Philos. Technol. (2015) 28:107–124
DOI 10.1007/s13347-014-0156-9

ORIGINAL PAPER

Moral Deskilling and Upskilling in a New Machine Age: Reflections on the Ambiguous Future of Character

Shannon Vallor

Received: 14 August 2013 / Accepted: 2 February 2014 / Published online: 21 February 2014

Abstract This paper explores the ambiguous impact of new information and communications technologies (ICTs) on the cultivation of moral skills in human beings. Just as twentieth century advances in machine automation resulted in the economic devaluation of practical knowledge and skillsets historically cultivated by machinists, artisans, and other highly trained workers (Braverman 1974), while also driving the cultivation of new skills in a variety of engineering and white collar occupations, ICTs are also recognized as potential causes of a complex pattern of economic deskilling, reskilling, and upskilling. In this paper, I adapt the conceptual apparatus of sociological debates over economic deskilling to illuminate a different potential for technological deskilling/upskilling, namely the ability of ICTs to contribute to the *moral deskilling* of human users, a potential that exists alongside rich but currently underrealized possibilities for moral reskilling and/or upskilling. I flesh out this general hypothesis by means of examples involving automated weapons technology, new media practices, and social robotics. I conclude that since moral skills are essential prerequisites for the effective development of practical wisdom and virtuous character, and since market and cultural forces are not presently aligned to bring about the more salutary of the ambiguous potentials presented here, the future shape of these developments warrants our close attention—and perhaps active intervention.

Keywords Deskilling · Virtue ethics · Automation · Artificial intelligence · Robots · New media · Autonomous weapons

1 Introduction

In the twentieth century, emerging technologies profoundly reshaped human practices and institutions around the globe, the economic, political, environmental, cultural, and moral consequences of which we are still struggling to understand. Among the many concepts developed by sociologists, philosophers, economists, and other scholars in an

S. Vallor (✉)
Department of Philosophy, Santa Clara University, 500 El Camino Real, Santa Clara, CA 95053, USA
e-mail: svallor@scu.edu

effort to grapple with these changes, the concept of "deskilling" has been used to frame the way in which twentieth century advances in machine automation resulted in the economic devaluation of practical knowledge and skillsets historically cultivated by machinists, artisans, and other highly trained workers (Braverman 1974). Rooted in a Marxist analysis of capitalism's division of skilled labor as a means of lowering costs while increasing productivity and managerial control (Marx 1990), Braverman's original deskilling thesis has come under considerable critical pressure (Friedman 1977). Yet the information revolution has spawned a new wave of worries about deskilling, as doctors, nurses, pharmacists, teachers, social workers, artists, lawyers, and librarians have seen core aspects of their skillsets made redundant by ICTs (see Apple 1994; Rinard 1996; Carey 2007; Davis 2008; Roberts 2010). The psychological, cultural, and economic impacts of deskilling on workers and their communities can be devastating. Yet technological deskilling is a highly contested and deeply *ambiguous* phenomenon (Wood 1987; Spencer 2000); consider the way in which the computer revolution freed legions of white collar workers from mindless, repetitive tasks such as filing, copying, and collating, arguably creating rich opportunities for reskilling or even upskilling as workers in many fields assumed more knowledge-laden and creative tasks in the information and service economy (Attewell 1987; Adler 1990; Gallie 1991; OECD 2001).

The twenty-first faces its own challenges with respect to technological deskilling; even as the revolutions of the prior century continue to play out, advances in biotechnology, nanotechnology, robotics, neurotechnology, and artificial intelligence are generating new pressures on traditional professions and their associated skillsets. Robotic surgeons are likely to further transform the responsibilities and skillsets of their human counterparts. Will this make the skills of human surgeons redundant, or will it free them to cultivate new and improved capacities for patient care? The improvement of software algorithms may render redundant many of the skills of human actuaries, market traders, and forecasters. Or will this enable humans to not just predict, but better understand future events and their complex causes? Efforts to develop automated systems for grading essays, moderating class discussions, and delivering instructional content invite the question of the future role of the human teacher in the classroom. On one view, teachers may be rendered superfluous, or relegated to the role of mere support techs; on another, teachers may be freed to devote more time to developing creative and stimulating curricula, or to pay more attention to the psychological, emotional, and developmental needs of individual students. Economists, sociologists, and psychologists are continuing to track and reflect upon these pressures and the ambiguous potentials they generate (Ritzer 2008; Heisig 2009). Yet this paper focuses on an important aspect of deskilling/upskilling that has yet to be fully articulated in the literature on the social impacts of emerging technologies, namely the potential in many professional and private contexts for an ICT-driven *moral deskilling* of human beings, along with an equally powerful potential for moral reskilling or upskilling.

In what follows, I adapt the conceptual apparatus of sociological debates over economic deskilling to this presently undertheorized potential for ICT-driven moral deskilling.[1] My general hypothesis is that the potential for moral deskilling is

[1] While the concept of moral deskilling is alluded to in Whitby (1996, 2008), Manders-Huits (2006), and Coeckelbergh (2013, discussed further below), it has not been a central theme of analysis in the literature on the ethics of emerging technologies.

significant and, despite the more optimistic possibilities that manifest alongside it, very likely to present a serious ethical problem in several areas of technomoral practice if current trends prevail. This hypothesis will be fleshed out by means of examples involving automated weapons technology, new media practices, and "carebots." I conclude with the argument that since moral skills are essential prerequisites for the effective development of practical wisdom and virtuous character, and since market and cultural forces are not presently aligned to bring about the more ethical of the ambiguous potentials presented here, the future shape of these developments calls for our closer attention—and perhaps active intervention.

2 The Role of Moral Skills in the Cultivation of Character

Viewed through the lens of a virtue-ethical approach oriented toward questions of character, moral skills appear just as vulnerable to disruption or devaluation by technology-driven shifts in human practices as are professional or artisanal skills such as machining, shoemaking, or gardening. This is because moral skills are typically acquired in specific practices which, under the right conditions and with sufficient opportunity for repetition, foster the cultivation of practical wisdom and moral habituation that jointly constitute genuine virtue. The driving concern of this paper is that profound technological shifts in human practices, if they disrupt or reduce the availability of these opportunities, can interrupt the path by which these moral skills are developed, habituated, and expressed.

