

## Technology and the Virtues: a Response to My Critics

Shannon Vallor<sup>1</sup> 

Received: 2 August 2017 / Accepted: 27 September 2017 / Published online: 7 October 2017  
© Springer Science+Business Media B.V. 2017

The thoughtful and highly constructive commentaries from Don Howard, Emily McRae and Howard Curzer on my book, *Technology and the Virtues: A Philosophical Guide to a Future Worth Wanting*, are a pleasure to grapple with. I sincerely thank Diane P. Michelfelder for organizing their publication, and for proposing and moderating the special author-meets-critics session at the American Philosophical Association's Central Division meeting in Kansas City, Missouri in March 2017, at which these commentaries were originally given. I also wish to thank all three contributors for their careful critical insights, which I will examine in turn and then offer some overarching thoughts and questions in reply.

I will begin with Emily McRae's remarks because they have a very specific focus from which the broader critical discussion can expand, namely the role that Buddhist virtue ethics plays in what she sees as a primarily neo-Aristotelian philosophical project. I will then offer a response to Howard Curzer's remarks which are somewhat broader-ranging in their critical scope, and then finally to Don Howard's remarks, which propose a road forward from *Technology and the Virtues* rooted in a deeper exploration of the civic virtues, a proposal that I enthusiastically embrace.

McRae's critique centers on a pointed question: is the Buddhist ethical perspective in *Technology and the Virtues* more than just an ornamental feature, and possibly an ill-matched one, on what is ultimately a robustly neo-Aristotelian project? The question is fair and deserves a detailed response. I will also comment briefly on her highly constructive and welcome suggestion regarding the Buddhist practices of *bodhicitta* that might anchor Buddhist ethical thought more securely to the aims of my book.

I will begin with a response to McRae's most sweeping critical concern, and then move on to more technical objections. McRae asks whether Buddhist ethics is truly necessary to my project and whether it is an appropriate comparison with Aristotelian ethics, given that the latter is oriented to the contributions of a single ancient Greek

---

✉ Shannon Vallor  
svallor@scu.edu

<sup>1</sup> William J. Rewak, S.J. Professor, Department of Philosophy, Santa Clara University, 500 El Camino Real, Santa Clara, CA 95053, USA

thinker, and the former to a vastly richer philosophical lattice of contributions spanning continents and millennia. The critique is well-motivated; however, I would suggest that the force of the criticism depends in part on the purpose and scope of the comparison that is its target. If I am doing comparative philosophy for purely historical and theoretical purposes, then the criticism lands with far greater force, I think. For then, we really are comparing apples and oranges, or rather, as McRae might claim, an apple and an orange grove.

However, given the more limited and practical aims of my comparison, the critique may not be so devastating. First, contemporary neo-Aristotelian virtue ethics has diverged considerably from the views of Aristotle himself in a number of ways, for example, by rejecting his views on women and slavery, by calling into question his insistence on the reciprocity of the virtues, or by setting aside his teleological premises concerning a natural human function (*ergon*). Thus one finds quite a rich plurality of theoretical claims and commitments among virtue ethicists today, and in the past, who have counted themselves influenced by the Aristotelian model. It is not quite true, then, that a neo-Aristotelian approach is confined to the thought of a single thinker, Aristotle. While Aristotelianism has no doubt remained more conceptually unified than has Buddhism, neo-Aristotelian thinkers, too, span continents, and millennia. *Technology and the Virtues* employs not just the views of Aristotle but also those of Aquinas, McIntyre, Nussbaum, McDowell, and others who do not merely reproduce nor slavishly defend the classical Aristotelian position on ethics.

That said, I do not wish to dispute McRae's claim that the Buddhist ethical tradition, culturally and historically, is of a different magnitude of diversity and richness than the Aristotelian one. So that being said, what can be the value of a comparison such as I make in *Technology and the Virtues*? Is it necessary to my project, and if so, why? Is my book's treatment of Buddhist ethics (and perhaps likewise my treatment of Confucian virtue) indeed merely 'ornamental,' or does it have a legitimate purpose?

