Screening War: Moral skills and mediation theory

Edward Windle

December 2024

1 Introduction

There is a long tradition assessing technology's effects on human skills (Braverman 1974; Marx 1887; Smith 1976). Technological advancements in modes of production has cultured effects on wealth, employment, consumption, and skilled labour. By automating tasks, the skills necessary to complete said tasks are lost. Analogously, Vallor's (2015) thesis maintains that information and communication technologies (ICTs) morally deskill by automating tasks required to cultivate moral skills: learned and habituated abilities to identify morally salient features of a scenario and determine the correct means to deal with those features. In this essay, I broaden this thesis. Specifically, I propose that, by acknowledging technological mediation of human perception of reality as also affecting moral skills, moral deskilling produces more nuanced assessments of the moral effects of technology. Mediation theory will be defined as a descriptive and normative theory of human-technology-world relations, wherein technologies always mediate the relationships humans have to the world. The descriptive part is a continuation of the postphenomenological tradition; the normative part examines both how technologies currently mediate ethical life and how they should mediate ethical life. Moral deskilling (or upskilling) is the negative (or positive) relation between a loss (or increase) of the conditions required to engender moral skills. Deskilling[1] will be assessed from a virtue-ethical approach, such that moral skills are a prerequisite to the cultivation of a virtuous character.

Firstly, §2 will unpack the virtue-ethical foundation and moral deskilling thesis. I present Vallor's dialectic and expose the assumption that automation is exclusively necessary and sufficient to morally deskill. §3 will make clear that Vallor's thesis is narrow and lacks the necessary descriptive apparatus to produce nuanced assessments of *how* moral deskilling occurs, specifically in relation to semi-autonomous technologies. I thus propose a broader account of moral deskilling in §4. By establishing hermeneutic mediation as another means to morally deskill individuals, I augment Vallor's analysis with its descriptive toolkit. Essentially, I make clear that automation is a necessary, but not always sufficient, effector of moral deskilling

[1] Hereafter, deskilling will be the primary focus

2 Moral Skills and Virtue Ethics

To get good at something, one needs to develop the necessary skills relevant to that task. To gain skills, one has to practice a task: 'the things we have to learn before we can do [a craft], we learn by doing [the craft]' (Aristotle & Ross 1963, 1103a). Essentially, we learn how to do something by doing it. For example, the skill of, say, typing is cultivated by repeatedly typing; 'good' typing is cultivated by repeatedly typing while following principles, such as resting one's index fingers on the 'F' and 'J' keys, only using certain fingers for certain keys, etc.. Deskilling suggests that the skill of typing, facilitated by technological advancement, has likely produced a devaluation of another skill: writing—i.e., fewer people use longhand; even fewer use longhand well. Equally, however, because shorthand is more efficient than longhand, this devaluation freed people's time to spend on other tasks, such as planning what they write. The deskilling thesis maintains that a skill's vulnerability is characterised by the destabilisation of its value in relation to *automation*. As typewriters automated the process of writing letters to the press of a button, and computers facilitated more efficient automation, the actual process of handwriting letters is devalued. This is the standard deskilling thesis, whereby the 'labor process is to be rendered independent of the craft, tradition, and the works knowledge', which technological change, coupled with the actions of management, rationalise for efficiency by automating it (Braverman 1974, p.113).

Applying this conception of deskilling to a virtue-ethical theory implies the same vulnerability for *moral skills*. Specifically, moral deskilling can occur by undermining (some of) the conditions conducive to the cultivation of a moral skill. Moral skills are cultivated abilities to (F) recognise the 'morally salient features' of a situation, and (M) know both what means to use and why one is using them in said situation. One acquires moral skills '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' (Vallor (2015, p.109); emphasis added). So, if a technology negatively affects the right conditions for cultivating moral skills or makes insufficient the opportunity for repetition, then the potential for moral deskilling is introduced. For the conditions to cultivate moral skills have been reduced. To have been *reduced* is a necessary condition for moral deskilling, insofar as morally salient features must have previously existed to facilitate the acquisition of skills. (Analogously, it would be impossible to become a builder without materials to build with!)

