Techné: Research in Philosophy and Technology ISSN: 2691-5928
26:2 (2022): 1–189 DOI: 10.5840/techne2022727162

"Structural Ethics" as a Framework to Study the Moral Role of Non-Humans

Ehsan Arzroomchilar

Abstract: A challenging issue within the philosophy of technology is the moral relevancy of artifacts. While many philosophers agree that artifacts have moral significance, there are numerous positions on how moral relevancy ought to be understood, ranging from scholars who argue that there is no room for artifacts in moral debates to those who argue for the moral agency of artifacts. In this paper, I attempt to avoid extreme positions; accordingly, I reject both the neutrality thesis and the moral agency of artifacts thesis. Instead, I propose finding a compromise for describing their moral role. In doing so, I take Philip Brey's idea of developing a new framework, called 'Structural Ethics,' as my point of departure. Although the structural ethics proposed by Brey needs some revisions, it may serve as a proper metaethical theory to account for the role of non-humans.

Key words: structural ethics, moral agency, mediation theory, morality

1. Introduction

In recent decades, there have been extensive debates on how to study the moral impacts technological artifacts may have on human beings. Some scholars have argued that there is nothing special about artifacts compared to other objects, concluding that they are morally irrelevant. However, other scholars have ventured to ascribe a kind of agency to artifacts. In this account, technological artifacts condition humans' lives and thereby effect their morality to such an extent that one can sensibly count them as agents. For instance, philosopher of technology Peter-Paul Verbeek (2005; 2011) speaks of 'distributive agency.' Accordingly, agency is not the property of an isolated human. Instead, it is a combination of human-plus-

Ehsan Arzroomchilar Department of Philosophy and Religious Studies, University of South Bohemia, Faculty of Theology Kněžská 8 370 01 České Budějovice Czech Republic; arzroe00@tf.jcu.cz

technology.[1] Bruno Latour (e.g., 1993; 2005), a STS scholar, openly treats non-humans on a par with humans in his 'symmetrical' approach. Further, Luciano Floridi and J. W. Sanders (2004) take some sophisticated 'artificial' technologies to be real agents.

Given that the literature comprising arguments for and against these extreme positions is fairly rich, I will not include this discussion in detail within this article.[2] Instead, I will argue that neither the neutrality thesis nor the moral agency thesis regarding the moral relevancy of artifacts can do justice to their actual role in the daily lives of humans. I suggest that we can find a compromise between these two extremes. With Philip Brey, I believe that the best way to address the moral relevancy of artifacts is to develop a new framework to study their moral impacts alongside pervasive, individualistic ethics (Brey 2014, 135). Within contemporary ethics, human behavior is studied in itself, regardless of the influences of non-human factors. A sentient human's action, regardless of her connection with her context, is the only thing that deserves our attention, from the contemporary ethics' point of view.

I will try to show that developing a structural ethics parallel to individualistic ethics would overcome some of the difficulties that artifacts may impose on morality. As my point of departure, I will elaborate upon Brey's recent proposal of 'structural ethics.' Such an approach will pave the way to study the moral role of non-human factors more fruitfully, without falling into the contradictory trap between moral agency and the neutrality. In Section 2, I will critically evaluate Brey's proposal of structural ethics and try to fine-tune it by substantiating its content and removing some problematic dimensions. In Section 3, I will go through the advantages of such an approach. Finally, Section 4 will sketch out the relationship between structural ethics and traditional ethics, namely deontological ethics and consequentialist ethics.

2. Elaborating Brey's Structural Ethics

Leaving out some details,[3] I will take up and critically elaborate upon Brey's approach to structural ethics. Its core idea is that interacting humans and non-humans may constitute a network whose moral status may be explored. At the beginning of his paper, he describes this approach as follows:

> Structural ethics studies social and material arrangements as well as components of such arrangements, such as artifacts and human agents. It has three aims: (1) to analyze the production of moral outcomes or consequences in

> existing arrangements and the role of different elements in this process; (2) to evaluate the moral goodness or appropriateness of existing arrangements and elements in them; and (3) to normatively prescribe morally desirable arrangements or restructurings of existing arrangements. (Brey 2014, 135)

Then, he proceeds to construct a new vocabulary needed for such a framework. As a preliminary step, he starts by defining 'network' as " structures of interacting entities that together determine outcomes or actions that are the subject of moral evaluation" (Brey 2014, 137). This network or structure can involve many elements; he calls elements inside of a network 'factors' to differentiate them from agents.