My response is that the inclusion of Buddhist ethics (along with Confucian ethics) is entirely *essential* to the project that *Technology and the Virtues* proposes, as I argue throughout Chapter 2: "The Case for a Global Technomoral Virtue Ethic." First, the inclusion serves a remedial function. The scholarly literature on virtue in the English-speaking world, and especially scholarship in *applied* virtue ethics, strongly identifies the provenance of virtue ethics with ancient Greek philosophy.<sup>1</sup> This risks creating the false impression that the Greeks, and especially Aristotle, were the original or exclusive progenitors of this way of thinking about the nature of ethics. Only in a minority of articles on applied virtue ethics are virtue traditions *not* indebted to the Aristotelian influence acknowledged or explored. One chief function of including treatments of

<sup>1</sup> For example, in *Moral Machines: Teaching Robots Right from Wrong*, a foundational and highly influential work in the rapidly growing field of applied machine ethics, it is proposed that we look further than top-down, rule-based ethical theories to "a very different conception of morality that can be traced to Aristotle, namely, virtue ethics" (Wallach and Allen 2009, 10, emphasis added). My own early articles in applied technology ethics, through their omission of the broader context offered in *Technology and the Virtues*, unwittingly reinforce the perception that virtue ethics just *is* an Aristotelian view (see Vallor 2010, 2012). Such strong identifications of virtue ethics with Greek, and especially, Aristotelian provenance are also reinforced by many commonly searched online reference sources such as the *Internet Encyclopedia of Philosophy*, which states in the opening of its entry on Virtue Ethics (<http://www.iep.utm.edu/virtue/>) that "Most virtue ethics theories take their inspiration from Aristotle...." and Wikipedia, which asserts in its entry on Virtue Ethics ([https://en.wikipedia.org/wiki/Virtue\\_ethics](https://en.wikipedia.org/wiki/Virtue_ethics)) that "the theory of virtue ethics was born with Plato and Aristotle."

Confucian and Buddhist notions of virtue in my book is to highlight the culturally myopic assumptions behind such omissions. Aristotle, and (neo)-Aristotelians, may well offer the most *familiar* model of virtue ethics to many scholars in the West, but there is no supportable historical claim that his is the *only* workable model of virtue, or that all future conceptions of virtue ought to derive from it.

Second, given that *Technology and the Virtues* begins by proposing to explore a subset of those cultural traditions that might inform and enrich a contemporary technomoral virtue ethic that can be applied to problems and concerns of *global* scope, to have *not* addressed Buddhist or Confucian moral philosophies and practices, or any other comparably robust virtue traditions, would have been as unsound as an architect arbitrarily restricting a review of contemporary global design possibilities to columns, lintels, and friezes, and their later modulations and influences.

Third, the exploration of Buddhist and Confucian resources for thinking about virtue in a contemporary context is motivated by the need to recognize the impossibility of constructing a contemporary global virtue ethic that is entirely cohesive and without irreconcilable conflicts and tensions. Conceptions of virtue and human flourishing are never universal. There have always been, and will always be, coherent accounts of the good life that cannot be reduced to or fully reconciled with others. All virtue traditions, including the Aristotelian one, are culturally inflected, historically contingent, and contested both within and externally to their local instantiations in particular communities. In contrast to premodern eras, however, when many societies could limit their interactions with other cultures to expedient but limited forms of trade and information exchange and thus maintain *relatively* insular and stable moral cultures, today humanity's increasing global interdependence and shared technoscientific power mean that distinct and frequently irreconcilable conceptions of the good life *must* find new ways to enter into fruitful conversation and projects of cooperative deliberation if we are to flourish together on this planet. Acknowledging the existence of traditions of virtue just as robust and enduring as the Aristotelian one, or even more so, is an essential step toward grounding those challenging but necessary conversations and collaborations.

Finally, by bringing Buddhist ethics into view, my book is better able to highlight its central concept of moral cultivation, and the cross-cultural force of this ideal. This is the unifying thread of Part II of *Technology and the Virtues*. As McRae notes, the concept of moral cultivation is a core element of the Buddhist tradition, and while its focus and aims diverge significantly in the Theravada and Mahayana schools, for example, the notion remains fundamental. It is precisely the vast historical scope and robustness of the Buddhist tradition that led me to highlight it in its totality as a compelling evidence of the profound and enduring force of moral cultivation as a practical ideal. While I take pains in the book to point out that specific conceptions of the good life and the virtues themselves are irreducibly diverse and historically and culturally fluid, even *within* particular traditions, there is a robust resonance between otherwise irreconcilable virtue traditions insofar as they emphasize the malleability of the human person through effective moral *practices*.