On the Aristotelian view that informs much of contemporary virtue ethics, virtues are cultivated rather than inborn states of character. Whether or not a person develops a particular virtue will largely depend on whether they engage repeatedly in the kinds of practices that cultivate it. Yet not all repeated practice leads to virtue, only those practices that successfully engender certain *skills* of acting rightly in particular moral contexts. Such skills can be understood as necessary but not sufficient conditions for full-blown virtue. Many contemporary virtue ethicists have reconstructed Aristotle's account of virtue as a skill model, where the relevant skill is an acquired quasi-perceptual sensitivity to the morally salient features of particular circumstances, or to put it another way, to the moral reasons that such circumstances generate (Annas 1993, 2011; McDowell 1998; Hursthouse 2002; Jacobson 2005). Contrary to the account of moral expertise given by Dreyfus and Dreyfus (1990), in which the relevant skills of the expert are displayed in the non-conceptual activity of "absorbed coping," the Aristotelian account emphasizes the *intelligent* and *rational* dimension of moral skill. While both accounts deny that moral skill can be equated with knowledge of codified moral rules or principles, only the Aristotelian account holds that moral skill entails "a true course of reasoning," ((1984) 1140a20-21) or as Julia Annas puts it, that "the person with skill knows what she is doing and why, and can explain this to others" (1993, p. 67).

It may seem odd that Aristotle would consider virtue to be founded on skill, since he is known for emphasizing habit (*hexis*) as the prerequisite of virtue. However, Aristotle makes clear that the sorts of moral habits he describes are nothing like rote or mindless reflexes, but are rather *skilled practices* guided by an intelligent grasp of the moral demands manifested by particular situations (Lockwood 2013). More specifically, Aristotle claims in Book II, Section 6 of the *Nicomachean Ethics* that virtuous persons

possess the skill-like ability to reliably discern the "intermediate" or "mean" course between an excessive and a deficient response, relative to circumstances. This amounts to an ability to reliably judge how to feel and act "at the right times, with reference to the right objects, towards the right people, with the right aim, and in the right way" ((1984) 1106b20-25). This distinctive ability is acquired through habituation, but the result is much more than a "mere" habit.

While Aristotle describes this moral ability as *analogous* to a skill, he does insist upon a distinction between true skills of craft (*technai*) associated with material production (*poiēsis*), and the moral ability which is far more refined in its accuracy, and which produces nothing beyond moral action (*praxis*) itself ((1984) 1140b6-7). As Annas notes, however, Aristotle could easily have classified the moral ability as a distinctive *species* of skill, a strategy which she and many other virtue ethicists adopt, and which I follow here (Annas 1993, p. 68; Jacobson 2005, p. 389). Once habituated, this moral skill provides a sort of scaffold or stable grafting site upon which virtue can (but may or may not) take hold; for genuine virtue is something more than moral skill or know-how, it is a state in which that know-how is reliably put into action when called for, and is done with the appropriate moral concern for what is good. Someone could have moral skills but fail to be virtuous because they are unreliable in enacting them, or because they act well only for non-moral reasons.

As Aristotle explains in Book Six of the *Nicomachean Ethics*, genuine virtue combines a standing desire to act rightly for its own sake (which the *merely* skilled moral practitioner acting for other motives will lack), the ability to judge well even in novel moral circumstances for which one has not practiced (in which the morally rigid or imperceptive agent will often fail), and a holistic grasp of the general sorts of goods that are worth aiming for in life (missing from a short-sighted or ethically compartmentalized mindset).[2] Virtue must therefore be conceived as a habituated skill of discerning moral judgment *joined* with a moral motivation and aim that guarantees the goodness of its use: "The agent also must be in a certain condition when he does [virtuous acts]; in the first place he must have knowledge, secondly he must choose the acts, and choose them for their own sakes and thirdly his actions must proceed from a firm and unchangeable character" ((1984) 1105a30-35). The end result of this unification of moral skill, moral choice, and moral knowledge is practical wisdom or *phronesis*, the encompassing moral excellence that defines a person of genuinely virtuous character.

Whether the moral skill prerequisite for practical wisdom is successfully cultivated depends on a number of associated factors: whether the agent is exposed to good models of skillful practice from whom to learn; whether the culture's laws and other norms are well-designed to reinforce skillful moral practice; whether the agent possesses the basic motivation and cognitive and emotional resources to acquire the requisite skills; and whether the practical environment supplies sufficient opportunities for habituation and sufficient feedback for the agent to learn from failure.[3] How do we know whether a certain quality or ability of a person constitutes a moral skill? An effective test can be drawn from the key passage of the *Nicomachean Ethics* quoted

[2] A similar account can be found in the writings of Confucian virtue ethicists, in which the skillful moral practice of the "village honest man" (*xiangyuan*) is contrasted with the genuine moral virtue of the exemplary person (Confucius 1998; Mencius 1970).

[3] See Coeckelbergh (2012) for a related discussion of the importance of cultivating moral skills.

earlier: if it is challenging to practice towards the right people, at the right times and places, and in the right manner, then it is a moral skill, not merely a reflex, attitude, belief, or value.

Take the virtue of honesty as an example. A person who has cultivated this virtue is not merely *inclined* to consistently tell the truth—they have learned through repeated practice the *skill* of truth-telling in the relevant social environments. Ideally, such practice will be guided by a virtuous model, i.e., someone who is already good at being honest—that is, someone who already knows when to be honest, and to whom, at what times, and in what way. Yet observation of good models of honesty is not enough. Equally important are repeated opportunities to practice communication in a range of contexts that specifically invite or warrant honesty (such as intimate conversation with close friends and family members, therapeutic discussion, contract writing, religious confession, or scientific/legal testimony). Over time, such practices can lead a person to see for themselves what honest communication is, to do it better and more easily, even in novel circumstances, and eventually, to see being honest as good and desirable in itself.

Moral skills are thus necessary (while not sufficient) conditions of genuine moral virtue or practical wisdom. Moral virtue assumes the availability of prior opportunities to cultivate, through stable and repeatable practices of the right sort, the prerequisite moral skills. Without them, even a person who sincerely wishes to do well consistently and for its own sake will be incapable of doing so. It follows that if new technological practices disrupt the cultivation of moral skills on a large enough scale, the future of human character may be profoundly affected. For this reason, moral skills are even more crucial than other skillsets to shield from widespread loss and cultural devaluation.

3 Moral Deskilling

As noted in the introduction, "deskilling" is a familiar sociological tool of analysis arising from twentieth century studies of the way in which certain skills of machinists and other classes of mechanical labor were made redundant and subsequently devalued by widespread factory adoption of automated machine tools (Braverman 1974). While its merits as a general empirical thesis about the modern capitalist workforce have been widely criticized (Friedman 1977; Attewell 1987; Adler 1990), it remains a credible conceptual tool for the analysis of local socioeconomic phenomena and trends. Recent sociological analyses of deskilling emphasize trends such as those in the nursing profession, where many of the most highly skilled nursing practices are being transferred to advanced medical monitoring and medication delivery systems (Rinard 1996). The deskilling question has been used to analyze the impact of case management and tracking techniques using ICTs on social workers (Carey 2007). Deskilling has also been recognized as a growing phenomenon in the military profession, where rarified abilities such as those cultivated by elite military snipers are becoming increasingly obsolete (Townsend and Charles 2008).