For clarity, let's recharacterize this definition: if, for a right practice, both (i) right conditions and (ii) sufficient opportunity for repetition are fulfilled, then one can (iii) engender moral skills that are necessary (but not sufficient) for cultivating virtue. Firstly, it is necessary that the conditions concern the *right practice* – 'practices that successfully engender certain skills of acting rightly in particular moral contexts' (*ibid.*). Virtues, as understood in the Aristotelian sense, are stable and lasting (i.e., not temporary) dispositions to excel in specific contexts of action (e.g., courageously, justly). Dispositions are the product of

habituated learning to act, think, and respond (e.g., courageously, justly) to the morally salient features of particular scenario. A virtuous disposition is also active and persistent, insofar as one is disposed to act in certain ways despite challenges and difficulties (Annas 2011, p.8). We can thus define moral virtue as being a stable, persistent and active state of character; when a person has acquired a moral virtue, she is disposed to excel in moral matters. Importantly, moral virtue is considered intrinsically valuable: therefore, virtue ethics makes the normative claim that 'one should promote the means to cultivate virtue'. The upshot of this is the following normative claim: 'one should promote the conditions to cultivate moral skill'.

Attaining the moral skills necessary for virtue requires (i) and (ii) to be fulfilled. Fulfilling (i) is twofold. Firstly, conditions are facilitated by a virtuous model – i.e., someone who is already virtuous. So, the initial 'right conditions' to be courageous require a model who is already courageous[2]: they can, for example, identify all the morally salient features and explain why they are morally salient and how to deal with them. There is more to (i) than merely copying this model. For, secondly, (i) requires also the '*drive to aspire*' – the active drive to learn discerningly, and to understand why and how one should improve upon themselves and their model (Annas 2011, p.17). If a person is what I term a moral copyist, then she lacks the drive to aspire – i.e., a rote copier of moral acts. As such, (i) is satisfied by a virtuous model and the drive to aspire. Nevertheless, fulfilling all conditions for (i) will not cultivate moral skill without being habituated through sufficient repetition (ii). If this were not a necessary condition for attaining moral virtue, then one lacks the repeated practice that satisfies what I term the knowledge factor: this engenders specific skills necessary for attaining virtue. By the knowledge factor, I mean an epistemological dimension of a task that can only be acquired by doing that task[3]. Therefore, (ii) is characterised by the potential attainment of the knowledge factor by sufficient opportunity for repetition and practice. The concomitant and *continued* engagement of the antecedents (i) and (ii) eventually engender the consequence, (iii), which then entails the cultivation of (F) and (M) that constitute a moral skill.

Clearly, (i) and (ii) are analogous to the aforementioned process of acquiring skills. This is not to claim that moral skill is the same as a 'skill of craft associated with material production', but that moral skills are a distinctive type of skill. Once a moral skill has been established through (i) and (ii), it provides a 'scaffold or stable grafting site upon which virtue can (but may or may not) take hold' (Vallor 2015, p.110). By taking our fleshed out characterisation of moral skill and applying it to the virtue of courage, we can see each of these factors at work. Initially, a person satisfies the drive to aspire by active engagement and learning from someone who already embodies the virtue of courage (i). Then, through repeated practice that forms the habituated disposition (ii), she

[2]Such presupposes the association of extant and specific socio-cultural context

[3]I presume here the correctness of the Rylean argument that 'knowing-that' (i.e., knowing that something is that case) presupposes 'knowing-how' (i.e., knowing how to do things) Ryle (1945).

begins to engender the moral skills necessary for courage (iii). By doing this continuously, she learns how to excel at responding appropriately to fear, insofar as discerning the right features of a situation to fear or brave (F), and knows the right means to display courage in a dangerous situation and why to use those means (M). Though not sufficient for the virtue of courage, one still has to act courageous to eventually attain courage.

Importantly, Vallor maintains that the way contexts are affected is by automating the tasks that allow for participation in these processes, hence the ability to satisfy (i) and (ii). So, if all decisionmaking processes could be automated by technologies, humans would be morally deskilled, left diminished: 'utterly helpless to act justly and compassionately without [technology's] assistance' (2015, p.113). To relate deskilling in the traditional sense to this conception of moral skills, a critical implication can be highlighted: that 'we may have reason to regret the loss of the professional context for cultivating the given skills because we think the skills themselves are intrinsically valuable' (Vallor 2013, p.477). Switching *professional context* for the context outlined in claims (i) and (ii) entails this implication for moral skills also. To lose moral skill is a regrettable outcome; just as the loss of traditional manufacturing methods that require artisanal skills is regrettable. The negativity behind the outcome is that it detracts from the value of a product – a value (partly) derived from the arduous journey to cultivate a skill that enables its production[4]. Analogously, impeding the conditions that facilitate the satisfaction of either (i) or (ii) can be characterised as moral deskilling. The negativity behind the outcome of moral deskilling is due to the deficient development of an intrinsically valuable virtue, necessarily produced by moral skill.

