Philosophy   College of Arts & Sciences

9-19-2016

Introduction: Envisioning the good life In the 21st century and beyond

Shannon Vallor
Santa Clara University, svallor@scu.edu

Introduction

ENVISIONING THE GOOD LIFE
IN THE 21ST CENTURY AND BEYOND

IN MAY 2014, cosmologist Stephen Hawking, computer scientist Stuart Russell, and physicists Max Tegmark and Frank Wilczek published an open letter in the UK news outlet *The Independent*, sounding the alarm about the grave risks to humanity posed by emerging technologies of artificial intelligence. They invited readers to imagine these technologies "outsmarting financial markets, out-inventing human researchers, out-manipulating human leaders, and developing weapons we cannot even understand."[1] The authors note that while the successful creation of artificial intelligence (AI) has the potential to bring "huge benefits" to our world, and would undoubtedly be "the biggest event in human history . . . it might also be the last." Hawking echoed the warning later that year, telling the BBC that unrestricted AI development "could spell the end of the human race." While some AI enthusiasts dismiss such warnings as fearmongering hype, celebrated high-tech inventors Elon Musk, Steve Wozniak, Bill Gates, and thousands of AI and robotics researchers have joined the chorus of voices calling for wiser and more effective human oversight of these new technologies.[2]

How worried should we be? More importantly: what should we *do*?

AI is only one of many emerging technologies—from genome editing and 3D printing to a globally networked "Internet of Things"—shaping a future unparalleled in human history in its promise *and* its peril. Are we up to the challenge this future presents? If not, *how can we get there*? How can humans hope to live well in a world made increasingly more complex and unpredictable by emerging technologies? Though it will require the remainder of the book to fully respond to that question, in essence my answer is this: we need to cultivate in ourselves, collectively, a special kind of moral character, one that expresses what I will call the *technomoral virtues*.

What do I mean by *technomoral* virtue? To explain this concept will require introducing some ideas in moral philosophy, the study of ethics. At its most basic, ethics is about what the ancient Greek philosopher Socrates called the "good life": the kind of life that is most worthy of a human being, the kind of life worth choosing from among all the different ways we might live. While there are many kinds of lives worth choosing, most of us would agree that there are also some kinds of lives *not* worth choosing, since we have better alternatives. For example, a life filled mostly with willful ignorance, cruelty, fear, pain, selfishness, and hatred might still have some value, but it would not be a kind of life worth *choosing* for ourselves or our loved ones, since there are far happier choices available to us—better and more virtuous ways that one can live, for ourselves and everyone around us. But what does ethics or moral philosophy have to do with *technology*?

In reality, human social practices, including our moral practices, have always been intertwined with our technologies.[3] Technological practices—everything from agriculture and masonry to markets and writing—have shaped the social, political, economic, and educational histories of human beings. Today, we depend upon global systems of electronic communication, digital computation, transportation, mass manufacturing, banking, agricultural production, and health care so heavily that most of us barely notice the extent to which our daily lives are technologically conditioned. Yet even our earliest ancestors used technology, from handaxes and spears to hammers and needles, and their tools shaped how they dealt with one another—how they divided their labor, shared their resources and living spaces, and managed their conflicts. Among our primate cousins, female chimpanzees have been observed to stop fights among males through technological disarmament—repeatedly confiscating stones from an aggressor's hand.[4]

Ethics and technology are connected because technologies invite or *afford* specific patterns of thought, behavior, and valuing; they open up new possibilities for human action and foreclose or obscure others. For example, the invention of the bow and arrow afforded us the possibility of killing an animal from a safe distance—*or* doing the same to a human rival, a new affordance that changed the social and moral landscape. Today's technologies open their own new social and moral possibilities for action. Indeed, human technological activity has now begun to reshape the very planetary conditions that make life possible. Thus 21st century decisions about how to live well—that is, about *ethics*—are not simply moral choices. They are *technomoral* choices, for they depend on the evolving affordances of the technological systems that we rely upon to support and mediate our lives in ways and to degrees never before witnessed.