Yet the concept of deskilling has declined in academic use, for several reasons. First, the information revolution has problematized any simple causal link between machine automation of human tasks and human skill loss or devaluation. Computers have

liberated many workers from relatively low-skilled work like filing, copying, and collating, and allowed to them to take on more creative and intellectually demanding tasks in the knowledge and service economy (OECD 2001). ICTs have thus seemed to deliver at least as much upskilling as deskilling. As a result, the impact of emerging technologies on human skill development has come to be seen as a deeply ambiguous phenomenon, something which a lopsided conceptual emphasis on deskilling seems to miss. A second reason for the decline in usage is that the concept of deskilling is strongly associated with the historical debates about labor, commerce, and economic justice in which it has been used, and is therefore heavily charged with political and ethical implications. Given that many sociologists today try to avoid infusing their work with normative or polemical commitments, the concept of deskilling has lost some favor as a tool of neutral scientific analysis.

The concept of *moral deskilling* is used even less frequently in contemporary academic discourse. One reason is the aforementioned avoidance of normative commitments in social science research. The other is the tendency to associate fears of moral deskilling with unreasoned "moral panics" in response to technological change. New technologies have often given rise to such panics, especially when those technologies alter established social norms and practices. Yet emerging developments in twenty-first century ICTs may warrant renewed concerns about deskilling, especially moral deskilling. Concerns about moral deskilling need not reflect reactionary politics, unexamined dogma, or unreasoned panic. They may be rooted instead in a serious and critically minded investigation of the developmental conditions of human character, virtue, and flourishing. Such an investigation would be at home within the normative structure of philosophical research into ethics, but it could also draw upon a richer interdisciplinary foundation combining virtue ethics, moral psychology, and rigorous empirical study of the material and social conditions of human flourishing. It must also address the inherently ambiguous character of technology's impact on moral skills. Yet as I will show, there are several areas of technological advancement in which the risks of moral deskilling are not presently balanced by forces driving moral upskilling or reskilling. Thus one objective of such an investigation could be to identify specific institutional and cultural changes that might help to remedy this imbalance.

Why should we care about technology's impact on moral skill development enough to study it? One reason was offered in the prior section; if moral skills are a prerequisite to the cultivation of virtuous character, then we should be greatly concerned about the impact of emerging technologies upon them. But perhaps the reader is skeptical about the existence of genuine virtue, or the human need for it. For that reader, I offer a second reason. Moral skills are *intrinsically valuable*—even when they do not lead to the development of truly exemplary persons of virtue. Consider how many have lamented the skills of artistic handicraft lost to mass manufacture of ready-made objects (Roberts 2010), and the resulting resurgence of interest in "handmade," "custom," or "artisanal" products. I submit that our interest is not only in the economic welfare of the artisan, or the quality of their products, but also in the connection between an artifact and a human whose excellence and skill was responsible for its production. We think that it is *good* that we humans are good at making beautiful and commodious objects for living. Even if machines could produce all such goods for us, it would be sad and regrettable if humans were no longer capable of doing the same. Wouldn't we want to say the same in the case of moral skills? Even if intelligent machines could somehow

direct all human interactions to produce the most just, harmonious, and compassionate outcomes possible, *we would be diminished as creatures* were we utterly helpless to act justly and compassionately without their assistance.

This view is inspired by and largely congruent with Albert Borgmann's discussions in *Technology and the Character of Contemporary Life* of the good life for a human being as presupposing the skill of "intensive and refined world engagement," which is "bound up with social engagement" and which "molds the person and gives the person character" (1984, p. 42). Borgmann is explicitly concerned with deskilling, both in the context of modern labor practices (p. 118) and in the broader philosophical context of the *device paradigm*, in which technology makes increasingly fewer demands upon our "skill, strength and attention" degrading the quality of our being in the process (p. 42). Yet my view departs from Borgmann's in an important way: while Borgmann does acknowledge that technology can and should play a supportive and enabling role in the good life, his program for reform of technological culture repeatedly emphasizes "limitation" and "restraint" of technology, the value of removing "technological clutter" from our lives, and the ability to use technology only "at certain times and up to a point, one that is left behind when we reach the threshold of our focal and final concerns" (p. 220). Borgmann maintains a view of embodied human excellence and technological excellence as separate spheres, the latter of which must contract and withdraw as the former is engaged. On the view I develop here, this rests on a false dichotomy between technology and the human lifeworld, for as philosopher of technology Peter-Paul Verbeek rightly notes (2011), technology has always conditioned our humanity, and our morality as well. We might say that we are *technomoral* creatures to the core.

Thus as we will see, not *all* forms of moral assistance from technology diminish us as creatures. The conceptual tools of philosophical ethics are themselves a technology for amplifying our moral capacities, and we are not ordinarily diminished by our use of these.[4] Yet we *could* be, if these tools somehow obviated the need for us to employ our own skillful moral judgment, rather than aiding us in its successful exercise. Fortunately even fixed principles such as Kant's categorical imperative, the greatest happiness principle, and the "golden rule," if they are to be usefully applied, demand considerable interpretive effort and our discerning attention to the morally salient features of the given situation. An automated system that produces reliable and situation-specific moral judgments in a particular domain by employing its own perceptive and interpretative capacities, with minimal human input, is an entirely different story.

4 Concrete Technological Risks of Moral Deskilling

Below I identify three potential hazards of moral deskilling posed by today's emerging technologies, intended to supply only a representative sample and not an exhaustive list: the risk that military drones and other autonomous weapons systems may engender moral deskilling of soldiers in the use of military force; the risk that new media practices of multitasking may engender deskilling in the realm of moral attention; and the risk that social robotics may engender moral deskilling in practices of human

[4] Thanks to an anonymous reviewer for recommending this clarification.

caregiving. Building on my prior analyses of the potential for technological disruption of such practices, as well as the moral virtues they can foster (Vallor 2010, 2011, 2013), I argue here that these risks are real and of profound ethical significance.

4.1 Autonomous Weapons Systems

Present trends in the development of military technology reflect the growing autonomy of weapons systems based on information and communications technologies, from self-piloting unmanned aerial vehicles (UAVs) like the US X-47B and UK Taranis armed long-range drones, to the "fire and forget" Israeli Harpy drone which can loiter and destroy radar installations independently of human control, to South Korea's Samsung SGR-1 armed sentry robot, to automated algorithims for detecting, repelling, and disabling hostile computer agents. Many have suggested that due to the calculational speed and algorithmic precision of these systems, the human soldier is already, or will be soon, "the weakest link in the military arsenal" (Heyns 2013, p. 10); the conclusion of many observers is that the goals of military efficiency, force protection, and mission success will almost inevitably result in the large-scale deployment of fully autonomous weapons, unless strong global measures to renounce such weapons are adopted (Sparrow 2009).