This is to be embraced if we take the view I defend in the book, namely that at the level of moral concreteness, there is unlikely to be one single best way to live, even from a single historical and cultural standpoint. The comparisons upon which the coherence of my book rests most heavily, then, are at a higher level of abstraction with

respect to the importance of undertaking general moral practices of self-cultivation, as opposed to practices of moral rule-giving or consequentialist analysis.

At the level of moral concreteness, that is, of determining how we ought to live and how we ought to cultivate our virtues as we stand within a particular human context or lifeworld, I argue in *Technology and the Virtues* that neither Aristotle, nor the Confucians, nor the ideals of the countless thinkers who have articulated the teachings of the Buddha, can be our immediate guides to living well with new technologies. Instead, I argue that we need to develop some *new* guidance on the foundations of virtue practices, guidance that responds to the unprecedented challenges ahead of us, many of which virtue ethicists of earlier eras could have scarcely conceived.

That said, I *do* draw upon both Buddhist and Confucian traditions as sources of additional moral imagination about the virtues and the forms they might take in our future, as a corrective to some of the narrow assumptions about the good life that remains embedded in Eurocentric forms of moral philosophy that still bear deep traces of the classical Greek inheritance. In particular, Aristotle's relative disinterest in the ethics of family life receives an important corrective from Confucian moral thought, just as Buddhist thought offers to other virtue traditions a much-needed centering of the importance of a properly scaled moral perspective that takes a view of the moral whole, as well as the importance of moral attention, especially to suffering and vulnerability—elements of moral life to which Aristotle is damnably inattentive. I argue in *Technology and the Virtues* that both of these elements of moral life are of increasing importance in a contemporary technosocial context, where the vulnerability, interdependence, and fragility of life—including that of non-humans—are rapidly increasing due to technosocial change, and becoming ever more important to attend to in our technomoral practices.

My aim, then, in drawing upon the Buddhist tradition was to begin to decenter the Aristotelian perspective from which an understanding of virtue ethics so often departs among contemporary moral philosophers in the West, and which too often constrains our moral imagination of the ways in which humans and others might flourish in the future. In that sense, to answer McRae's question, my exploration of the Buddhist perspective *is* most certainly necessary. As a moral philosopher originally educated within the European tradition, my understanding of Buddhist thought is inevitably constrained and impoverished by that point of departure. Yet from a hermeneutic standpoint, this limited perspective does not foreclose the opportunity to engage in meaningful dialog with other philosophical traditions; indeed, it opens it.

Next, let me address two important technical criticisms. The first is the question of whether the Buddhist tradition is truly exemplarist in the same sense as Aristotle or Confucian ethics. I do not think that it is exemplarist in the same *way*. In contrast to the Aristotelian and Confucian models, in Buddhism, there is no one concrete model that all moral aspirants can look to from any stage of moral development. As McRae rightly notes, and as I try to acknowledge in *Technology and the Virtues*, the Buddhist tradition takes seriously, and to its credit, the notion that in her words, "what is best, morally, is a moving target, depending on what is psychologically possible for you at that moment." Not only do I think this is broadly right, it perhaps provides a better, richer form of exemplarism than the other two traditions I consider. For I still do think that moral exemplars play a key role in Buddhism.

As McRae herself notes, many of the stories told in the literatures of Buddhism are intended to stretch our conventional assumptions about what it is possible for a moral being to do. The message is *not* that the ordinary Buddhist aspirant should directly imitate the behavior in the story. Rather, through narrative projection into the stories and their examples, one may gradually come to conceive and discern a wider field of moral possibilities for noble action than before. Thus, the stories in which the exemplar's behavior seems particularly radical or unfathomable, for example, the story of the hungry tigress, or a similar story I reference in the book in which in a former life, the Buddha gives away his own wife and children to wandering brahmins (Vallor 2016, 115)—are not action scripts to copy. They are examples of a different sort—examples that serve the purpose of breaking through our often narrow and self-serving conventions about what is possible in the domains of duty, compassion, and love.