So, according to Vallor, moral skills act as a necessary (but not sufficient) part of cultivating virtues. Moral skills themselves are produced by the concomitant enaction of (i) and (ii) to the right practice that then (iii) engender the specific conditions to develop the means–(F) and (M)–that constitute moral skill, which is necessary (but not sufficient) for cultivating virtue. (i) is satisfied by following a virtuous model and maintaining a drive to aspire; (ii) is satisfied by repetition to form habituation, concomitantly evolving the knowledge factor. Moral deskilling is caused by the automation of practices¸ resulting in either a dearth of (F) or insufficient opportunity to carry out (i) and (ii). Therefore, (iii) cannot be engendered and the specific means and situations necessary for cultivating virtue are not achieved. Let us now see Vallor's assessment in action.

3 A Narrow Scope

The only cause of moral deskilling Vallor addresses is automation. But is this always sufficient? Consider the following case, expanding upon a scenario acknowledged by Vallor:

The Drone Operator

[4]Consider here contemporary yearning for 'artisanal' and 'craft' products.

> Sylvia is a sensor operator for an MQ-1 Predator drone. The role of a sensor operator is to control the sensors and the weapons. 400 miles away from Sylvia, the drone circles 25,000ft above her target: an opposition's military installation. Sylvia trains the camera directly on a building confirmed to house a high-value target. Individuals are identifiable; their movements and routines have been tracked over a long period of time. After receiving the go-ahead from a chain-of-command, Sylvia presses a button: a missile is fired. With (literal) laser precision, the missile strikes the building. The screen turns white; the scene emerges and a pile of rubble is displayed. Confirming that the strike is successful, operators pilot the drone back to base (most of the flying is, of course, carried out on autopilot).

Firstly, drones are *semi-autonomous*: i.e., they can carry out tasks 'with some degree of independence but require human oversight and do not target or fire autonomously' (Vallor 2016, p.212). They facilitate, importantly, killing. Right military conduct of killing in war is guided by just war theory's principle of *jus in bello* (see Walzer 2006). *Jus in bello* can be conceived as a form of military virtue, ratifying both (F) and (M) in relation to force: it 'must be applied only at the right times, towards the right persons, in the right places, and in the right manner' (Vallor 2015, p.114). Drones are clearly a method of applying force, so incorporating the assessment of moral deskilling should determine how drones affect the necessary conditions for engendering moral skills with regards to cultivating (F) and (M) for military virtue

Following Coeckelbergh (2013), Vallor identifies drones as exampling the ambiguous nature of technology. The dichotomic nature of ICTs can act to facilitate new kinds of moral skills or knowledge through providing new perspectives or depriving an agent of other types of moral skills or knowledge. Vallor would argue that Sylvia is deprived of the moral skills cultivated through practising military virtues acquired by lived experience in the battlefield (e.g., bravery and courage). Essentially, she does not have to enter the field of war because the physical tasks she would carry out there are automated by the drone, and thus cannot engage (i) or (ii). However, she is still involved in the processes: experiencing reality through a novel perspective, viewing the lives, routines, and habits of her targets. This arguably provides a richer moral context concerning the target, enhancing Sylvia's moral knowledge of the context of the action, which can cause 'greater psychological and moral trauma' because the automation provides a detailed perspective but 'no other mode of practical engagement with their targets than passively observing or killing them' (2015, p.116). This assessment is *prima facie* accounted for by the means of automation. The ambiguous influence of certain ICTs and their automation of tasks can incur other consequences for morality, such as the moral-epistemological dimension or moral deskilling shown here.

However, analysing exclusively the relation between automation and deskilling restricts nuanced analysis regarding *how* Sylvia's interactions this technology may perpetuate moral deskilling. Let me clarify: by automating the tasks

and removing Sylvia from the environment, neither (i) nor (ii) can be satisfied because the prerequisite morally salient features are no longer part of her environment. She cannot culture the moral skills necessary for, say, courage because she is located 400 miles away from the morally salient features (i.e., objects that incite a rational fear), such that she cannot learn and practice the means that allow her to know how to deal with fear, when and why to be selfless and face fear, etc.. She is, explicitly, deskilled by automation. But the task has only been semi-automated: she still controls the weapons, sensors, etc. She is afforded only a passive mode of killing. She presses the button and hence bears the responsibility of killing. Here, Vallor's descriptive powers wane. For moral skills account only for the loss of a limited set of products resulting from interactions with technologies, not, however, the loss of other skills (potentially unrelated to that task) through interacted with the new technology. Surely, then, we can delve deeper into the unique deskilling that arises from her continued interaction with the reality represented by the technology, or even from a person's relation both with the technology and what the technology does. So, Vallor's analysis seems to leave open the question: 'how can *using* a semi-autonomous technology *continue* to morally deskill?' What I propose, then, is an analysis of the other half of the equation not covered by Vallor's *narrow thesis*, namely the human relation to the technology. I aim to do this with the toolkit of mediation theory. This is not to say that automation is not a catalyst or necessary condition, but that mediation can act to deskill further.

