Martin Peterson. *The Ethics of Technology: A Geometric Analysis of Five Moral Principles*. Oxford University Press 2017. 264 pp. $74.00 USD (Hardcover ISBN 9780190652265).

Martin Peterson's *The Ethics of Technology: A Geometric Analysis of Five Moral Principles* offers a welcome contribution to the ethics of technology, understood by Peterson as a branch of applied ethics that attempts 'to identify the morally right courses of action when we develop, use, or modify technological artifacts' (3). He argues that problems within this field are best treated by the use of five domain-specific principles: the Cost-Benefit Principle, the Precautionary Principle, the Sustainability Principle, the Autonomy Principle, and the Fairness Principle. These principles are, in turn, to be understood and applied with reference to the *geometric method.* This method is perhaps the most interesting and novel part of Peterson's book, and I'll devote the bulk of my review to it.

Peterson's principle-based approach to the ethics of technology shares some significant similarities with Beauchamp and Childress's Principalist approach to bioethics, but also some notable differences (see Tom L. Beauchamp and James F. Childress, *Principles of Biomedical Ethics*, Oxford University Press 2012). Like Beauchamp and Childress, Peterson advocates the use of mid-level, domain-specific principles in applied ethics, as opposed to both theory-first and casuist approaches. So, Peterson argues that general purpose ethical theories often fail to determine the answer to individual cases, and that there is no satisfactory method for incorporating uncertainty or disagreement about ethical theory into our decision-making process (7-8). However, he also rejects the strong casuist claim that there are no general principles for evaluating or comparing cases, and argues that 'sharp and precise moral principles' are needed for applied ethics (13).

This demand for precision and clarity in applied ethics also helps motivate Peterson's argument for a 'geometric' construal of the principles. Unlike Beauchamp and Childress, who describe mid-level principles that hold only *prima facie,* and which must be weighed and balanced against one another, Peterson argues that the morally right course of action in a particular case should (ideally) be determined by exactly *one* mid-level moral principle, with the other principles being given no weight. The geometric method purports to do this by identifying each principle with one or more paradigm cases, each of which can be represented as a point in moral space. This space is, in turn, partitioned by a *Voronoi tessellation* which 'divides moral space into a number of regions such that each region consists of all points that are closer to a predetermined seed point (paradigm case) than to any other seed point' (15). Which moral principle should be used in a particular case is determined by which principle's paradigm case the case being considered is *most similar* to, with the other principles being given no weight (45-46). So, for example, a paradigm case for the Cost-Benefit principle concerns prioritizing design improvements in cars (88), while that for the Autonomy Principle concerns the 'Great Firewall of China' (158).

In chapter 2, Peterson offers a nuanced discussion of the mechanics of the geometric method, including the choice of a distance measure to depict the similarity between cases (30-36), and the appropriate way of interpreting the dimensions of the moral space associated with a geometric construal of the principles (37-40). I found several claims especially interesting and provocative. First, Peterson suggests that the paradigm case *x* for a given principle *p* can be identified in two ways: *ex-ante* or *ex-post.* In the former approach we 'have some (sufficiently good) direct reason for thinking that *x* is the most typical case to which *p* can be applied' (40). In the latter approach, we proceed by identifying a number of non-paradigm cases to which we are confident that the principle applies, and then calculate the 'center of gravity' of these cases (41). Second, Peterson suggests that some principles may have more than one paradigm case associated with them, either because the principle

itself can interpreted in more than one way, or because of uncertainty about the location of a paradigm case identified *ex-post.* The possibility of multiple paradigm cases for a single principle, when combined with Peterson's contention that all paradigm cases for a principle are equally important, in turn, allows for the existence of 'overlapping' regions of moral space where *multiple* principles apply. In these regions, Peterson argues for a 'gradualist' analysis, according to which actions are right to a certain degree, and wrong to a certain degree (52-55).