In this way such examples serve a purpose not unrelated to Aristotle's advice that to find the appropriate path for ourselves, we must often pull our thoughts and feelings toward the one extreme that runs directly counter to our natural inclinations, for it is through confronting such seemingly alien possibilities that the more modest path to our own moral cultivation becomes easier to envision (Aristotle 2011, 1109b1–5.) Moreover, as I note in my discussion of the role of the monastic *Sangha* community as exemplary, in Buddhist practice, what is being modeled for the aspirant is often not the outward *behavior* of the exemplar, since a layperson is unlikely to take up the monastic way of life. Rather, what is often being modeled is proper *intention* and *discipline*. These are just as essential to the cultivation of virtue as proper ritual and habit, and regardless of the particular social role or stage of spiritual development, one occupies the *Sangha's* dedication to monastic life, and practice provides a particularly vivid example of “extended commitment to the practice” of moral cultivation and the uncommon forms of life that may be driven by a deepened spiritual insight (Vallor 2016, 73). Here, the *Sangha* functions as a body of moral exemplars *collectively*—not as individual models or moral “experts” in the way each *phronimos* can be assumed to be, but as a group who represents the moral transformation and expansion that uncommon spiritual discipline, right intention, and right belief make possible.

Next, there is the question of whether the notion of “character” has any place in a Buddhist account which, through the doctrine of *anātman*, denies the substantial and enduring reality of a personal self. I explicitly anticipate and address this criticism on page 88 of *Technology and the Virtues*, where I point out that the Buddhist tradition does not deny the robust, and more or less stable, *appearance* of a personal self that during one's lifetime displays characteristics of continuity, such as memory and traits of character. Nothing essential to practical ethics is lost if these features of the apparent self fail to persist beyond their worldly appearances. For the purposes of practical ethics to which *Technology and the Virtues* is oriented are *purely* worldly, concerned only with our flourishing on this planet in the not-too-distant future, and can accommodate or be divorced from any notion of the survival of a personal self after death. Indeed, the Buddhist goal of spiritual liberation, which McRae argues that I conflate with self-cultivation, is entirely distinct from but in no way incompatible with the latter aim. It is entirely consistent to strive throughout one's life to cultivate one's apparent, worldly personality in ways that make one increasingly less invested in the idea of its substantiality and permanence.

Moreover, the framework of technomoral virtues I am attempting to construct is, quite clearly, not *itself* a Buddhist ethic. Thus McRae's worry that the required loss of *belief* in a substantial self will destabilize any commitment to technomoral self-cultivation seems unwarranted. All that is required to take the concept of virtuous character seriously within ordinary moral practice (as opposed to the detached theoretical standpoint of a moral philosopher) is a belief that committed practice can make a noticeable difference to the sort of person you are, gradually improving your repertoire of moral responses to the world and those with whom you share it. I know few practicing Buddhists who would resist that claim on a mundane level, but even if Buddhism *were* incompatible with that claim, it would not follow that a contemporary technomoral virtue ethic can learn nothing of value from Buddhist virtue ethics.

Finally, let me take up McRae's helpful suggestion to adopt the Mahayana concept of *bodhicitta* as a way of better grounding my accounts of the virtues of care and moral attention, as well as empathy, perspective, and others. This is welcome advice. This is a concept I do reference briefly and indirectly in *Technology and the Virtues* with respect to the Buddhist mental practices of perspective-taking that I describe on page 116, where I point out that such practices help to "recruit moral attention (hence the emphasis on 'viewing' from the 'eyes' of another being,)" while broadening the scope of our ordinary empathic practices with the help of moral imagination. I go on to suggest on pages 117, 150, and 173 that *new* forms of such practices can and should be developed within the technosocial context. Such practices can cultivate the ability to attend to phenomena like the fragility of the planetary climate or oceanic ecosystems, systems that we did not evolve to pay close *moral* attention to, since until very recently, we lacked the technological powers to knowingly imperil or care for systems at that scale. I welcome McRae's helpful suggestion to go beyond my brief treatment of these practices and to more fully develop the rich resources in the Mahayana Buddhist tradition for such a project.