I thus maintain that Vallor's account is too narrow. Of course, a proponent of this narrow thesis could try account for the proposed question by considering this further loss of skill as reducible to automation: that, for instance, the conditions for other (previously unaffected) skills are being recontextualised by automation. This, however, misses the core argument of my position: moral skills are made vulnerable by automation, which in itself can act to morally deskill (shown above); however, assessment of moral skills with automation does not answer questions concerning ways in which human interactions with technologies can further deskill. Moral deskilling can benefit from investigating how people interact with the technologies that mediate the user's experience of the conditions necessary for the cultivation of moral skills. Essentially, I aim to aggregate this analysis upon the narrow thesis, thus assessing more deeply the relation between performing the same task but with different or limited modes. At the core of these points are moral skills: we perform morally relevant tasks and technologies can change how we complete these tasks.

Before exploring this development, it is worthwhile noting that Vallor suggests a movement away from conceiving of technologies as neutral and instrumental (i.e., only automating tasks). She suggests that we should consider the ontological claim that humans are technological, that technologies 'merge with, mediate and continually transform our own agency and capacities (2015, p.116). Essentially, we are *technomoral creatures* – what we produce shapes us and our morality. This is moving in the right direction. However, this characterisation misses an important point. For there is another relation between humans and the world, namely between humans and the social. Humans are

social beings–Beings-with (see *Mitsein* in Heidegger 1962)–and we can make a normative claim in relation to moral deskilling: 'technologies should make us engage more with other humans in reality'. By emphasising how human relations to technologies actively preface and contextualise the ways we engage with the world, we can acknowledge also how they actively influence how we interact with other people. As with our drone operator, her interactions with the humans she is killing are entirely prefaced and contextualised *through* technology, where the relation between her and reality is mediated by a screen.

4 Screening War

To reiterate, mediation theory analyses technologies in terms of their role as mediating the relations between human and reality. It maintains that technologies always instantiate a relation between humans and reality when used. As such, a technology can, for example, introduce novel perspectives, lastingly alter how we interpret the world, and alter social orders or norms[5]. There is a descriptive and normative aspect here. The former consists of two dimensions: a pragmatic dimension, concerning technologies designed to influence action, and a hermeneutic dimension, concerning how technologies can shape interpretations of reality. The hermeneutic part is more relevant to our purposes here. The normative part uses the descriptive part to assess how technologies already mediate and should mediate life, which is exampled in the application of mediation to moral deskilling.

Don Ihde (1990; 1998) explores the hermeneutic dimension by analysing how technologies provide a *representation* of the world; one experiences reality by means of a technology that produces a relation between humans and reality. This relation is structured such that technologies make present a representation of reality that is not transparent. By transparent, I mean that reality is perceived by means of (not through) technology. Verbeek (2008*b*) illustrates hermeneutic mediation with the technology of ultrasound. Ultrasound acts as a *functional* means to make visible, for example, an unborn foetus. As such, this shapes how the unborn foetus is perceived. Firstly, this has an *epistemological dimension*, insofar as the representation, though markedly different from reality, constructs a mind-dependent appearance that allows inference to certain qualities about the reality being represented, thus producing certain beliefs about the reality. Secondly, the visual relation provides an *ontological reality* of the object being represented – a strengthening of the belief about the *existence* of the object. In this way, technologies can translate reality, visually introducing epistemological and ontological dimensions to a reality otherwise (mostly) imperceptible.

Because hermeneutic mediation forms a relation between humans and reality by means of technology, this introduces the opportunity for *relation issues*. If, for example, one cannot interpret correctly the ultrasound, then one may make

[5]Notably, this can be in the sense of social epistemology or ethnomethodology.

an incorrect judgement about the unborn child or simply be unable to distinguish the child in the representation. Moreover, though some technologies will provide a higher degree of fidelity, the mere fact that it is a representation of reality means that at least some percepts are lost. In this sense, we can say that representations are 'semi-opaque', such that different technologies that enable hermeneutic mediation can lose or misrepresent something a reality.