There are many reasons to be concerned about such a trajectory, including but not limited to potential conflicts with the requirements of just war theory (Asaro 2008), the lowering of barriers to armed conflict (Sharkey 2010), the problematic assignment of responsibility to electronic agents (Sparrow 2007), and the potential for abuse of such systems by rogue states or agents. Each of these considerations was taken up at a May 2013 meeting of the United Nations to address growing calls for a moratorium on these technologies. Yet there is another reason for concern about these developments that is less likely to find its way onto the international agenda: the potential for a debilitating moral deskilling of human military forces. One may wonder what moral skills we can be talking about here; soldiers are in the business of, among other things, killing people, and killing people does not immediately jump to mind when we think about morally skillful actions. But this is a mistake; war itself may be immoral, but as I have argued elsewhere (2013), the conduct of killing in war demands considerable moral skill if it is not to descend into utter moral chaos, in which the lines between civilian and combatant, friend and foe, military necessity and mindless vengeance do not just get blurred (as they do in all wars), but are wholly abandoned. The right use of military force meets our test of a moral skill; force is something that must be applied only at the right times, towards the right persons, in the right places, and in the right manner. Indeed, the principles of *jus in bello* function primarily to formalize this requirement.

There exists a rich institutional tradition of education for virtue in most modern professional armies, especially within the officer or other leadership corps where the broad concept of "military virtue," and the specific norms such as courage, loyalty, discipline, and service that go along with it, carry special resonance (Miller 2004). But as with all virtues, they cannot be acquired in a classroom, or even in a simulator. Only in the actual practical context of war, where situations are neither stable nor well-defined and where success and failure have lifelong moral consequences, can words like "courage" and "discipline" be more than empty slogans or aspirational terms that cannot by themselves direct one to their achievement. The possession of virtue in

combat, as opposed to the mere celebration of it, presupposes the prior cultivation of moral *skills* relative to combat, through some form of practice that is capable of developing those skills.

It is well known that such skills are hard to acquire, and even harder to preserve in the madness of war, and that on the whole human soldiers do not perform nearly as well as we would hope in the moral use of force. This is why roboticists like Ron Arkin (2009) suggest that making robot soldiers more ethical than their human equivalents may not be too high a bar to clear. But it would be a mistake to draw the conclusion from our aggregate failure to fight morally that military virtue is a chimera. A careful study of military history will reveal countless astonishing displays of moral restraint and practical wisdom, both great and small, alongside every atrocity. Many of these moments are lost to cultural memory, some are ignored or concealed, but a few broaden our consciousness of why it matters whether or not we fight with moral character. Consider, for example, the actions of US Army Warrant Officer Hugh Thompson Jr.'s famous intervention in the 1968 massacre at My Lai, in which Thompson and his helicopter crew shielded and rescued Vietnamese civilians being willfully massacred by American soldiers, even aiming their weapons at their fellow men and risking their lives, reputations, and careers in the process. But many more quiet moments of moral skill in war transpire daily all over the globe; they occur every time a civilian crossing the field of fire is discerned and warned away, every time a call is correctly made as to whether a given response is proportional or necessary, and every time a gravely wounded and disarmed enemy is spared by opposing forces.

Now imagine a future in which the "dirty and dangerous" business of killing in war is farmed out to robots, drones, and other artificially intelligent and/or autonomous weapons systems. Imagine for the sake of argument, though of course this is implausible to some, that the autonomous weapons fulfill the roboticists' wildest hopes and fight more ethically than humans do. Humans in this future would not even be in the position to serve as supervisors of such agents, except in a very remote sense; for they will lack comparable battlefield experience to recognize whether a given robot or computer's call is the morally right one. A supervisor, in order to be a legitimate authority, must have more experience and practical wisdom than the supervisee; in this scenario, what wisdom will future humans "on the loop" be able to offer the machines in this regard? Furthermore, the human supervisor will lack equivalent computational powers of the supervised agent, and unless machine agents run on brain-emulating software, the human supervisor will not share with them a comparable decision architecture. On what grounds will she examine and evaluate the moral wisdom of a machine agent's decision?

Now ask yourself a variant of the question I posed earlier; in a world where machines fight more ethically than we do, and in which as a result humans are no longer skilled in the moral uses of force, should we regret that result? Set aside for now the obvious and considerable benefits, from a consequential perspective, of reduced civilian losses in this scenario. Humans in this world are still responsible for wars; they still start them, fund them, and continue them. But they no longer really know, except in the most abstract and impotent sense, how to fight them morally. Is this itself a loss? Is it one we have reason to care about? I assert that we have at least a *prima facie* basis for thinking so.

Are there other possibilities for the development of intelligent military systems in which these would actually enhance, rather than make redundant, the moral skill and

practical wisdom of human soldiers? Certainly, this is no exception to the rule of ambiguous technology. For example, rather than developing artificial moral intelligence that supplants human decision-making in the use of lethal force, artificially intelligent systems might instead be usefully deployed to provide soldiers with enhanced information about morally salient features of the battlefield for use in their deliberations. An artificially intelligent miniaturized drone might be able to record conversations in a multitude of language and dialects, and decode more quickly than a human whether the conversation pattern in a targeted vehicle indicated a friendly or hostile presence. Or intelligent software agents might offer improved feedback concerning the alignment of soldiers' habits and decision patterns with norms of military honor, courage, and restraint—a kind of moral "biofeedback" system.

Moreover, the ambiguous nature of technology guarantees that systems depriving soldiers of one kind of moral skill or knowledge may in fact enhance another, in ways that have complex and unintended ethical consequences. Mark Coeckelbergh's analysis of drone pilots notes that while technologies that distance soldiers from killing appear to "de-skill" them by depriving them of "knowledge that is grounded in lived bodily experience, in handling things on the ground, in skillfully engaging with what happens on the battlefield and with others" (2013, p. 94), the extended surveillance and acute visual imaging of drone targets actually end up *enhancing* the pilot's moral knowledge of what they are doing, and to whom. The outcome, however, is hardly desirable. Drone pilots are arguably exposed to greater psychological and moral trauma from killing than they would if they lacked this moral knowledge of their targets' daily habits, work, and family life, but this richer moral context arguably does little good, since the technology affords them no other mode of practical engagement with their targets than passively observing or killing them. One might regard this as the worst of both worlds.