While ethics has always been embedded in technological contexts, humans have, until very recently, been the primary authors of their moral choices, and the consequences of those choices were usually restricted to impacts on individual

or local group welfare. Today, however, our aggregated moral choices in technological contexts routinely impact the well-being of people on the other side of the planet, a staggering number of other species, and whole generations not yet born. Meanwhile it is increasingly less clear how much of the future moral labor of our species will be performed by human individuals. Driverless cars are already being programmed to make 'ethical' driving decisions on our behalf while we relax and daydream, even as other cars roll out of the factory programmed to commit the *unethical* act of cheating on their innocent owners' emissions tests.[5] High-frequency trading algorithms now direct the global flow of vital goods and wealth at speeds and scales no human observer can follow. Artificially intelligent life coach apps are here to 'nudge' us when we need to lower our voices, call our mothers, or write nicer emails to our employees. Advanced algorithms inscrutable to human inspection increasingly do the work of labeling us as combatant or civilian, good loan risk or future deadbeat, likely or unlikely criminal, hireable or unhireable.

For these reasons, a contemporary theory of ethics—that is, a theory of what counts as a *good life* for human beings—must include an explicit conception of how to live well with technologies, especially those which are still emerging and have yet to become settled, seamlessly embedded features of the human environment. Robotics and artificial intelligence, new social media and communications technologies, digital surveillance, and biomedical enhancement technologies are among those emerging innovations that will radically change the kinds of lives from which humans are able to choose in the 21st century and beyond. How can we choose wisely from the apparently endless options that emerging technologies offer? The choices we make will shape the future for our children, our societies, our species, and others who share our planet, in ways never before possible. Are we prepared to choose *well*?

This question involves the future, but what it really asks about is our readiness to make choices in the *present*. The 21st century is entering its adolescence, a time of great excitement, confusion, and intense anxiety, an age both wildly hopeful and deeply troubled. As with many adolescents, our era is also deeply self-absorbed. In popular and scholarly media, we find both historical consciousness and the 'long view' of humanity giving way to an obsessive quest to define the distinctive identity of the present age, an identity almost always framed in technological terms. Whether we claim to be living in the 'Age of Information,' the 'Mobile Era,' the 'New Media Age,' or the 'Robot Age,' we seem to think that defining the technological essence of our era will allow us to better fathom the course of its future—*our* future.

Yet in one of those cruel paradoxes of adolescence, all our ruminations and fevered speculations about the mature shape of life in this century seem only to

make the picture *more* opaque and unsettled, like a stream bottom kicked up by shuffling feet. Among all the contingencies pondered by philosophers, scientists, novelists, and armchair futurists, the possibilities presented by emerging technology have proved to be the most enticing to the imagination—and the most difficult to successfully predict. Of course, early visions of a postindustrial technological society were strikingly prescient in many respects. Debates about today's emerging technologies echo many of the utopian and dystopian motifs of 20th century science fiction: fears *and* hopes of a 'brave new world' of bioengineered humans constructed by exquisite design rather than evolutionary chance; of humans working side-by-side with intelligent robotic caregivers, surgeons, and soldiers; of digitally-enabled 'Big Brothers' recording and analyzing our every act; and of the rise of a globally networked hive mind in the 'cloud' that radically transforms the nature of human communication, productivity, creativity, and sociality.

Still, we cannot help but smile wistfully at the lacunae of even our most far-seeing science fiction visionaries. In the classic Ray Bradbury tale 'The Veldt,' first published in 1950, we encounter the existential and moral dilemma of the Hadley family, whose complete surrender to the technological comforts of the 'Happy-life Home' has stripped their lives of labor, but also of joy, purpose, and filial love. In a present marked by the increasingly sophisticated design of 'smart homes,' Bradbury's story resonates still. It may have taken a few decades longer than he expected, but affluent modern families can now, just like the Hadleys, enjoy a home that anticipates their every personal preference for lighting, room temperature, music, and a perfectly brewed cup of coffee—and the 'smart homes' of the future will even more closely approximate Bradbury's vision. We also recognize all too well the Hadleys' parental anxiety and regret when their children, irretrievably spoiled by the virtual world of their interactive playroom, fly into an incandescent rage at the thought of having their electronic amusements removed.