In the next step, he introduces three twofold terms, indicating types of factors involved in a network: 1. Positive factor versus negative factor; 2. Accidental factor versus intentional factor; and finally 3. outcome-oriented factor versus behavior-oriented factor. The characteristics of these notions, in turn, go as follows:

> [M]easured against a moral rule or principle, a positive moral factor is one that contributes positively to a moral principle being upheld, whereas a negative moral factor contributes negatively. . . . An accidental moral factor is one that happens to contribute towards a moral outcome in a particular arrangement. An intentional moral factor [by contrast] is one that has been intended to contribute to an outcome in a particular way. (Brey 2014, 137f.)

He takes a speed bump or a traffic controller as examples to make the definitions tangible: "relative to the moral outcome of cars driving safely, a speed bump and a traffic controller would both be intentional moral factors, whereas a pothole that causes cars to drive slowly would be an accidental moral factor" (Brey 2014, 137).

Concerning the third set of terms, Brey continues:

> An unjust distribution of goods that results from an action . . . is a moral outcome. . . . Various moral factors can be identified as having caused these outcomes [which are called outcome-oriented]. A behavior-oriented moral factor is one that influences the moral behavior . . . of an agent. For example, my wearing a seat belt is a moral action that is influenced by various moral factors, such as blinking lights on my dashboard and police officers who check on seat belt use.
>
> Putting it simply then, a factor is a behavior-oriented one if it co-shapes the behavior of agents—who are exclusively human—and it is an action-oriented one if it directly contributes to the outcome of the networks. (Brey 2014, 138)

Having outlined Brey's notion of structural ethics, I can now critically analyze it to show further opportunities for elaboration and enhancement. I suggest two ways to expand on and improve Brey's structuralist ethics: The first is a critique of his distinction between outcome-oriented and behavior-oriented factors. The second is meant to develop a new concept to better illustrate the role of factors within networks.

On the first issue, Brey's conceptualization of the structural framework seems contrived. As defined, any factor in a network would be both action-oriented and behavior-oriented simultaneously. To explain, a gun, to use a paradigmatic example, within a network comprised of a gun and a person, not only helps to co-shape the outcome of the network, namely killing someone, but it also influences another factor involved: mediating the human agent who is taking revenge. This also holds true for the human agent. While she contributes directly to the outcome, killing someone, she is simultaneously manipulating the gun to do its inscribed job—shooting. In this sense, there seems to be some symmetry between agents and non-agent factors, contrary to what Brey maintains. All mediators within a network contribute to the ultimate outcome while mediating other factors. Thus, it seems that this distinction is bereft of any use in practice and, accordingly, may be discarded.

My second point concerns the degree of embeddedness of factors in a network. To morally appraise a network, and consequently mitigate its undesirable outcomes, it is crucial to provide a benchmark by which we can measure the extent any factor may contribute to the network's outcome. Clearly, some factors are more vital than others for generating an outcome. Therefore, the significance of the factors, concerning their contribution, is a matter of degree. In taking revenge on a person by using a gun, the weight of these two factors is not the same. The human agent, initiating the network into the action, after all should be placed on a higher level regarding her contribution to the outcome, namely killing someone, compared to the gun.

I propose introducing a new concept, 'degree of embeddedness,' which indicates the level of a network's dependency on a specific factor. The initial description of the concept can go as follows:

> *We say, in a network N and dependent on the outcome O, the degree of embeddedness, D, of the factor F1 is greater than that of the factor F2, if and only if, all other conditions remaining the same, F1 is less negligible than F2. That is to say, if the outcome of a network due to removing F1 changes to a*

higher degree in comparison to F2, then the degree of embeddedness of F1 is greater than F2.

Two points should be noticed here regarding this definition:

1. This definition is a relational description rather than an absolute one—meaning that, it determines a factor's degree of contribution compared to others. It is not supposed to report the absolute degree of a factor's contribution.

2. Provided that removing a factor from an actual working network in reality is impossible on most occasions, this definition should be taken as a thought experiment rather than a real-life experiment. Through careful imagination and brainstorming, one may understand the significance of each factor.