Yet as we noted above, we *can* imagine a host of other possibilities in which autonomous and/or artificially intelligent agents deliberately support and reinforce, rather than diminish or disrupt, the cultivation and use of moral skill and practical wisdom in war. What stands in the way of these possibilities being realized, as opposed to the possibilities currently receiving active military funding? This is a question that cannot be fully answered here; but perhaps the largest obstacle is a dangerous and deep-seated cultural tendency for narrowly instrumental thinking when it comes to realizing the ambiguous potential of technology. We still think of technologies as neutral tools for accomplishing things, rather than as systems that merge with, mediate and continually transform our own agency and capacities (Verbeek 2011). Thus our first thought is always "How can I use this new, more powerful technology to do something *for* me?" Very rarely is our first question, "How should this new technology be related *with* me?" or more precisely, "How can I relate to this technology in a way that helps me become the sort of human being I would like to be?" Military technologies are not *merely* tools to accomplish practical objectives; increasingly, they are the medium that defines what *kind* of soldier a soldier can choose to become, and what moral virtues she can develop through her service.

4.2 New Media Practices and Multitasking

A related pattern can be seen in the case of new media practices, which I have argued elsewhere (Vallor 2010) have the ambiguous potential to either weaken the moral skills

and virtues engendered by communicative practices, or to encourage and strengthen their cultivation. A range of morally significant skills could be impacted positively or negatively by new media practices. But following up on clues provided by a growing body of empirical research, I wish to draw attention here to one in particular: the skill of *paying attention*. Attention is not just a cognitive ability, it is a moral one as well. This should not surprise us; cognitive skills should be expected to be partly constitutive of moral intelligence as well as other types of intellect. More specifically, knowing when to pay attention, to whom, for how long, and in what manner (and being able to do so successfully and habitually) is just as critical to cultivating a virtuous character as knowing when to be angry, or when to forgive. A person who cannot be counted on to pay attention when you tell her about the recent death of your closest friend, or who is unable to stay focused on the grave and imminent danger to which you're trying to alert her, or who cannot attend to the expressions on your face during an intense conversation, is not someone who can be said to be virtuous. This is true even of a person who makes a sincere effort to pay attention to her social environment but who has unwittingly lost the cognitive ability to succeed in this task. This possibility is especially worrisome for reasons I am about to explain.

Empirical studies continue to bear out the common perception that compared with other populations, new generations of digital natives are increasingly prone to media multitasking (Carrier et al. 2009; Rideout et al. 2010). "Media multitasking" is the practice of consuming multiple information streams at once (social media feeds, text messages, photos, Internet videos, television, homework, music, phone calls, reading, etc.), and/or carrying on multiple information exchanges simultaneously. However, an impressive and growing body of evidence suggests that the habit of multitasking brings with it considerable cognitive costs; not only does multitasking significantly impede performance of tasks drawing upon working memory and attention, but chronic multitasking may have lasting negative effects on human cognitive abilities, leaving us more distractible and less efficient at refocusing our attention (Ophir et al. 2009; Wang and Tchernev 2012).

In their recent article "The Myth of Media Multitasking," researchers Wang and Tchernev set out to answer the obvious question that follows. If media multitasking is so damaging to our cognitive abilities, why do we (and especially young people) continue to do it? The answers they found are unsettling on multiple levels. The researchers first assumed that young people must be drawn to multitask because they see it as a way to satisfy some perceived need or needs, and they must continue to do it because they receive some gratification from it. But what kind of needs, and what kind of gratifications? What they found was that their subjects were motivated to multitask primarily by two kinds of needs: cognitive (the perceived need to strengthen "information, knowledge, and understanding") and habitual ("ritualized media use" that seeks to reinstate a familiar pattern) (p. 495). But as we have explained, the benefits multitaskers receive are *not* satisfactions of a cognitive sort; media multitasking actually *impeded* subjects' ability to meet those needs. Yet the multitasking did provide strong gratifications of a habitual sort—it "scratched the itch," if you will. And "like a locomotive picking up steam," (p. 511) the more subjects multitasked on a given day, the more they found themselves compelled to do it the next day. Unexpectedly, the multitasking also gave many subjects *emotional* gratification, even though this was not one of the needs they sought to meet by multitasking. Furthermore, the subjects who

got the most emotional gratification were the ones with the *lowest* preexisting emotional need—on the other hand, multitasking subjects with genuinely high emotional needs received far less emotional gratification. Media multitasking, then, appears to rob us of what we do really need (more cognitive strength), while satisfying only superfluous or "phantom" needs largely generated by the practice itself.

Unless these data or conclusions are deeply flawed, then it is hard to not to see media multitasking as a highly pernicious trap that we have allowed ourselves, and worse, our children to fall into. We have accepted as "normal," even "optimal," a technological practice that undermines both our cognitive and moral ability to pay attention, while offering us nothing useful in return. Why have we allowed this, and how can we turn the ship around? It is not as if the problem has been entirely invisible to us. Aside from empirical studies and other works targeting the "new distractibility," hyper-connected adult users often trade anecdotal complaints of noticeable declines in their ability to concentrate on a difficult text or an important conversation for an extended period of time; the compulsion to check Facebook messages, update a Twitter status, or check out a new Tumblr page is simply too great for many of us. Wang and Tchernev describe the feedback structure of multitasking's emotional rewards as following the structure of classical conditioning (2012, p. 509) so perhaps we should not be surprised to find ourselves with these unexplained compulsions to do what we neither truly want to do nor should do.

Many who have become frustrated with a perceived ICT-related decline in their own cognitive and moral functioning have sought solace in a surprising place—ICTs themselves. A variety of apps and software programs with names like *Freedom*, *AntiSocial*, *Blinders*, *Self-Control*, and *Concentrate* are now being marketed to save you from your digital self—often by locking out access to social media and other distractions for a set period of time. This is arguably an imperfect solution, however. Why must we choose between tools that debilitate us and digital lockboxes to keep us away from these tools? Why is it not possible to have *useful tools that do not debilitate us?* In principle, nothing stands in the way. Rather than offering a constantly refreshed torrent of 140-character tidbits and a hydra's head of unrelated digital content streams, new media could in principle facilitate creative collaborations or extended conversations about mutually engaging topics that do not encourage the endless intrusion of messages and viewing opportunities of lesser importance.

And, of course, many of us try to will ourselves to use the tools for these purposes, and sometimes we even succeed. But this is a fight against the current; the tools are not designed for such purposes and must be repurposed or patched in ad hoc fashion, often with great effort and cost to usability. What stands in the way of it being easier? Once again, there is a complex, more adequate answer than what I am about to say; but it is worth pointing out that developers, media experts, and users have continually asked what these technologies can do *for* us, but not what we want them to do (or what they are currently doing) *with* us. The choice is not between surrendering to technology or liberating ourselves from it. We are *technomoral creatures* to the core; that is, we allow and have always allowed the things we make to reshape us. The only question is whether this process is deliberate and wise or unreflective and reckless. Our present course of technological development has no moral teleology; how could it, as a process directed by free market forces committed to nothing but the satisfaction of immediate and unreflective desire? Regardless of your view of free market ideology, it is one thing

to let market forces determine how best to serve the genuine needs and desires humans have. It is quite another to let the free market determine what quality of moral beings we are, or shall become.