Yet today we can only laugh or cry when Lydia, the children's mother, complains that her surrender to domestic technology has left her without "enough to do," and too much "time to think." No technologically-savvy 21st century parent can identify with Lydia Hadley's existential plight.[6] Rather, the promised land of unlimited technological leisure has given way to a reality of electronic overstimulation and hypersaturation, a 24-hour news cycle, and smartphones on which your boss texts you from the 18th hole in Dubai while you sit at the dinner table wolfing down take-out, supervising your child's Web research on whale sharks, feverishly trying to get caught up on your email, responding to your Facebook invitations, and updating the spreadsheet figures your colleagues need for their afternoon presentation in Seoul. Leisure is one thing our age does *not* afford most modern technology consumers, who struggle each night to ignore the

incoming status updates on their bedside devices so that they may grab a few precious hours of sleep before rejoining that electronic day that knows neither dusk nor dawn.

Indeed, the contemporary human situation is far more complicated, dynamic, and unstable than any of the worlds depicted in our first imaginings of a high-tech future. Today, exponential leaps in technological prowess and productivity are coeval with widespread economic stagnation, terrestrial resource depletion, and rising ecological instability. A global information society enabled by a massive electronic communications network of unprecedented bandwidth and computing power has indeed emerged; but far from enabling a 'new world order' of a utopian *or* dystopian sort, the information age heralds an increasingly *dis*ordered geopolitics and widening fractures in the public commons. The rapid amplification of consumerism by converging innovations and ever-shorter product marketing cycles continues apace; yet far from ensuring the oft-predicted rise of technocratic states ruled by scientific experts, the relationship between science, governance, and public trust is increasingly contentious and unsettled.[7] Paradoxically, such tensions appear to be greatest where scientific and technical power have been most successfully consolidated and embedded into our way of life; consider that the nation that gave birth to Apple, Microsoft, Google, Intel, Amazon, and other tech behemoths has slashed federal funding for basic science research, struggled with declining scientific literacy and technical competence among its population, and adopted increasingly ambivalent and politicized science policy—even as it continues to shower the tech industry with tax loopholes and political access.[8]

Such complexities remind us that predicting the general shape of tomorrow's innovations is not, in fact, our biggest challenge: far harder, and more significant, is the job of figuring out what we will *do* with these technologies once we have them, and what they will do with *us*. This cannot be done without attending to a host of interrelated political, cultural, economic, environmental, and historical factors that co-direct human innovation and practice. Indeed, a futurist's true aim is not to envision the technological future but our techno*social* future—a future defined not by which gadgets we invent, but by how our evolving technological powers become embedded in co-evolving social practices, values, and institutions. Yet by this standard, our present condition seems not only to defy confident predictions about where we are heading, but even to defy the construction of a coherent narrative about where exactly we *are*. Has the short history of digital culture been one of overall human improvement, or decline? On a developmental curve, are we approaching the next dizzying explosion of technosocial progress as some believe, or teetering on a precipice awaiting a calamitous fall, as others would have it?[9]

Should it matter whether our future can be envisioned with any degree of confidence? Of course we might *want* to know where we are and where we are heading, but humans characteristically want a lot of things, and not all of these are necessary or even objectively worthwhile. Could it be that our understandable adolescent curiosity about what awaits us in our century's adulthood is, in the grand scheme of things, unimportant to satisfy? Let us imagine for the sake of argument that given certain efforts, we could better predict the future shape of life in this century. Other than idle curiosity, what reason would we have to make such efforts? Why not just take the future as it comes? Why strain to see any better through the fog of technosocial contingencies presently obscuring our view? There is a simple answer. Our growing technosocial blindness, a condition that I will call *acute technosocial opacity*, makes it increasingly difficult to identify, seek, and secure the ultimate goal of ethics—a life worth choosing; a life lived *well*.