Let me turn now, then, to Curzer's critical response, which like McRae's, begins from the perspective of comparative philosophy. Curzer points out that the doctrine of the mean in Aristotle differs considerably from the notion of the mean in Confucian thought or from the Middle Way in Buddhism. I think that his criticisms here are well founded, though perhaps I might be able to dilute their force with a few short remarks. First, I will note that even Aristotle's doctrine of the mean is murkier than one might assume. Aristotle (2011) says in Book II, Chapter 6 of the *Nicomachean Ethics* that a virtue is a mean between extremes of deficiency and excess, but then goes on to point out that this can seem problematic with virtues such as temperance, since we cannot be "too temperate" (or too just, or too wise). He goes on to reconstruct with some difficulty the two extremes opposed to each virtue, but notes that some lack easily identifiable names, and that some virtues lie nowhere near the "middle" but remarkably close to one of their associated extremes.

Aristotle does describe the mean as directly concerned with taking appropriate measures of pleasures and pains, as I note on p. 68 of *Technology and the Virtues*. So, the link between Aristotle and the Buddhist concern for moderation in the realm of pleasures and pains is not wholly untenable, even if their motivations are different. Still, the doctrine of the mean functions in Aristotle as a strong heuristic rather than a fixed metaphysical principle—the virtuous mean is a useful way of framing a common pattern of moral life, but one that resists universal or precise specification by reasoning



(Aristotle 2011, Book II, Chapter 9). The fact then, that Confucians and Buddhists adopt other heuristics that operate rather differently from Aristotle's in their scope or application is perhaps not much of a problem for my purposes, which are to show that virtue traditions in general do tend to concern themselves with certain norms and heuristics of "appropriate measure" and "balance" in moral life.

As I state in Chapters 2 and 3, the cross-cultural comparisons I draw in *Technology and the Virtues* do not aim to find *points of contact* between different virtue traditions so much as conceptual *resonances* (Vallor 2016, 64); that is, concepts which are not equivalent but which may nonetheless have sufficient conceptual *proximity* to foster meaningful discourse across virtue traditions. I take pains throughout *Technology and the Virtues* to note that virtually all *direct* comparisons between culturally distinct virtue concepts and ideals break down at some level, or fail outright. It does not follow that we learn nothing by analyzing the *relative* conceptual proximity and resonance between irreconcilable notions. If that conclusion did follow, the art of communication between speakers of different native languages would be largely impossible to explain.

With respect to the notion of moral extension in Aristotle, which Curzer finds absent in Aristotle, I do note in the book that "while the extension of moral concern is a central and explicit theme in both Confucian and Buddhist thought, its role in Aristotle's ethics is more obscure" (Vallor 2016, 111). I go on to offer a speculative reconstruction of Aristotle's notion of complete friendship that, if tenable, might show that Aristotle at least *implicitly* held a notion of the appropriate extension of moral concern. But if that speculative interpretation is ultimately untenable, and Aristotle held no such notion implicitly or explicitly, then I say, so much the worse for Aristotle. The project of *Technology and the Virtues* is presented as an entirely pragmatic and forward-looking one, not wedded to any particular classical view of the good life. On page 35, I note that we have good reason to think that each of the classical perspectives from which I draw include more than few commitments that are "simply wrong," factually and morally, as well as other blindspots and biases that ought to be subjected to critical pressure and rejected.

The second part of Curzer's critique takes a different theme; there he challenges my claim that the virtues I describe are *technomoral* virtues, that is, particularly necessary for flourishing in the world shaped by modern technology. Here, he and I simply have some differences of opinion, although part of the dispute can perhaps be resolved by pointing out two senses in which a virtue can be "particularly essential" for flourishing in technosocial contexts. Curzer takes this to mean that the virtue must be *more* essential than it was in the past. However, it need not mean that—it can just mean that among all the traits of character that we need at the present moment, to face the present and future challenges, these stand out more than some others. They may have *also* been essential in previous eras, to meet challenges to human flourishing that presented themselves *then*. That is, they need not be objectively *more essential* today than in the past, to be particularly essential today among all the positive traits available to the human personality.