4.3 Robot Caregivers

Here I will offer just a few brief words about our final example, a topic I and others have discussed elsewhere—the ethical implications of social robotics, especially "carebots": robots being developed for the purposes of caregiving (Sparrow and Sparrow 2006; Decker 2008; Vallor 2011; Parks 2010; Sharkey and Sharkey 2012; van Wynsberghe 2013). Amidst mounting demographic and economic pressures in many nations to supply more and better care for the aged and other dependent persons, a new market opportunity has arisen for developers of robots that can potentially provide such care. One unresolved question is how *good* or even adequate robotic care could be; yet my claim is that *even if* robots turn out to be effective and welcome caregivers for human beings, that does not resolve the moral quandary they present. For in a possible scenario where robots *are* more effective and/or welcome than human caregivers, the capacity for the human cultivation of caring practices is that much more endangered, along with the attendant moral skills and virtues that caring engenders. Caring, after all, is recognized as a burden that humans (especially women) must bear, and what are technologies for if not to ease or even deliver us from our burdens? But this again is to ask the question of what technologies can do *for us*, rather than the more illuminating question of what they may do *with us*.

First, let us clarify why caring is relevant to worries about moral deskilling. "Caring" is an ambiguous term; it can refer simply to an attitude of appropriate concern, which would not be appropriate to characterize as a skill: ("He really cares."). But caring as an active practice of offering support to someone who needs it *is* a skill. It is difficult to know how to care for people *well*—emotionally, physically, financially, and otherwise, in the right ways, at the right times, and for the right persons.[5] Efforts at care often go awry despite our best intentions; to use a Heideggerian phrase, we may "leap in" for someone and prematurely appropriate their agency, or we wait far longer than we should to help, out of a misplaced sense of respect and deference. It takes time, and practice, and intelligence, and good will to become even moderately good at it, and even then we may botch the job. But persons who care for others well, or are perceived as such, are among the models of virtue that today we recognize most quickly and resolutely. Why, after all, did Mother Theresa's name essentially become a modern placeholder for the concept of a good person, when in historical fact her life and choices were hardly free of moral controversy or immune from challenge? The answer is that to many people, her name was a kind of moral shorthand for the apex of skillful caring, and even those who knew almost nothing about her life would have explained their attribution of virtue to her in terms of care. The skills of caring are important for humans to have because a person who cannot care well for others is incomplete in a fundamental way; whatever other virtues they may possess, the one that is lacking precludes our regarding such a person as a model for our own life, or for that of our children.

[5] See Held (2006), Kittay (1999), and Tronto (1993) for a fuller articulation of the skills implicit in successful caring practices, and Parks (2010) for a discussion of these as related to carebots.

Jennifer Parks, in her feminist analysis of the prospect of carebots, paraphrases care ethicist Annette Baier (1985) as noting that we become "skilled in the arts of personhood through relationships of dependency that we share with others" (2010, p. 111). Parks holds that this relational concept of the human person entails a moral problem with the development of carebots designed to take over, rather than *assist* in, human caring relations. Yet her focus is the danger that robotic carers may deprive those cared for of the essential relational benefits of "conversation and human touch" (p. 112). This is indeed a legitimate fear; but it is too easily blunted by developers' promises (however unlikely to be realized) that robotic conversation and touch will someday be as good as their human equivalents. The question that is not often asked, but should be, is how the skills cultivated by the caring arts of human conversation and touch can be preserved for the *carers*. For conversing *well* with the aged, or sick, or developmentally challenged, or otherwise vulnerable person demands far greater caring skill than any other kind of conversation, and the same is certainly true for the case of touch. It is thus more *difficult* to cultivate these caring skills, and the rewards of their development are sometimes less tangible and immediate than in other relational contexts, despite their profound importance. We would be wise to be wary of carebots that offer to deliver us from such "burdens."

Aimee van Wynsberghe has developed a similar relational approach to carebot ethics, but which offers concrete design guidance to robot developers. Following care ethicist Joan Tronto (1993), van Wynsberghe (2013) frames the goal of caring practices in terms of four moral elements: attentiveness, responsibility, competence, and reciprocity. She claims that robotic developers ought to consider how their designs impact these elements of caring relations in a variety of caring contexts, with the goal of developing carebots only to "safeguard the manifestation of care values," rather than displacing them (2013, p. 424). Like Parks, van Wynsberghe emphasizes the importance of eye contact and touch in attentive, responsible, competent, and reciprocal caring relations, but in terms of skills to be preserved, not mere actions to be performed. Using the example of an automated system for lifting patients, she notes that a nurse "lifting" a patient with the help of a machine is hardly engaged in skillful caring if the nurse's attention throughout must be devoted to monitoring the machine controls and not the patient—even if the end result is a safe and efficient lift. Far better would be a design such as a mechanical exoskeleton that ensures a safe outcome while still allowing the nurse to maintain eye contact with the patient, verbally inquire after the patient's comfort and sense of security, and offer reassuring touch throughout the process. Thus the goal here is not to design a robot that will perform or even assist humans in isolated caring *tasks*, but to design a robot that will support the skillful carrying-out of holistic caring *practices* (p. 429). In this way, a carebot may not only benefit patients but also help meet the moral needs of caregivers, by allowing them to become more skillful carers (p. 430).

But how plausible is it that this concern will be adopted by carebot developers? Parks notes that the capitalist forces currently driving carebot development favor far narrower goals (2010, p. 112). Cutting health care costs (while minimizing liability exposure from injury to patients or property) seems likely to be the overwhelming priority. Is it therefore *inevitable* that we will face a future of robotic carers who make human caring skills redundant? Hardly and not just because of the technical challenges posed by their development. There is simply no good reason (*including* market forces)

for robots to be designed as full-fledged surrogates for human care—*unless we are already resigned to the idea that care is no longer an ineradicable human responsibility.* Instead of quality substitutes for human care, the design and implementation of "carebots" could be narrowly targeted at the relief of overburdened caregivers from those specific tasks that drain their capacities for care or lead to carers' physical and mental deterioration, ultimately supporting and encouraging caring skills and relations of greater consistency and higher quality. Advanced software agents could serve as inexpensive support for human carers, checking in with them, answering questions, linking them to needed resources and to support groups of other caregivers, or taking care of legal or financial matters that overwhelm and distract them from their care.