Ethics, defined broadly as reflective inquiry into the good life, is among the oldest, most universal, and culturally significant intellectual preoccupations of human beings. Few would deny that humans have always and generally preferred to live well rather than badly, and have sought useful guidance in meeting this desire. Yet the phenomenon of acute technosocial opacity is a serious problem for ethics—and a relatively new one.[10] The founders of the most enduring classical traditions of ethics—Plato, Aristotle, Aquinas, Confucius, the Buddha—had the luxury of assuming that the practical conditions under which they and their cohorts lived would be, if not wholly static, at least relatively stable. While they knew that unprecedented political developments or natural calamities *might* at any time redefine the ethical landscape, the safest bet for a moral sage of premodern times would be that he, his fellows, and their children would confront essentially similar moral opportunities and challenges over the course of their lives.

Without this modest degree of foresight, ethical norms would seem to have little if any power to guide our actions. For even a timeless and universally binding ethical principle presupposes that we can imagine how adopting that principle today is likely to sustain or enrich the quality of our lives *tomorrow*. Few are moved by an ethical norm or ideal until we have been able to envision its concrete expression in a future form of life that is possible for us, one that we recognize as relevantly similar to, but qualitatively better than, our current one. When our future is opaque, it is harder to envision the specific conditions of life we will face tomorrow that can be improved by following an ethical principle or rule today, and such ideals may then fail to motivate us.

While philosophical ethics first emerged in Greece and Asia in the 6th–4th centuries BCE, the need for ethical guidance as we face our future applies equally to modern systems of ethics. Yet modern ethical frameworks often provide *fewer*

resources for mitigating the difficulty posed by an uncertain future than do classical traditions. For example, the ethical framework of 18th century German philosopher Immanuel Kant supplied a single moral principle, known as the *categorical imperative*, which is supposed to be able to resolve any ethical dilemma. It simply asks a person to consider whether she could will the principle upon which she is about to act in her particular case to be universally obeyed by all other persons in relevantly similar cases.[11] If she *can't* will her own 'subjective' principle of action to function as a universal rule for everyone to follow, then her act is morally wrong. So if I cannot will a world in which *everyone* lies whenever it would spare them trouble, then it cannot be right for me to lie.

Although it can be applied to any situation, the rule itself is highly abstract and general. It tells us nothing specific about the shape of moral life in 18th century Europe, nor that of any other time or place. At first we might think this makes the principle *more* useful to us today, since it is so broad that it can apply to any future scenario we might imagine. Yet this intuition is mistaken. Consider the dutiful Kantian today, who must ask herself whether she can will a future in which *all* our actions are recorded by pervasive surveillance tools, or a future where we *all* share our lives with social robots, or a future in which *all* humans use biomedical technology to radically transform their genes, minds, and bodies. How can any of these possible worlds be envisioned with enough clarity to inform a person's will? To envision a world of pervasive and constant surveillance, you need to know what will be *done* with the recordings, *who* might control them, and *how* they would be accessed or shared. To know whether to will a future full of social robots, you would first need to know what *roles* such robots would play in our lives, and how they might transform human interactions. To will a world where all humans enhance their own bodies with technology, won't you first need to know which *parts* of ourselves we would enhance, in what ways, and what those changes would *do* to us in the long run, for example, whether we would end up improving or degrading our own ability to reason morally? Once even a fraction of the possible paths of technosocial development are considered, the practical uncertainties will swamp the cognitive powers of any Kantian agent, paralyzing her attempt to choose in a rational and universally consistent manner.

Modern utilitarian ethics of the sort promoted by 19th-century British philosophers Jeremy Bentham and John Stuart Mill fares little better by telling us that we may secure the good life simply by choosing, among the available courses of action, that which promises the greatest happiness for all those affected. The problem of discerning *which* course of action promises the greatest overall happiness or the least harm—among all the novel paths of biomedical, mechanical, and computational development open to us—is simply incalculable. The technological potentials are too opaque, and too many, to assign reliable probabilities

of specific outcomes. Moreover, technology often involves effects on humanity created by the aggregate choices of *many* groups and individuals. When we factor in the interaction effects between converging technologies, social practices, and institutions, the difficulty becomes intractable.