Incorporating such a concept into the vocabulary of structural ethics would eventually enable us to determine the degree of moral contribution each factor can make through generating an outcome in a network. The greater degree of embeddedness a factor gains, the higher moral share it contributes to the ultimate outcome. In the extreme case, removing the human agent, the factor with the greatest degree of embeddedness will lead to the greatest changes in the outcome of the network. Applying the concept 'degree of embeddedness' to a network results in a hierarchy: the hierarchy may benefit us in identifying the more influential factors when we face a network generating morally undesirable outcomes. As a result, to morally mitigate a network, we would first correct the factors at the top of the hierarchy. In other words, the 'degree of embeddedness' provides practitioners with specific instructions to approach any structure to decrease undesirable consequences of the interactions within that particular network.

The degree of embeddedness may also serve a further application: As I have said earlier, a multitude of different factors are enrolled in networks, ranging from "humans, artifacts, animals and natural objects," as Brey notes, to "larger structures composed of such entities" (Brey 2014, 137). If this is indeed the case, the number of players in a network would soon get out of hand. One might even say, after all, that every single factor is a representative of a prior network. For example, it might be argued that the network of a man with a gun is not just a matter of human-plus-gun. Rather, it comprises a human and a prior network of which the gun is an outcome, namely the manufacturers, the laws governing its usage, the people engaged along the way of manufacturing it, and many others. If the network generated an unethical outcome, like killing an innocent person, all of those prior networks' elements would also bear responsibility. Hence, this would lead to rapid growth in the number of elements involved. The chain would be endless. How

may one put an end to this chain? Where can the elements be cut and the networks closed? The concept of 'degree of embeddedness' might be useful: The degree of embeddedness may help one to identify the threshold where the significance of the individual element's contribution is sufficiently low to be negligible. We would not need to go down to the end of the seemingly endless sequence of elements to analyze the networks. We would be able to disregard the elements from a certain point. However, the positioning of this threshold is a pragmatic issue: It depends on one's vision as to how thoroughly the moral condition of a network ought to be described. Put simply, when applying the degree of embeddedness, practitioners may disregard the elements with the lowest contribution and then look into the rest of the network. In doing so, one may avoid infinite networks and arbitrarily closing the networks.

3. Merits of a 'Structural Ethics'

As mentioned earlier, Brey's novel proposal to establish structural ethics is a way to shed some light on the moral relevancy of non-humans. In this section I try to highlight its advantages in more detail:

1. Structural ethics, as it has been shown, paves the way for addressing the moral relevancy of artifacts without reducing the human or the artifact to the other. Further, it does not exaggerate either role. Although there are few who believe that technology does not impact us in contemporary scholarship, there is no consensus among those who do believe in technology's potential on how to articulate its impacts. Generally speaking, there are some scholars who consider artifacts as playing no significant role in morality, a position called instrumentalism. Other scholars assume a kind of symmetry between humans and artifacts concerning their agency. According to the latter, artifacts, on a par with humans, should be incorporated into contemporary ethics. Instrumentalism, however, overlooking the impacts of artifacts thoroughly, does not feel a need for any change in the structure of contemporary ethics. According to the moral agency thesis, artifacts are the 'missing masses'[4] in contemporary moral considerations. For instrumentalists, they maintain that 'Guns don't kill, people kill.'

Structural ethics can appropriately settle the fight between these two positions by proposing a compromise. While it believes in the contribution of non-humans to morality, structural ethics does not put non-humans and humans on the same level.

2. Despite taking the significant role of non-humans into moral consideration, structural ethics, as I have mentioned earlier, recognizes the fundamental differ-

ences between humans and non-humans. Traditionally, the notion of 'agency' has been used to refer to intentional human beings. In a different context, to address non-humans, for instance, the notion starts to lose its original meaning, which, in turn, implies putting humans and non-humans at the same level of moral consideration. Non-humans can never qualify as agents. According to the structural approach to ethics, in contrast, while there seems to be an urgent need to revise some basic moral notions to acknowledge the role of non-humans, one should not go as far as demoting the humans' role to that of non-humans. That being said, we need to preserve the key notion of 'agency' in its purely humanistic meaning.

3. By applying a structural framework, we may also clarify some traditionally problematic cases, such as Plato's example of Cephalus. Cephalus defines 'justice' as speaking the truth and paying one's debts. Socrates quickly refutes this account by suggesting that it would be wrong to repay certain debts, such as returning a borrowed weapon to a friend who is not in his right mind.