So, how can we account for moral deskilling by means of hermeneutic mediation? I will propose three ways, though application is not confined to these. (1) Firstly, a technology's semi-opaque representation can affect the right conditions (i) through relation issues in conveying the salience of moral features that are a prerequisite for the cultivation of moral skills: without morally salient features, the cultivation of a skill is *nearly* impossible[6]. A human's sensory modalities are contingent upon the reality they experience; if the reality of an object is, for example, represented by a dot on a sonar screen, there is little hope of constructing a morally significant experience of that dot. Maybe the dot is identified as a submarine—from that one could build a mental representation of that reality, but this may still provide insufficient information to construct a morally significant structure of the environment. Even though the ontological reality of the object is introduced, one's experience of that reality is wholly constituted by the technology. In turn, one's 'drive to aspire' may be muted because of a lack of moral features to stimulate its active engagement with the world. One may even make morally correct decisions concerning the dot, yet lacking the drive to aspire renders them a mere moral copyist. They are thus deskilled by their interactions with a representation contextualised through a technology because it reduces the moral context compared to an unmediated perception of the object.

What about the individually-relative potential that represented features are morally (in)significant? This requires a slight detour to technological intentions: a technology's intentions are the action-guiding role they play in shaping actions and experiences of humans. These can be conceived as *derivable intentions* (Searle 1983, cited in Verbeek 2008*a*, p.95)—i.e., entirely reducible down to human intentions. Firstly, there are *designed intentions*, which is the way the technology was designed to be used. Secondly, there are *projected intentions*, which is the projection of use case on the material makeup of the technology. Both can produce unforeseen mediations that can act to morally deskill. Notably, as long as someone is interacting with the technology (e.g., semi-autonomous technologies), then both intentions will be active.

In turn, (2) technologies can morally deskill through their designed intentions by, for instance, a the intended mediation emphasising parsimony of details, which thus removes morally salient features. Equally, a technology may be designed to limit what courses of action one may take, which can habituate (ii) regarding incorrect means (M), or limit engagement (i) with morally salient

[6]Recall Aristotle's conception of skills being learned through performing the appropriate task.

features (F). Their unintended mediations can also morally deskill. For example, even though mobile phones were designed for ease of social interactions, the mediated experience has resulted in decreased social interaction. This clearly goes against the fundamental human condition as Beings-with and could (in a wider sense) restrict the cultivation of the necessary social skills to even engage one's virtuous model; hence, causing widespread deskilling concerning (i).

Thirdly, projected intentions can morally deskill by (3) producing unexpected mediations because the user projects both other contextual information or sense-datum and their intellectual, moral, and epistemic history onto the mediation with reality. So, for example, a sonar dot that only represents a submarine can be engaged with as if it were morally significant. Similarly, if one has significant experience with a simulation game that accurately represents a fictional reality on a screen and eventually joins the profession that the game simulates, then one may not be actively sensitive to the relevant features because they have habituated a disposition to fictitious features that translate to reality. This results in an inability to properly cultivate (i) and (ii): their *knowledge factor* will be inadequately satisfied by a fictional object[7] (ii) and the fictional object may reduce the salience of the moral features when faced with the same representation but of a real object (i). Additionally, this projection results from the user, and thus carries the potential of personal differences concerning the extent of deskilling or which way one is deskilled.

To see this analysis of mediated experience and moral skills, let us return to the drone operator, Sylvia. It is uncontroversial that the targets represented on the screen can be morally salient. We can address why they persist by considering the degree of *relation issue*. Certain qualities–humanness, visual acuity, colours, etc.—all contribute to building a salient experience. However, if some of the features are not adequately translated, then Sylvia may be unable to discern morally salient features (F). As such, (1) is in action: she cannot engage sufficiently discern the morally salient feature and thus cannot cultivate all the relevant conditions necessary to produce the moral skill previously available to her.