In their book *Unfit for the Future*, philosophers Ingmar Persson and Julian Savulescu note that the technological and scientific advances of the 20th century have further destabilized the traditional moral calculus by granting humans an unprecedented power to bring about "Ultimate Harm," namely, "making worthwhile life *forever* impossible on this planet."[12] We might destroy ourselves with a bioengineered virus for which we have no natural defenses. Carbon dioxide, nitrogen, and phosphorus from large-scale industry and agriculture may acidify our oceans and poison our waterways beyond repair. Or we might unleash a global nuclear holocaust, a risk that experts warn is once again on the rise.[13] How can existential risks such as these, scenarios that would ruin *any* future possibility for happiness, possibly be factored into the calculation?

Moreover, emerging technologies such as nanomedicine and geoengineering in theory have the potential to forestall 'Ultimate Harm' to humanity *or* to cause it; and not enough is known to reliably calculate the odds of either scenario. Add to this the fact that engineers and scientists are constantly envisioning new and untested avenues of technological development and the insolubility of the moral calculus becomes even more obvious. John Stuart Mill himself noted that the practicality of utilitarian ethics relies heavily upon our collective inheritance of centuries of accumulated moral wisdom about how to maximize utility in the *known* human environment.[14] Even on the timescale of our own lives, this environment is increasingly unstable and unpredictable, and it is not clear how much of our accumulated wisdom still applies.

Given this unprecedented degree of technosocial opacity, how can humans continue to do ethics in any serious and useful way? The question compels an answer; to abandon the philosophical project of ethics in the face of these conditions would not only amplify the risk of 'Ultimate Harm,' it would violate a deep-seated human impulse. Consider once again Ray Bradbury, whose stories are still among the most widely read and appreciated in the tradition of science fiction. What drives the imagination of a storyteller like Bradbury, and what makes his stories resonate with so many? Reading his most lauded works *Fahrenheit 451, The Martian Chronicles,* and the collection *The Illustrated Man* (which leads with 'The Veldt'), one notices how closely Bradbury's vision tracked human beings of a future Earth, or human descendants of Earth. Why this anthropological fidelity in a writer hardly wanting for imaginative horsepower?

Even the Martians in Bradbury's stories serve as literary foils who expose and reflect upon the distinctive powers, obsessions, and weaknesses of human beings.

And why is the human future usually envisioned on a time scale of fifty years, or a hundred and fifty? Why not a thousand years, or ten thousand? Why do so many of Bradbury's tales have a patently *ethical* arc, driven less by saintly heroes and dia-bolical villains than by ordinary, flawed humans working out for themselves how well or how poorly their lives in an era defined by rockets, robots, and 'televisors' have gone? Here is one plausible answer: Bradbury seemed compelled to imag-ine how human beings *more or less like himself*, and those he cared about, would fare in the not so distant technological future—to envision the possibilities for *us* living well with emerging technologies, and more often, the possibilities for our *failing* to live well.

All of this is meant to suggest that the ethical dilemmas we face as 21st cen-tury humans are not 'business as usual,' and require a novel approach. Now, it is a common habit of many academics to roll their eyes at the first hint of a suggestion that the human situation has entered some radically new phase. As a prophylactic against overwrought claims of this kind, these sober-minded individuals keep on hand an emergency intellectual toolkit (which perhaps should be labeled 'Break Glass In Case of Moral Panic') from which they can readily draw a litany of exam-ples of any given assertion of transformative social change being trumpeted just as loudly a century ago, or five, or ten. This impulse is often well-motivated: librar-ies worldwide are stocked with dusty treatises by those who, either from a lack of historical perspective or an intemperate desire to sell books, falsely asserted some massive seismic shift in human history that supposedly warranted great cultural alarm.