In light of this new context, we can now address the dilemma more fruitfully: Here, Socrates is approaching this dilemma from a structural perspective by regarding the weapon and the person as a network—the result of which would be murdering someone. Cephalus, on the other hand, approaches it from an individualistic ethics perspective, concentrating on the mere person, regardless of her relationship to other factors. For Socrates, returning the weapon is morally wrong, because the outcome of the network would be unethical. Simply speaking, in this dispute, the debaters are standing on different levels of analysis, that is, the micro and macro levels, respectively.

4. What Is 'Structural Ethics'? A Moral Theory or Metaethics?

Now, we should be asking: Is the structural account a moral theory? To answer this question, we need to take one step back and ask what a moral theory, in general, is like. In Mark Timmons's view, any conceptual framework that is supposed to serve as a moral theory pursues two aims: one practical and the other theoretical (Timmons 2012, 3). While the practical goal "has to do with the desire to have some method to follow when, for example, we reason about what is right or wrong," the theoretical aim "has to do with coming to understand the underlying nature of right and wrong, good and bad" (Timmons 2012, 3). Besides, in his view, ordinarily, there is a 'moral principle,' in the shape of a 'decision procedure,' which serves as a bridge between these two aims (Timmons, 2012, 3). Firstly, the 'moral principle' points to the underlying feature of all malicious actions, making all of them morally wrong. Secondly, it is supposed to deliver an instruction through which one

comes to know, in certain circumstances, what would be morally wrong, obligatory, or optional to do. For example, suppose a moral principle of a specific moral theory goes as follows:

An action is right if and only if the action does not interfere with the wellbeing of those individuals who are likely to be affected by the action.

Here, we encounter two distinct precepts: first, an intimate connection between a right action and the feature of 'not interfering with the well-being of certain individuals'; second, a criterion by which one can recognize, in a certain situation, the right thing to do.

Now, we can investigate structural ethics against this schema. Initially, the question, 'is structural ethics a moral theory?' breaks down into two further questions: a. Can structural ethics bring a theoretical aspect accounting for the underlying feature of all morally wrong actions? b. Can structural ethics offer a practical aspect pinpointing the morally right actions under certain circumstances?

Brey's elaboration of structural ethics suggests neither a theoretical nor a practical goal. Which means that it cannot response to the two questions above. Further, it does not offer a moral principle as a decision-making procedure. Consequently, structural ethics does not qualify as a moral theory. So, arguably, none of the modern ethical theories can be replaced with structural ethics. As Brey admits, this framework can only serve in a complementary role to the existing moral theories rather than rivaling them (Brey 2014, 135).

In contemporary ethics, two major opposing approaches are commonly distinguished—deontological ethics and consequentialist ethics. While the latter bases morality on the consequences of actions, the former emphasizes the very action or rule itself regardless of the consequences. How can one situate structural ethics in relation to deontology and consequentialism if it is not their competitor? Given that the structural account is not a moral theory on its own, maybe it could be integrated into existing ethical theories and enrich contemporary ethics. Thus, our first inquiry examines whether structural ethics can be consolidated into the existing approaches.

To incorporate the concept of network into these approaches, we would have to replace the concept of 'human individual' with the concept of 'network' or 'structure.' Let me first begin with a simple formulation of both contemporary ethical lines of thought to see if they can be cast into a structure-oriented mode. One may present 'consequentialist ethics' in the following two formulations:

a. An action A is obligatory if and only if (and because) A would produce a higher level of utility than would any other alternative action that the agent could perform instead; and,

b. An action A is wrong if and only if (and because) A would produce less utility than would some other alternative action that the agent could perform instead.

And 'deontological ethics' may be articulated in the following way:

c. Act only in accordance to a maxim that you would accept as a universal law.

However, translating existing ethics into structural ethics is not that easy. Most modern terms bear humanistic connotations: to highlight, terms such as 'action' and 'responsibility,' not to mention 'will' and 'intention,' are heavily bound to human agents. These notions have traditionally been applied to describe individual human affairs rather than arrangements or structures. In this sense, modern ethics is inherently individualistic, one might say. Central to both consequentialism and deontological ethics is an intentional, autonomous human being who holds full responsibility for her free 'actions.' One's 'action' is exactly what is supposed to be, subject to moral evaluation either on its own (deontological approach) or in its results (consequentialist approach). Donald Davidson, for example, famously asserted that an 'action' is something an 'agent' does that is 'intentional' (1980). In this sense, being aware of what is being done and making a decision to get something done would be needed whenever an action is going to be taken. In this case, one might not easily replace the 'intentional agents' with supposedly 'unintentional structures.'