Can we be assured that developing these forms of moral support will be primary goals of those developing and marketing carebot services? Unfortunately not. *Why* not? Let us avoid the lazy answer: "Because capitalism." For to blame market forces is to conceal our own responsibility for how we presently conceive our relations to technology, markets, and human goods. The problem is not some mysterious, disembodied "force" of markets, but the *cultural* forces driving their operation and the realization of our technology's ambiguous potentials. To say that the relationship between emerging technologies and moral skills or virtue is an ambiguous one does not mean that the relationship will turn out to be half bad and half good, or that even good options will have ethical trade-offs (though they certainly will). It means only that there are potentials on both sides of the ethical spectrum, but the cultural values presently guiding the realization of those potentials are not neutral or random. They are the narrow, short-sighted values of technological instrumentalism presently permeating technoscientific societies, and as things stand, these are more likely to lead to suboptimal ethical results than to salutory ones because such aims are largely indifferent to ethical values. In my concluding remarks, I offer some passing thoughts on what it would take to transform those aims and values.

5 Conclusion

I have argued that the ambiguity that has long characterized the phenomenon of industrial and professional deskilling, and that continues to drive robust economic and sociological debate about that phenomenon's shape and trajectory, turns out to be a structural feature of the phenomenon of moral deskilling as well. Philosophers and ethicists of technology reflecting on the risks of moral deskilling must explicitly anticipate this ambiguity and be prepared with the conceptual tools to analyze it. The proper aim of such analysis will not be the elimination of the phenomenon's basic ambiguity, but rather its effective use as a stimulus for our practical and moral imaginations. Understanding the ambiguous and open horizon of emerging technologies' impact on moral skills can help us not only to better identify the moral risks and benefits of such technologies in the abstract but also to envision new developmental and design possibilities that concretely realize potential moral benefits of these technologies, while effectively minimizing the losses they might otherwise inflict on our own moral capacities and excellences.

Yet envisioning those developmental and design possibilities is not enough. One could fill a library with all the possibilities for ethical technologies that have been

envisioned and never built or sold. Movements that privilege "values by design" or "participatory design" are a fine start, but will not take us all the way. What would be required to feed our moral imaginations more reliably into the process of technological development, manufacture, and mass marketing? I assert that what would be needed is a deeper shift of cultural values in technological societies, one that would reflect a new global awareness that humans are technomoral artificers *and* artifacts. We can no longer afford to view technological development as something to be delegated to engineering wizards and marketers, who make tools to be adopted by human consumers; human beings must acquire a sense of collective responsibility for intelligently directing the course of our own technomoral evolution—for consciously using technologies to aid us in becoming the practically wise beings that we want to be, beings who can live the kinds of lives we genuinely want to live. The imagination of the moral self must lead the way for the "technologies of the self."[6]

I am hardly the first to point this out; Hans Jonas called this challenge philosophy's first "cosmic task" (1979). But philosophers alone cannot carry out the task; in a society that is increasingly unphilosophical, we cannot count on the culture to follow our lead. In his remarks on artificial intelligence and the associated threat of moral deskilling, Blay Whitby notes that we should not be surprised if the majority of humans show no interest in cultivating or preserving their own skills of moral judgment, and prefer to defer to authority. Yet he claims that fears about moral deskilling are overstated, since "as with various other outmoded skills, some humans will choose to keep this skill alive" (2008, p. 559). Far from reassuring us that the concerns I have raised in this article are exaggerated, this remark, which so casually dismisses "the ability to form moral judgments" as an "outmoded skill," should remind us that our present cultural vulnerability to moral deskilling is all too great. The Enlightenment ideal of *Bildung* is a fading cultural echo, and the call for a new commitment to moral self-cultivation has yet to be heard. The awareness of and desire to meet this challenge, to actively cultivate our technomoral selves, must somehow be introduced within technological culture itself, not only through philosophy but through art, literature, politics, and other means of empowering the moral imagination.

The future of moral character thus remains ambiguous, not least because the horizon of our technological choices remains open to the possibility of a new era of technomoral responsibility. For this hope, and our remaining chance to realize it, we ought to be profoundly and enduringly grateful.

References

Adler, P. S. (1990). Marx, machines and skill. *Technology and Culture, 31*(4), 780–812.
Annas, J. (1993). *The morality of happiness*. Oxford: Oxford University Press.
Annas, J. (2011). *Intelligent virtue*. Oxford: Oxford University Press.

[6] Here I rely on a double-meaning that alludes both to Foucault's (1988) use of the phrase "technologies of the self" to refer to classical philosophical techniques of moral self-cultivation, as well as a more literal reference to the technologies that mediate our worldly relations and hence guide our self-formation in the world. My claim is that we retain the ability to imagine and develop new and alternative forms of these relations, even as our present powers of moral imagination continue to be shaped by our existing technology relations. The future trajectory of human moral abilities is technologically conditioned, but not technologically determined.

Apple, M. W. (1994). Computers and the deskilling of teaching. *Computer Professionals for Social Responsibility Newsletter, 12(2)*. http://cpsr.org/prevsite/publications/newsletters/issues/1994/Spring1994/apple.html/. Accessed 13 January 2013.

Aristotle. (1984). *The complete works of Aristotle.* Revised Oxford translation. Edited by J. Barnes. Princeton: Princeton University Press.

Arkin, R. (2009). *Governing lethal behavior in autonomous robots*. Boca Raton: CRC Press.

Asaro, P. (2008). How just could a robot war be? In P. Brey, A. Briggle, & K. Waelbers (Eds.), *Current issues in computing and philosophy* (pp. 50–64). Amsterdam: IOS Press.

Attewell, P. (1987). The deskilling controversy. *Work and Occupations, 14*(3), 323–346.

Baier, A. (1985). *Postures of the mind: essays on mind and morals*. Minneapolis: University of Minnesota Press.

Borgmann, A. (1984). *Technology and the character of contemporary life*. Chicago: University of Chicago Press.

Braverman, H. (1974). *Labor and monopoly capital: the degradation of work in the twentieth century.* New York: Monthly Review Press.

Carey, M. (2007). White-collar proletariat? Braverman, the deskilling/upskilling of social work and the paradoxical life of the agency care manager. *Journal of Social Work, 7*(1), 93–114.

Carrier, L. M., Cheever, N. A., Rosen, L. D., Benitez, S., Chang, J. (2009). Multitasking across generations: Multitasking choices and difficulty ratings in three generations of Americans. *Computers in Human Behavior, 25*(2), 483–489.

Coeckelbergh, M. (2012). Technology as skill and activity: revisiting the problem of alienation. *Techne, 16*(3), 208–230.

Coeckelbergh, M. (2013). Drones, information technology, and distance: mapping the moral epistemology of remote fighting. *Ethics and Information Technology, 15*, 87–98.

