abortion and the doctrine of double effect. *Oxford Review*, 5, 5–15.
- Goodall, N. J. (2016). Away from trolley problems and toward risk management. *Applied Artificial Intelligence*, 30(8), 810–21. <https://doi.org/10.1080/08839514.2016.1229922>
- Hansson, S. O. (2017). The ethics of doing ethics. *Science and Engineering Ethics*, 23(1), 105–20. <https://doi.org/10.1007/s11948-016-9772-3>
- Himmelreich, J. (2018). Never mind the trolley: The ethics of autonomous vehicles in mundane situations. *Ethical Theory and Moral Practice*, 21(3), 669–84. <https://doi.org/10.1007/s10677-018-9896-4>
- Kamm, F. M. (2016). In E. Rakowski (Ed.), *The trolley problem mysteries*. New York, NY: Oxford University Press.
- Keeling, G. (2019). Why trolley problems matter for the ethics of automated vehicles. *Science and Engineering Ethics*, 7(3), 149. <https://doi.org/10.1007/s11948-019-00096-1>
- Kitcher, P. (2011). Philosophy inside out. *Metaphilosophy*, 42(3), 248–60. <https://doi.org/10.1111/j.1467-9973.2011.01684.x>
- Lin, P. (2016). Why ethics matters for autonomous cars. In M. Maurer, J. C. Gerdes, B. Lenz, & H. Winner (Eds.), *Autonomous driving* (pp. 69–85). Berlin: Springer.

- Musschenga, A. (2005). Empirical ethics, context-sensitivity, and contextualism. *The Journal of Medicine and Philosophy*, 30(5), 467–90. <https://doi.org/10.1080/03605310500253030>
- Nordmann, A. (2007). If and then: A critique of speculative nanoethics. *NanoEthics*, 1(1), 31–46. <https://doi.org/10.1007/s11569-007-0007-6>
- Nyholm, S., & Smids, J. (2016). The ethics of accident-algorithms for self-driving cars: An applied trolley problem? *Ethical Theory and Moral Practice*, 19(5), 1275–89. <https://doi.org/10.1007/s10677-016-9745-2>
- Popper, K. (1945). *The open society and its enemies*. Princeton, NJ: Princeton University Press.
- Richardson, H. S. (1997). *Practical reasoning about final ends*. Cambridge: Cambridge University Press.
- Rittel, H., & Webber, M. (1973). Dilemmas in a general theory of planning. *Policy Sciences*, 4, 155–69.
- Simon, H. A. (1973). The structure of ill structured problems. *Artificial Intelligence*, 4(3–4), 181–201. [https://doi.org/10.1016/0004-3702\(73\)90011-8](https://doi.org/10.1016/0004-3702(73)90011-8)
- Simon, H. A. (1979). *The sciences of the artificial*. Cambridge, MA: MIT Press.
- Stoker, G., & John, P. (2009). Design experiments: Engaging policy makers in the search for evidence about what works. *Political Studies*, 57(2), 356–73. <https://doi.org/10.1111/j.1467-9248.2008.00756.x>
- Thomson, J. J. (1985). The trolley problem. *The Yale Law Journal*, 94(6), 1395. <https://doi.org/10.2307/796133>
- van de Poel, I. (2013). Translating values into design requirements. In D. P. Michelfelder, N. McCarthy, & D. E. Goldberg (Eds.), *Philosophy and engineering: reflections on practice, principles and process* (pp. 253–66 15). Dordrecht: Springer.

van de Poel, I. (2017a). Dealing with moral dilemmas through design. In J. van den Hoven, S. Miller, & T. Pogge (Eds.), *Designing in ethics* (pp. 57–77). Cambridge: Cambridge University Press.

van de Poel, I. (2017b). Moral experimentation with new technology. In I. van de Poel, L. Asveld, & D. C. Mehos (Eds.), *New perspectives on technology in society* (pp. 59–79). New York, NY: Routledge.

van den Hoven, J., Lokhorst, G.-J., & van de Poel, I. (2012). Engineering and the problem of moral overload. *Science and Engineering Ethics*, 18(1), 143–55.
<https://doi.org/10.1007/s11948-011-9277-z>

van den Hoven, J., Miller, S., & Pogge, T. (Eds.). (2017). *Designing in ethics*. Cambridge: Cambridge University Press.

1 This publication has been written as part of the project ValueChange that has received funding from the European Research Council (ERC) under the European Union’s Horizon 2020 research and innovation programme under grant agreement no 788321. I would like to thank Michael Klenk for comments on an earlier version; all claims and possible mistakes are my own responsibility. The text contains some (very) short excerpts from an earlier published book chapter on ‘dealing with moral dilemmas through design’ (Van de Poel, 2017a).

2 These authors have different disciplinary backgrounds and offer different accounts of what I call ill-structured problems (their terminology is sometimes different). The discussion below does not pretend to offer an account of ill-structured problems that all these

authors would agree with. Rather, it presents what I take to be the main characteristics of such ill-structured problems.