Still, I think some of the virtues I describe really *are* more essential than ever before. For reasons of space, I will not argue for them all, but considering just empathy and perspective, I think the fact that emerging technologies create and magnify our power to destroy the viability of life on the planet for future generations, and the fact that technoscientific practices originating in one country can severely impact the availability

of clean air, water, energy, or lifesaving vaccines in another, make evident an even stronger need for a holistic moral perspective and for more expansive capacities for empathy than were needed by humans living in the classical age.

I also challenge some of Curzer's other claims. I think that some forms of effective communication *are* harder to manage today, not easier, in part, because of new technologies that we have yet to cultivate good moral habits and skills around. I think that mobile devices which provide portable porn, games, and news on demand, not to mention the ability to access your workload from anywhere at anytime, *do* create more acute challenges for self-control than the technologies they replaced. I flatly reject Curzer's claim that "the more technology takes over the tasks of caring, the less necessary it is for people to possess the virtue of care." As I argue extensively in Chapter 9 of *Technology and the Virtues*, the virtue of care is a fundamental moral skill without which we cannot hope to flourish with others, and we can only hope for automation to allow us to care *better* and more readily, rather than to care for one another less often or less skillfully. Curzer later asserts (somewhat in tension with the quote above) that carebots are likely to enhance rather than reduce human caring. I state in Chapter 9 that this is a real and welcome possibility. But we cannot simply *assume* that will be the case, for there are other plausible scenarios that run the other way, and it will be the technomoral virtues of the designers, users, and regulators of carebots that will determine which scenarios obtain. With respect to military robotics, Curzer claims that remote drones bring us closer to killing than missiles, but does not mention the growing pressure to develop drones and other potentially armed autonomous weapons that do not require human remote operators, and that can operate without real-time human guidance and oversight of targeting and firing decisions.

In general, Curzer adopts the view that technology's benefits consistently outweigh its associated risks and costs. Though *Technology and the Virtues* does resist that generalized form of techno-optimism, Curzer distorts the tone of the book by describing my view of technology as "rather pessimistic." Optimism and pessimism are not the only options on the table, and certainly they are not the best options. Curzer seems to miss the explicitly constructive function of my book, which is to develop a practical framework for building *better technologies* that promote human and planetary flourishing. If I were indeed *pessimistic* about technology's prospects for enabling human flourishing, I would not have bothered to write *Technology and the Virtues*, or I would have written a very different book that counseled us to eschew technological advancements. Instead, at every turn in the book I counsel precisely the *opposite* strategy—the active embrace and skillful, thoughtful engineering of better, more virtuous technosocial designs and practices. As I note in my opening chapter,

...however widely we share this part of ourselves with other creatures, humanity without technology is not a desirable proposition—it is not even a meaningful one. The only meaningful questions are: which technologies shall we create, with what knowledge and designs, affording what, shared with whom, for whose benefit, and to what greater ends? These are the larger questions driving this book. Yet humans lacking the technomoral habits and virtues described within its pages could, I think, never hope to answer them. Let us not surrender that hope. (Vallor 2016, 13).



The constructive pro-technology aim of *Technology and the Virtues* is also embedded in each of the more narrowly focused and applied chapters of Part III. There I do examine ethically problematic designs and uses of specific classes of technologies, from digital media and surveillance to robotics and biomedical enhancement, but in each case, I also discuss existing and potential applications of the same technologies that can promote human flourishing, if designed and deployed under the right conditions and with the necessary skill and moral wisdom.

While Curzer seems to read into my work, a form of techno-pessimism that I repeatedly and explicitly eschew, he and I do agree that “technology demands new versions of the old virtues,” as I state in the book (Vallor 2016, 119). The technomoral virtues are not intended to be *new* character traits, but rather pre-existing human potentials of character that have been specifically *adapted* to the new challenges increasingly presented by emerging technologies (Vallor 2016, 10, 32, 50). So Curzer and I share the view that the challenge of cultivating the technomoral virtues requires figuring out how to take older habits of moral life that worked for virtuous people (for example, staying properly informed about morally important matters of civic and global life by simply turning on any of the nightly news channels), and figuring out how to adapt or replace them with new moral habits better suited to the new affordances of technology and its rapidly changing social expressions.