Does this mean that we have failed to develop a structural ethics? I do not think so. True, modern ethics' foundation is humanistic, but a competing view can be provided. Why not reconsider concepts such as 'action' and 'responsibility' and let structures play a role in ethical discussions? If technology, as well as other social and natural factors, effect human behavior, it seems possible to revise some foundational notions and recast ethics in a new shape. This kind of discussion is normally considered to be part of 'metaethics.' As Geoffrey Sayre-McCord (2012) articulates it, "metaethics is the attempt to understand the metaphysical, epistemological, semantic, and psychological, presuppositions and commitments of moral thought, talk, and practice." I assume one may imagine a structural variant of metaethics. This raises a question though. If, as argued, the metaethics of structural ethics is not compatible with that of modern ethics, how substantially different would the structural metaethics be from the current one? In response, I believe that one does not need to, or rather should not, redefine all metaethical terms attached

to modern metaethics. In particular, we do not need take non-humans to be moral agents for this to be possible. One should preserve the notion of 'moral agency' as an inherently human affair. In contrast, there are two notions in need of revision: 'responsibility' and 'action.' If non-humans influence humans, why not extend the notion of 'responsibility' to include non-humans as well? In this way, we would end up with a sort of 'distributive responsibility,' according to which responsibility would not be merely humanistic but rather a 'collective' affair. Structures, therefore, may be said to hold responsibility as well. On the subject of 'action,' too, one may extend the notion to cover not just intentional doings, but the outcomes of networks as well. In determining a location for building a house, after all, we normally take of variety of factors into consideration. House-making is not just a matter of 'human intention.' Several factors are involved ranging from the distance from a volcano, as a natural factor, to the technologies available for construction. Why not describe the 'act of building' as being facilitated by a multitude of factors, rather than merely by humans? Why ascribe the action exclusively to the human part? Not only humans act, but structures do as well. Before going any further to see how this new metaethics would support structural ethics, some points are in order.

First, structural ethics, as discussed before, does not aim to replace modern ethics. If this is the case, we would have different types of ethics at two different levels. At the macro level, the level of structures, and at the micro level, the level of human individuals.

Second, how are the macro and micro levels relevant, one might ask? And more generally, how are the outcomes of networks, on the one hand, and actions stemming from human individuals, on the other hand, related to each other? How are 'responsibility,' 'agency,' 'action,' and morality generally connected at these two levels?

As I have mentioned earlier, we have good reasons to retain the notion of 'agency' as purely humanistic. For one thing, intuitively speaking, we take humans to be capable of purposefully initiating an action. It is hard, if not impossible, to imagine a non-human as initiating an action or as being a prime mover. If we allowed 'agency' to be extended to cover non-humans as well, we would end up with the absurd conclusion that there would be no inherent difference between humans and non-humans. We tend to distinguish between intentionally 'closing a door' and the blind movement of the wind when it 'blows the door shut.' For this precise reason, we will always need individualistic, modern ethics. 'Agency' must retain its purely humanistic sense, while it is possible to reconceptualize other

notions to establish a new metaethical theory. In sum, while there may be 'collective responsibility' and 'collective action' (namely, the outcome of a network), pointing to what is meant to be structural, I cannot imagine 'collective or structural or distributive agency.' Besides, the responsibility of structures or 'collective responsibility,' and 'collective action,' are not reducible to human individuals, for the simple reason that non-human parts, too, 'contribute' to bringing about a state of affairs.

Now what about morality? If we could revise some metaethical notions, we would produce a revised morality. If the concepts of 'collective responsibility'[5] and 'collective action' are legitimate, one may refer to the 'morality of structures' as a sort of 'collective morality.' While agency, in the sense of intentionally performing an action, is an exclusively human affair, morality is not necessarily. Artifacts and generally non-humans, too, are morally relevant given their mediation of human moral actions.