Chapter 3 describes the results of several studies in which engineering students and trained philosophers were asked to apply the geometric method. Peterson argues that empirical results of this type can (given certain assumptions) help determine the location of moral principles, and that they allow for the measurement of agents' *moral coherence* (59). While the studies are certainly suggestive, Peterson's interpretations of the results are at times a bit speculative. In particular, while his proposed interpretations of the dimensions of two-dimensional and three-dimensional moral space (as representing roughly freedom, uncertainty, and time) may be plausible, I would have liked to see more space devoted to arguing for this, given the importance of these dimensions in understanding the relationship between the moral principles.

After this detailed introduction of the geometric method, Peterson dedicates a chapter to clarifying each of his proposed principles. While there is a variety of interesting content here, the discussion is at times disconnected from both the details of the geometric method and from particular, applied debates within the ethics of technology. So, for example, Peterson spends a significant amount of chapter 4 considering how best to render the Cost-Benefit Principle compatible with deontic constraints. He argues that this is best done by the use of input or output *filters*, which serve to alter either the description of the case or the possible actions, as opposed to assigning numbers directly to things like rights violations. In chapter 5, he goes on to argue for a similar conception of the Precautionary Principle, and articulates a total of four versions of it: a deliberative version pertaining to permissible actions, and three epistemic versions pertaining to permissible beliefs. Chapters 6 and 7, on the Sustainability Principle and the Autonomy Principle respectively, are largely devoted to arguing against claims that natural resources and autonomy possess non-instrumental value. Finally, in chapter 8, on the Fairness Principle, Peterson argues for a *multidimensional* interpretation of the principle, according to which the relevant meaning of 'fair' depends on a particular case's location in moral space (175).

In the conclusion to the book, Peterson argues that we ought to prefer 'the interpretation of each principle that best explains our considered intuition about the underlying paradigm case' (205). While this seems plausible enough, I wasn't always convinced that Peterson's proposed interpretations of his principles met this bar. For example, in chapter 6 he notes that only 41% of the professional philosophers in his study applied the Sustainability Principle to a climate change case he claims is an *ex ante* paradigm case for the principle, but suggests that this might be partially due to the fact that many of these philosophers were Americans, and that Americans are generally skeptical of climate change (140). Later in this same chapter, his argument against assigning non-instrumental value to natural resources depends crucially on his judgement that any moral theorist who attaches any importance to motives would agree that it would be wrong for the 'Last Man' to detonate a nuclear weapon in empty space simply because he wanted to (154). In both of the cases, I find it unlikely that professional philosophers and moral theorists do or would make the sorts of judgements Peterson assigns to them. This is, however, an empirical question that might well be resolved by future work.

More generally, I wasn't convinced that Peterson's treatment of the individual principles actually delivered on the sort of increased clarity and precision that the early chapters had promised,

especially when compared to something like a Principalist approach. Instead, when confronted with actual cases, the need to weigh and balance competing moral factors returned under a variety of new guises: the specification of input and output filters, disagreements about which paradigm case was most similar, shifts in meaning of terms appearing in the principle, and so on. Whatever the merits of the geometric method and the five principles might be in helping to frame certain debates in applied ethics, they (at least so far) don't seem amenable to the sort of algorithmic application Peterson suggests as a long-term possibility (207).

These small complaints aside, *The Ethics of Technology* is highly worthwhile, especially for readers interested in methodology in applied ethics. Its combination of empirical research, careful attention to current debates, and selective, well-motivated use of formal tools is a model for future work in the field. And while Peterson's approach is unapologetically 'analytic,' he even devotes a chapter to arguing against the 'artifact approach' to technology associated with figures like Latour and Heidegger, according to which technological artifacts 'have moral properties of their own' (185). I'll be interested to see whether and how Peterson develops and expands the geometric method in future work and, in particular, the extent to which it can be used to make progress on the sorts of outstanding debates in applied ethics that are missing from this book.

Brendan Shea, Rochester Community and Technical College