Curzer also notes helpfully that there are certain features of technosocial life that threaten to destabilize the very notion of a virtuous life that is reliably guided toward human flourishing. One is the risk of humans using genetic, biomechanical, or other forms of enhancement technology in ways that destabilize our own biological nature, such that the basic moral capacities and psychology taken for granted by virtue traditions can no longer be assumed. I confront this prospect in Chapter 10, and it remains a very real concern in an era of resurgent interest in gene editing with CRISPR and other powerful new techniques. Another equally pressing concern is the risk that if the technosocial contexts and structures of contemporary life continue to change too rapidly, we may be challenged to form *any* stable moral habits or rituals that remain well-adapted to our needs. This is a legitimate concern. It means, at a minimum, that habits of skillful improvisation and experimentation in moral life are going to be increasingly important to cultivate, and that our rituals will need to be even more dynamic and responsive. Fortunately, we can learn from the ways in which the Confucians emphasized the esthetic element of moral ritual which enables it to remain responsive and genuine, rather than rigid and rote.

Perhaps there is, however, an upper limit to how fluid and responsive moral habits can become and still be habits at all. In that case, perhaps there is an upper limit to how rapidly and often our social lives and institutions ought to be radically disrupted by new technologies. There is no law of nature or of economics that says that we *must* use technology to compulsively and indiscriminately destabilize our material and digital infrastructures, rather than to skillfully maintain those designs and artifacts that still do work well, while selectively refining or replacing those that do not. So I agree that we must be aware of the circumstances under which technosocial innovation interferes with *any* effective commitment to moral self-cultivation and human flourishing. But the pressure must be put on the mechanisms of innovation to leave room for the latter, not to sacrifice human flourishing on the altar of mindless technical compulsion.

Now, then, let me turn to Don Howard's remarks, which convince me that he has at some point loaded spyware onto my computer. For his proposal that I turn my project explicitly toward the civic virtues as means of developing the distinctively collective and distributed features of technomoral wisdom is one that echoes a sketch for another book that I have already begun to work on, although it is only in the very earliest stages of development.

Let me first, however, address Howard's two critical points. The first is a question about why I have chosen only Aristotelian, Buddhist, and Confucian traditions to inform my thinking about a global community of technomoral virtue. These three traditions were chosen not because they *exhaust* the global cultural resources we possess for thinking about virtue, but rather because an exhaustive account of those resources is impossible in a single book, and because my pragmatic aim was simply to deliver a "proof of concept" that virtue ethics offers a sufficiently broad conceptual footing for a new global ethic that can resonate across many different cultural traditions. I do think that the three traditions I chose share especially strong resonances with respect to the practice of moral self-cultivation. Islamic, Hindu, Orthodox Christian, Judaic, native American, African, and other moral traditions I might have mentioned have their own languages of moral development, some of which are easier to bring into resonance with virtue ethics than others. My book could have articulated this more clearly.

Howard's second critical point suggests that I might have focused more on the possibility that through a global discourse about human flourishing we might not merely come into meaningful conversation with other moral communities but construct, in his words, "new communities of belief and practice" and "new common ground" upon which to secure the good life. This is, in fact, precisely the goal I was hoping to promote with *Technology and the Virtues*. This is the focus of my discussion in Chapter 2 of the new global technomoral practices emerging that seem to seek distinctive goods internal to those practices: namely global community, intercultural understanding, global security, and global justice (46–47). I revisited this notion in my discussions in Chapter 6 when I raised the prospect of cultivating new norms of "global public character" (147–148), through the joint cultivation of virtues of technomoral justice, technomoral civility, and technomoral flexibility that enable collective deliberation and wise decision-making among a global network of actors.