Thus, we return to the claim from the first section: structural ethics stands in contrast with both the neutrality thesis and the moral agency of artifacts thesis. In this sense, provided structural ethics allows for 'collective' responsibility and action while simultaneously rejecting 'collective agency,' it sets its framework in the middle between these two positions. On the one hand, if it were possible to reduce the 'collective morality' of networks to the morality of the human(s) involved, there would be no need for structural ethics. If evaluating network outcomes were possible in terms of evaluating the human factors' actions, structural ethics could be reduced to existing ethics. Otherwise, if 'collective agency' were legitimate, we would end up holding the position of the moral agency of non-humans. What differentiates a proponent of structural ethics from one of the neutrality theses is a belief in a kind of 'collective morality'; whereas, what differentiates her from a proponent of the moral agency thesis is a rejection of 'distributed agency.'

If one finds for the previous argument convincing, one is prepared to take the final step in developing structural ethics. We may draw on the basic pillars of modern ethics and apply them to the structural account. It follows that moral evaluation can generally be imagined in terms of two distinct approaches: assessment of the consequences of an action or assessment of the very action or rule itself. When applied to the structural account, we will eventually attain two distinct strands: First, the consequentialist account of structural ethics, which may go as follows:

a'. An outcome O of a network N is obligatory if and only if (and because) O would produce a higher level of utility than any other alternative outcome that the network could perform instead. And,

b'. An outcome O of a network N is wrong if and only if (and because) O would produce less utility than some other alternative outcome that the network could perform instead.

For a deontological account, we need a foundation for which the expression 'will that it would become a universal law' in the original formulation makes sense. I have borrowed the expression 'people in a suitably described society'[6] from John Rawls (1999), whereby we ultimately come to this formulation:

c'. A network N should act only according to that maxim by which we as 'people in a suitably described society' would will that it become a universal law.

The structural ethics presented here deals with moral evaluation at a macro level. At the micro level, one can keep assessing human individuals from the moral point of view against the current ethical schools. Additionally, thanks to structural ethics, one is also equipped with a new framework to assess the complex combinations of humans and non-humans at the macro level. This new perspective seems to be crucial, especially for engineers, practitioners, manufacturers, or policymakers involved in developing or designing artifacts.

5. Conclusion

I began with the contention that neither the moral agency of artifacts thesis nor the ethical neutrality of artifacts thesis does justice to the role artifacts play in real life. Consequently, I tried to find a compromise; I endorsed Philip Brey's proposal to develop structural ethics in addition to individualistic ethics. I embraced his idea, with some modifications, as a departure point. More importantly, I introduced a criterion to measure every factor's 'degree of embeddedness' in a structure. This concept is needed particularly for networks generating morally undesirable outcomes. As a result, we would have to be aware of the portion of each factor contributing to the outcome to reshape the network.

I have also argued that structural ethics is neither a moral theory itself nor is it easily capable of being assimilated into modern ethics. As a result, structural ethics may serve as a metaethical theory by which one can create a more vivid picture of the moral status of networks whose constituents are both humans and non-humans. I also mentioned that while the notion of 'agency' needs to be sustained in its original humanistic sense, one may develop concepts of 'collective responsibility' and 'collective action,' and consequently of 'collective or structural morality,' alongside their 'individual' counterparts.

Eventually, I developed two parallel accounts for structural ethics: a consequentialist account of structural ethics that morally evaluates networks regarding

their outcomes, and a deontological account of structural ethics that morally evaluates networks regarding their actions. On the micro level, contemporary approaches can serve to morally evaluate the human agent, regardless of her interactions in a network; on the macro level, the subject of moral evaluation is a whole network whose constituents are both humans and non-humans.

The ubiquity of emergent technologies, especially sophisticated artificial intelligent artifacts, requires us to upgrade our notion of ethics as a subject matter exclusively pertaining to human affairs. Non-humans do affect our actions and accordingly mediate moral outcomes. To do justice to both the human, as the sole source of actions, and to non-humans as mediators of morality, we need to develop a framework to investigate morality in a broader sense.

Notes

The article is supported by the GAJU specific research grant No. 107/2019/H.

1. For an overview and also evaluation of Verbeek's view, see Arzroomchilar and Novotný (2018).

2. For some critiques about the 'moral agency thesis' one may see, among others, these: Arzroomchilar and Novotny (2018); Philip Brey (2014); Illies and Meijers (2014); Peterson and Spahn (2011).