Confucius. (1998) *Analects* (trans: Lau D. C.). London: Penguin Classics.

Davis, A. (2008). Outsourced radiology: will doctors be deskilled? *British Medical Journal, 337*(7669), 542.

Decker, M. (2008). Caregiving robots and ethical reflection: the perspective on interdisciplinary technology assessment. *AI & Society, 22*, 315–330.

Dreyfus, H. L., & Dreyfus, S. E. (1990). What is morality? A phenomenological account of the development of ethical expertise. In D. Rasmussen (Ed.), *Universalism vs. communitarianism: contemporary debates in ethics*. Cambridge: MIT Press.

Foucault, M. (1988). In L. H. Martin, H. Gutman, & P. H. Hutton (Eds.), *Technologies of the self: a seminar with Michel Foucault*. Amherst: University of Massachusetts Press.

Friedman, A. (1977). *Industry and labour: class struggle at work and monopoly capitalism*. London: Macmillan.

Gallie, D. (1991). Patterns of skill change: upskilling, deskilling or the polarization of skills? *Work, Employment and Society, 5*(3), 319–351.

Heisig, U. (2009). The deskilling and upskilling debate. In R. Maclean & D. Wilson (Eds.), *International handbook of education for the changing world of work* (pp. 1639–1651). New York: Springer.

Held, V. (2006). *The ethics of care: personal, political and global*. Oxford: Oxford University Press.

Heyns, C. (2013). Report of the Special Rapporteur on extrajudicial, summary or arbitrary executions. United Nations General Assembly, Human Rights Council. http://www.ohchr.org/en/issues/executions/pages/srexecutionsindex.aspx. Accessed 21 November 2013.

Hursthouse, R. (2002). *On virtue ethics*. Oxford: Oxford University Press.

Jacobson, D. (2005). Seeing by feeling: virtues, skills and moral perception. *Ethical Theory and Moral Practice, 8*(4), 387–409.

Jonas, H. (1979). Toward a philosophy of technology. *Hastings Center Report, 9*(1), 34–43.

Kittay, E. (1999). *Love's labor: essays on women, equality and dependency.* New York: Routledge.

Lockwood, T. (2013). Habituation, habit and character in Aristotle's ethics. In T. Sparrow & A. Hutchinson (Eds.), *A history of habit: from Aristotle to Bourdieu* (pp. 19–36). Lanham: Lexington Books.

Manders-Huits, N. (2006). Moral responsibility and IT for human enhancement. *Proceedings of the 2006 ACM Symposium on Applied Computing*, 267-271.

Marx, K. (1990). *Capital*. (trans: Fowkes B.). London: Penguin Books.

McDowell, J. (1998). *Mind, value and reality.* Cambridge: Harvard University Press.

Mencius. (1970). *Mencius*. (trans: Lau D. C.). London: Penguin Books.

Miller, J. J. (2004). Squaring the circle: teaching philosophical ethics in the military. *Journal of Military Ethics, 3*, 199–215.

OECD. (2001). Competencies for the knowledge economy. In *Education Policy Analysis*. OECD, Paris. http://www.oecd.org/innovation/research/1842070.pdf. Accessed 18 November 2013.

Ophir, E., Nass, C., & Wagner, A. (2009). Cognitive control in media multi-taskers. *Proceedings of the National Academy of Sciences, 106*(37), 15583–15587.

Parks, J. A. (2010). Lifting the burden of women's care work: should robots replace the "human touch"? *Hypatia, 25*(1), 100–120.

Rideout, V. J., Foehr, U. G., & Roberts, D. F. (2010). Generation M2: media in the lives of 8 to 18-year olds. Henry J. Kaiser Family Foundation. http://files.eric.ed.gov/fulltext/ED527859.pdf. Accessed 21 Nov 2013.

Rinard, R. G. (1996). Technology, deskilling and nurses: the impact of the technologically changing environment. *Advances in Nursing Science, 18*(4), 60–69.

Ritzer, G. (2008). *The McDonaldization of society*. Los Angeles: Pine Forge Press.

Roberts, J. (2010). Art after deskilling. *Historical Materialism, 18*, 77–96.

Sharkey, N. (2010). Saying 'no!' to lethal autonomous targeting. *Journal of Military Ethics, 9*(4), 369–383.

Sharkey, N., & Sharkey, A. (2012). Granny and the robots: ethical issues in robot care for the elderly. *Ethics and Information Technology, 14*(1), 27–40.

Sparrow, R. (2007). Killer robots. *Journal of Applied Philosophy, 24*(1), 62–77.

Sparrow, R. (2009). Predators or plowshares? Arms control of robotic weapons. *IEEE Technology and Society, 28*(1), 25–29.

Sparrow, R., & Sparrow, L. (2006). In the hands of machines? The future of aged care. *Mind and Machine, 16*, 141–161.

Spencer, D.A. (2000). Braverman and the contribution of labour process analysis to the critique of capitalist production: twenty-five years on. *Work, Employment and Society, 14*(2), 223–243.

Townsend, K., & Charles, M. B. (2008). *Jarhead* and deskilling in the military: potential implications for the Australian labour market. *Australian Bulletin of Labour, 34*(1), 64–78.

Tronto, J. (1993). *Moral boundaries: a political argument for an ethic of care*. New York: Routledge.

Vallor, S. (2010). Social networking technology and the virtues. *Ethics and Information Technology, 12*(2), 157–170.

Vallor, S. (2011). Carebots and caregivers: sustaining the ethical ideal of care in the 21st century. *Philosophy and Technology, 24*(3), 251–268.

Vallor, S. (2013). The future of military virtue: Autonomous systems and the moral deskilling of the military. In K. Podins, J. Stinissen, & M. Maybaum (Eds.), *2013 5th International Conference on Cyber Conflict* (pp. 471–486). Tallinn: NATO CCD COE Publications.

Van Wynsberghe, A. (2013). Designing robots for care: care centered value-sensitive design. *Science and Engineering Ethics, 19*, 407–433.

Verbeek, P. P. (2011). *Moralizing technology: understanding and designing the morality of things*. Chicago: University of Chicago Press.

Wang, Z., & Tchernev, J. M. (2012). The myth of media multitasking: reciprocal dynamics of media multitasking, personal needs, and gratifications. *Journal of Communication, 62*(3), 493–513.

Whitby, B. (1996). *Reflections on artificial intelligence: the legal, moral and ethical dimensions*. Oxford: Intellect Books.

Whitby, B. (2008). Computing machinery and morality. *AI & Society, 22*, 551–563.

Wood, S. (1987). The deskilling debate, new technology and work organization. *Acta Sociologica, 30*(1), 3–24.