I did suggest, in the spirit of a pluralism that I remain loathe to give up, that we might conceive of global public character as a body of norms and skills best suited to collective action in matters that directly impact broad cross-sections of the human family and systems on a planetary scale, while leaving more or less intact those local and regional conceptions of the good life that do not themselves impede their members' simultaneous pursuit of global civic character or participation in global civic life. We might call this the "Star Trek" ethos in which a unified civic body, the Federation, is made up of members who have a robust set of shared civic virtues cultivated within the practices of the interstellar Federation, but who preserve also, in some tension with these global practices, their own cultural norms and virtues. But perhaps Don Howard would prefer an even more universalist aim? I am not sure.

Howard's constructive project, to develop a more thorough and explicit analysis of the collective and distributed civic virtues that are needed to constitute wise communities of technomoral actors, is in fact in line with one of two projects I have already

envisioned as a follow up to this book. I jokingly attributed this to spyware, but in fact his mind-reading more likely reflects the fact that this is *Technology and the Virtues'* only novel and constructive proposal, yet one that is only sketched in the barest of outlines. Howard is absolutely correct: to merely cultivate ourselves as individuals who can flourish in local communities and can act as wise consumers and users of new technologies is not going to suffice to meet the challenges of preserving *together* the many technosocial systems of finance, transportation, defense, research, food production, and communication that already operate on intercontinental and even planetary scales—much less the challenge of building new and improved systems on those scales.

I therefore embrace Howard's suggestion to move forward with developing a fuller account of how distributed civic virtues and communities of practice might be cultivated on such scales, and how newly distributed civic virtues might be expressed. In particular, I welcome the opportunity to think through the civic virtues beyond those virtues of strong civic import that are already articulated in my book, such as honesty, justice, empathy, care, civility, flexibility, perspective, and magnanimity. Howard's suggestion of "Socratism" as an unarticulated virtue is an excellent example to follow. Another might include the virtue of mediation, as a skillful disposition to mend, mitigate, or bridge acute and seemingly irresolvable conflicts or gaps in communication that arise between third parties. Still another task might include a critical re-examination and adaptation of Walzer's civic virtue of loyalty (1974), which becomes newly problematic on a view of global civic character that includes fundamental duties to all, including future generations of humans or other forms of life, to whom I have no direct relations.

There are other daunting challenges with this project that I am still working out in my own thinking. Most prominent among them is the question of how far we can go with a notion of distributed virtue or character. When Howard asks, in reference to the HRI community, "what virtues distinguish [that community] qua community and a large proportion of its membership?", it opens up the question of whether the virtues of a community can be virtues *in the same sense* as the virtues that characterize its individual members, and how the precise relationship between those levels of virtue might be properly characterized.

Is a company or government agency that consistently demonstrates moral courage possible, and if so, will its virtue of moral courage consist only in the joint effects of its morally courageous employees and executives? Or is there a virtue of moral courage that attaches to the organization itself? Can an organization really have a *character*, since in most virtue traditions character requires not only a reliable disposition to appropriate action but also reliably appropriate internal states of belief, feeling, and motivation? Can institutions or other collectivities genuinely think, feel, and desire? Can organizations without a common mind have moral *integrity*? Can an organization as a *body* take up each of the fundamental practices of moral self-cultivation that I outline in *Technology and the Virtues*, and if so, how? Are there any special practices of moral self-cultivation that belong *strictly* to organizations?

The extent to which organizations can and should be personalized as virtuous or vicious gives rise to philosophical questions of some complexity and difficulty, and I am unsure how to resolve them. I will continue to rely on the civic virtues of generosity and Socratism, as demonstrated by my philosophical colleagues in this issue, to help me find my way.

## References

- Aristotle. (2011). *Nicomachean Ethics*, trans. Robert C. Bartlett and Susan D. Collins. Chicago: University of Chicago Press.
- Vallor, S. (2010). Social networking technology and the virtues. *Ethics and Information Technology*, 12(2), 157–170.
- Vallor, S. (2012). Flourishing on Facebook: virtue friendship and new social media. *Ethics and Information Technology*, 14(3), 185–199.
- Vallor, S. (2016). *Technology and the virtues: a philosophical guide to a future worth wanting*. New York: Oxford University Press.
- Wallach, W., & Allen, C. (2009). *Moral machines: teaching robots right from wrong*. New York: Oxford University Press.
- Walzer, M. (1974). Civility and civic virtue in contemporary America. *Social Research*, 41, 593–611.