3. For a comprehensive survey, see his original paper (Brey 2014).

4. A phrase initiated by Bruno Latour (1992) in his famed paper "Where Are the Missing Masses? The Sociology of a Few Mundane Artifacts."

5. I take 'collective responsibility' to be sensible next to individual responsibility. I think one may conceive a network morally responsible in the sense of blameworthiness, yet this blameworthiness does not consist in nor is based on individual responsibilities. Even one may think of 'collective responsibility' despite neither factors subjecting to blame. Consider a network of individuals generating collectively a morally undesirable outcome while individuals have not been *free* enough to avoid the task. The head of an organization, for example, with a gun to her head by outsiders, might have put pressure on employees to fulfill their part without questioning the whole project. Here, no one seems free enough to evade fulfilling the task. No one then may be subject to a moral condemnation accordingly. Still, the whole structure may be understood as blameworthy.

Or, as another example, one may imagine a condition where an evil project is split into bits in such a way that each piece is going to be taken over by an individual. But, no one has any idea of the evil nature of the whole vision. Each factor is engaged with a bit of it, which does not seem unethical to her. Here, one may take the network to be morally responsible and deserving to be blamed, but no individual may be assigned to be practicing something evil due to their lack of knowledge.

6. According to Rawls, a 'suitably described society' or a 'well-ordered society' is a society in which (1) everyone willingly accepts and affirms the same principles of justice; (2) these principles are successfully realized in basic social institutions, including laws and conventions, and are generally complied with by citizens; and (3) reasonable persons are morally motivated to comply by their own sense of justice—they want to do what justice requires of them (Freeman 2019).

References

Arzroomchilar, Ehsan, and Daniel Novotný. 2018. "Verbeek on the Moral Agency of Artefacts." *Organon F* 25(4): 517–38.

Brey, Philip. 2014. "From Moral Agents to Moral Factors: The Structural Ethics Approach." In *The Moral Status of Technical Artifacts*, ed. Peter. Kroes and Peter-Paul Verbeek, 125–42. Dordrecht: Springer.

Davidson, Donald. 1980. *Essays on Actions and Events*. Oxford: Oxford University Press.

Floridi, Luciano, and Jeff W. Sanders. 2004. "On the Morality of Artificial Agents." *Minds and Machines* 14(3): 349–79. https://doi.org/10.1023/B:MIND.0000035461.63578.9d

Freeman, Samuel. 2019. "Original Position." In *The Stanford Encyclopedia of Philosophy*, ed. Edward N. Zalta. https://plato.stanford.edu/entries/original-position/.

Illies, Christian. F. R., and Anthonie Meijers. 2014. "Artifacts, Agency, and Action Schemes." In *The Moral Status of Technical Artefacts*, ed. Peter Kroes and Peter-Paul Verbeek, 159–84. Dordrecht: Springer.

Latour, Bruno. 1992. "Where Are the Missing Masses? The Sociology of a Few Mundane Artifacts." In *Shaping Technology/Building Society: Studies in Sociotechnical Change*, ed. Wiebe E. Bijker and J. John Law, 225–58. Cambridge, MA: MIT Press.

Latour, Bruno. (1991) 1993. *We Have Never Been Modern*, trans. Catherine Porter. Cambridge, MA: Harvard University Press.

Latour, Bruno. 2005. *Reassembling the Social: An Introduction to Actor-Network Theory*. Oxford: Oxford University Press.

Peterson, Martin, and Andreas Spahn. 2011. "Can Technological Artefacts Be Moral Agents?" *Science and Engineering Ethics* 17(3): 411–24. https://doi.org/10.1007/s11948-010-9241-3

Rawls, John. 1999. *A Theory of Justice*. Cambridge, MA: Harvard University Press.

Sayre-McCord, Geoffrey. 2012. "Metaethics." In *The Stanford Encyclopedia of Philosophy*, ed. Edward N. Zalta. https://plato.stanford.edu/archives/sum2014/entries/metaethics/.

Timmons, Mark. 2012. *Moral Theory: An Introduction*. Lanham, MD: Rowman and Littlefield Publishers.

Verbeek, Peter-Paul. 2005. *What Things Do: Philosophical Reflections on Technology, Agency and Design*. University Park: Pennsylvania State University Press.

Verbeek, Peter-Paul. 2011. *Moralizing Technology: Understanding and Designing the Morality of Things*. London: University of Chicago Press.