What's more, we can expound on Vallor's point concerning harm through richer moral context. Firstly, we can say that the greater psychological harm is caused by the fidelity and novel perspective enabled by the system's representation of an ontological reality experienced and acted upon by Sylvia. Clearly, Vallor's analysis applies here, but is augmented by the descriptive part of mediation. Equally, (2) the designed intention of a drone (precision killing without risking human life) limits the modes of interaction with the represented reality. To further answer how this can affect moral skills, let us look at the projected intentions. For instance, (3) Sylvia can project her intellectual history (contextual information provided by a mission brief; prior experience) onto her interaction with the mediated reality; she can, in some sense, carry out a limited practice of (M) compared to the (M) acquired in lived-bodily combat. The practice of (M) is limited, for example, by oversensitivity or desensitisation to (F). Over-

[7] Unless reskilling occurred, but this is outside our present scope.

sensitivity may mean Sylvia has a rash and emotionally-charged moral deliberation concerning her action (hence, Vallor's psychological and moral harm). Her environment facilitates this more than a (personal) life-or-death combat situation – one that may act to culture moral skills concerning military virtue. Additionally, desensitization, through habituation with simulation, prior conditioning, or fictional objects, could limit her ability to properly conceive of the representations. For the moral salience of the reality they represent may be permanently dulled by habituation in the right practice with the wrong features (i.e., highly-accurate simulation). In either case, she will likely act as a moral copyist compared to a morally skilled person, following orders without the personal skill to individuate or evaluate the morality of her own actions.

Though brief, it is clear from this evaluation that the relation between skills and technology does not have to be affected exclusively by automation. Firstly, automation fails to capture the social dimensions rapt up in a technology's implementation. Moral deskilling certainly helps assess this, but fails to account for the moral nuance regarding the interactive relationship of an agent with a technology. It is beyond the scope here to evaluate the balancing of these two theses, except to say that mediation affects moral skills also—that automation can be considered necessary but not always sufficient to analyse a technology's affects on morals skills.

5 Conclusion

In this essay, I have built upon a comprehensive explication of both moral deskilling and hermeneutic mediation to account for further exploration concerning the relation between human moral skills and technology. I have developed a methodology, deepening the potential analysis to be parsed from moral deskilling. In detailing the relation between semi-automated technologies, human interaction, and technology in light of mediation, I maintain that automation is not always sufficient for analysis of moral deskilling. Such is highlighted by the more descriptive and nuanced application of this relation to the drone case.

Much more can be said about the ways in which technology may mediate and affect moral deskilling, especially in relation to, for instance, hermeneutical injustice. The aim here has been to propose a novel perspective concerning technology's effects on moral skills. Concerning this, mediation theory proposes a compelling toolkit that help disentangle the complex nature of technologies in themselves and understand our interactions with them, specifically in relation to moral skills.

References

Annas, J. (2011), *Intelligent Virtue*, Oxford University Press, Oxford.

Aristotle & Ross, D. (1963), *The Nicomachean Ethics of Aristotle*, Oxford University Press, Oxford.

Braverman, H. (1974), *Labor and Monopoly Capital: The Degradation of Work in the Twentieth Century*, Monthly Review Press, New York.

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

Heidegger, M. (1962), *Being and Time*, Blackwell Publishing, Oxford.

Ihde, D. (1990), *Technology and the Lifeworld: From Garden to Earth*, Indiana University Press, Indiana.

Ihde, D. (1998), *Expanding Hermeneutics: Visualism in Science*, Northwestern University Press, Evanston.

Marx, K. (1887), *Das Kapital*, Progress Publisher, Moscow. [Originally published in 1867].

Ryle, G. (1945), 'Knowing how and knowing that: The presidential address', *Proceedings of the Aristotelian Society* **46**, 1–16.

Searle, J. R. (1983), *Intentionality: An Essay in the Philosophy of Mind*, Cambridge University Press, Cambridge.

Smith, A. (1976), *An Inquiry into the Nature and Causes of the Wealth of Nations*, Clarendon Press, Oxford. [Originally published in 1776].

Vallor, S. (2013), The future of military virtue: Autonomous systems and the moral deskilling of the military, *in* '2013 5th International Conference on Cyber Conflict (CyCon 2013): Proceedings', pp. 471–486.

Vallor, S. (2015), 'Moral deskilling and upskilling in a new machine age', *Philosophy and Technology* pp. 107–124.

Vallor, S. (2016), *Technology and the Virtues: A Philosophical Guide to a Future Worth Wanting*, Oxford University Press, Oxford.

Verbeek, P.-P. (2008*a*), 'Cyborg intentionality: Rethinking the phenomenology of human-technology relations', *Phenomenology and the Cognitive Sciences* pp. 387–395.

Verbeek, P.-P. (2008*b*), Morality in design: Design ethics and the morality of technological artifacts, *in* P. Kroes, P. E. Vermaas, A. Light & S. A. Moore, eds, 'Philosophy and Design: From Engineering to Architecture', Springer, Dordrecht, pp. 91–103.