Yet sometimes things really *do* change in ways that we would be remiss to ignore, and which demand that we loosen up our scripted cultural patterns of response. At risk of inviting the scorn of the keepers of academic dispassion, I suggest that this is one of those times. The technologies that have emerged in the last half century have led to the unprecedented economic and physical interde-pendence of nations and peoples and an equally unprecedented transmissibility of information, norms, ideas, and values. A great many intellectual and cultural scripts are being rewritten as a result—scripts about modern state power, about socioeconomic development, labor and human progress, and about our relation-ship with our environment, to offer just a few examples. The conventional scripts of philosophical ethics must be rewritten as well. While an irreducible plural-ity of ethical narratives is both inevitable and desirable in a world as culturally diverse as ours, we need a common framework in which these narratives can be situated if humans are going to be able to address these emerging problems of *collective* technosocial action wisely and well. This framework must facilitate not only a shared moral dialogue, but also a global commitment to the cultivation of the specific technomoral habits and virtues required to meet this challenge.

Fortunately for us, a tradition already exists in philosophy that can provide such a framework. That tradition is *virtue ethics,* a way of thinking about the good life as achievable through specific moral traits and capacities that humans can actively cultivate in themselves.[15] Part I of this book explains the distinctive advantages of a virtue-driven approach to emerging technology ethics, and anticipates some of the challenges this project may face. Part II develops the theoretical foundations for our approach. Here we explore the rich conceptual resources of the classical virtue traditions of Aristotelian, Confucian, and Buddhist ethics, from which we construct a contemporary framework of *technomoral virtues* explicitly designed to foster human capacities for flourishing with new technologies. Part III applies the framework to four domains of emerging technology (social media, surveillance, robotics, and biomedical enhancement technology) that are likely to reshape human existence in the next one hundred years, assuming that we are fortunate and prudent enough to make it to the 22nd century.

No ethical framework can cut through the general constraints of technosocial opacity. The contingencies that obscure a clear vision of even the next few decades of technological and scientific development are simply far too numerous to resolve—in fact, given accelerating physical, geopolitical, and cultural changes in our present environment, these contingencies and their obscuring effects are likely to multiply rather than diminish. What this book offers is not an ethical *solution* to technosocial opacity, but an ethical *strategy* for cultivating the type of moral character that can aid us in coping, and even flourishing, under such challenging conditions.

The framework developed in the following chapters adapts Aristotelian, Confucian, and Buddhist reflections on moral development and virtue to our need for a profile of *technomoral* virtues for 21st century life.[16] These will not be radically new traits of character, for they must remain consistent with the basic moral psychology of our species. Rather the technomoral virtues are new alignments of our existing moral capacities, adapted to a rapidly changing environment that increasingly calls for collective moral wisdom on a global scale. In these challenging circumstances, the technomoral virtues offer the philosophical equivalent of a blind man's cane. While we face a future that remains cloaked in a technosocial fog, this need not mean that we go into it unprepared or ill-equipped, especially when it comes to matters of ethical life. The technomoral virtues, cultivated through the practices and habits of moral self-cultivation that we can learn from the classical virtue traditions examined in this book, are humanity's best chance to cope and even thrive in the midst of the great uncertainties and vicissitudes of technosocial life that lie ahead.[17] This hope will only be realized, however, if these virtues are more consciously cultivated in our families, schools, and communities, supported and actively encouraged by our local and

global institutions, and exercised not only individually but *together*, in acts of collective human wisdom. This is a tall order; but not beyond our capabilities.

There is, however, what philosophers call a 'bootstrapping problem.' Our hope of flourishing in this and coming centuries—or even of securing our continued existence in the face of species-level threats created by our present lack of technomoral wisdom—requires us to act very soon to commit significant educational and cultural resources to the local *and* global cultivation of such wisdom. The framework articulated in this book, which draws strength from multiple cultural sources, can help us accomplish just that. Yet our existing technomoral vices, along with the normal human range of cognitive biases and limitations, impede many of us from grasping the depth, scope, or immediacy of the threats to human flourishing now confronting us. Even among those who recognize the dangers, many fail to grasp that the solution must be an ethical one. We cannot lift ourselves out of the hole we are in simply by creating more and newer technologies, so long as these continue to be designed, marketed, distributed, and used by humans every bit as deficient in technomoral wisdom as the generations that used *their* vast new technological powers to dig the hole in the first place!

While the first step out of the hole requires reallocating individual, local, and global resources to technomoral education and practice, we can and must make wise and creative use of technology to aid in the effort. Each of the emerging technologies explored in the book has the potential to be designed and used in ways that reinforce, rather than impede, our efforts to become wiser and more virtuous technological citizens. Thus our way out of the hole is a *recursive* procedure, in which traditional philosophical and educational techniques for cultivating virtue are used to generate the motivation to design and adopt new technological practices that shape our moral habits in more constructive ways. These in turn can reinforce our efforts of moral self-cultivation, forming a virtuous circle that makes us even more ethically discerning in technosocial contexts as a result of increasing moral practice in those domains. This growing moral expertise can enable the development of still better, more ethical, and more sustainable technologies. Used as alternating and mutually reinforcing handholds, this interweaving of moral and technological expertise is a practical and powerful strategy for cultivating technomoral selves: human beings with the virtues needed to flourish together in the 21st century and beyond.

The Motivation of the Book

I was driven to write this book by a deep moral concern for the future of human character, one that arose over many years of watching my own moral and intellectual habits, and those of my students, be gradually yet profoundly transformed

by ever new waves of emerging digital technology. Far from regarding my initial classroom forays into this topic as silly technophobia, my students responded with overwhelming gratitude, even desperation, for a chance to talk about how their own happiness, health, security, and moral character were being shaped by their new technological habits in ways that often bypassed their understanding or conscious choice.

These concerns will be familiar to readers of popular writing on digital culture. Nicholas Carr, Evgeny Morozov, and Jaron Lanier are just a few of the prominent cultural critics who have recently expressed alarm at the possibility, even likelihood, that our mediatized digital culture is undermining core human values, capacities, and virtues. Carr's *The Shallows* warns us of deleterious cognitive and moral effects that our new digital consumption habits may be having on our brains. Morozov's *The Net Delusion* and *To Save Everything, Click Here* challenge our unreflective faith in technocratic 'solutionism.' From Lanier, a computer scientist and pioneer innovator of virtual reality technology, came the widely read humanistic manifesto *You Are Not a Gadget*, which laments the domination of contemporary technosocial life by the increasingly libertarian and antihumanistic values celebrated by many Silicon Valley technologists: unrestrained capitalism, consumerism, and reductive efficiency.

This book shares with these critics a deeply humanistic and explicitly moralized conception of value. It assumes that the 'good life,' by which we mean a human future worth seeking, choosing, building, and enjoying, must be a life lived *by* and *with* persons who have cultivated some degree of ethical character. It assumes that this is the *only* kind of human life that is truly worth choosing, despite the perpetual challenges we encounter in building and sustaining such lives. It also holds that a good and choiceworthy life has never been attained in any great measure by isolated individuals, but only by persons who were fortunate enough to enjoy some degree of care, cooperation, and support from other humans, and who were highly motivated to give the same. This book is therefore fundamentally inconsistent with antihumanistic and neoliberal philosophies, and if Lanier is right, inconsistent with the philosophy of many of those driving the emerging technological developments it proposes to examine.

Yet the reader will also find in this book a resolute hope for the future of human flourishing *with*, not without or in spite of, the technosocial innovations that will continue to shape and enrich our lives for as long as human culture endures. As a scholar who chose out of all possible specialties the philosophy of science and technology, who as a young girl wrote adventure games in BASIC for her Commodore PET and eschewed the Barbie Dream'Vette in favor of Star Wars AT-AT and X-Wing toys, it is simply impossible for me to be antitechnology, personally or philosophically. Indeed, to be antitechnology is in some sense

to be antihuman, for we are what we do and make, and humans have always engineered our worlds as mirrors of our distinctive needs, desires, values, and beliefs. Of course we are not alone—increasingly, researchers find other intelligent animals such as birds, elephants, and cephalopods reshaping their environments and practices in surprisingly skillful and creative ways. Perhaps to be antitechnology is also to be antilife, or antisentience. But 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.