



Neuromedia, cognitive offloading, and intellectual perseverance

Cody Turner¹ 

Received: 21 March 2021 / Accepted: 24 October 2021

© The Author(s) 2022

Abstract

This paper engages in what might be called anticipatory virtue epistemology, as it anticipates some virtue epistemological risks related to a near-future version of brain-computer interface technology that Michael Lynch (2014) calls 'neuromedia.' I analyze how neuromedia is poised to negatively affect the intellectual character of agents, focusing specifically on the virtue of intellectual perseverance, which involves a disposition to mentally persist in the face of challenges towards the realization of one's intellectual goals. First, I present and motivate what I call 'the cognitive offloading argument', which holds that excessive cognitive offloading of the sort incentivized by a device like neuromedia threatens to undermine intellectual virtue development from the standpoint of the theory of virtue responsibilism. Then, I examine the cognitive offloading argument as it applies to the virtue of intellectual perseverance, arguing that neuromedia may increase cognitive efficiency at the cost of intellectual perseverance. If used in an epistemically responsible manner, however, cognitive offloading devices may not undermine intellectual perseverance but instead allow people to persevere with respect to intellectual goals that they find more valuable by freeing them from different kinds of menial intellectual labor.

Keywords Neuromedia · Virtue epistemology · Intellectual perseverance · Cognitive offloading · Virtue responsibilism

1 The virtue epistemology of technology

The virtue epistemology of technology is a burgeoning philosophical subfield that examines the relationship between the use of cognitive artifacts and the cultivation and maintenance of intellectual virtues. A cognitive artifact is any technology used to help support cognitive processes. Intellectual virtues are character traits that make

✉ Cody Turner
cody.turner@uconn.edu

¹ Philosophy Department, University of Connecticut, Storrs, CT, USA

one an excellent person from an epistemic point of view or contribute to making one a superb thinker (e.g., open-mindedness, intellectual autonomy, intellectual humility, and intellectual perseverance).¹ Virtue epistemology examines the nature of intellectual virtue as well as how the virtues relate to one another and other epistemic elements like knowledge and justified belief (Battaly, 2008).² Unlike traditional analytic theorists in epistemology (e.g., Chisholm, 1989; Nozick, 1981; Goldman, 1967), virtue epistemologists treat intellectual virtues as foundational epistemic elements and hold that knowledge and justified belief can be understood in terms of the virtues and not vice versa.

The virtue epistemology of technology analyzes both (a) how different cognitive artifacts shape (or are prone to shape) the cultivation and possession of intellectual virtues and (b) how the cultivation and possession of different intellectual virtues shape (or are prone to shape) the manner in which agents approach and navigate their use of cognitive artifacts. Up until now, the majority of scholars working in this area have focused on contemporary cognitive artifacts like pharmaceutical ‘smart’ drugs (Carter, 2020; Ranisch, 2015), GPS navigation systems (Gillett & Heersmink, 2019), and more than anything else, the internet and online search engines (Sunstein, 2006, Greenfield, 2014, Heersmink, 2016, Smart 2018a, Schwengerer, 2021). Richard Heersmink (2018), for example, demonstrates how nine different intellectual virtues ought to be deployed within the context of online search engines. Regarding the virtue of intellectual carefulness, he says, “An intellectually careful person will avoid common mistakes when using Google. Such mistakes may include hastiness such as only clicking on the first-ranked page and not reading other pages as to be able to compare various sources of information” (Heersmink, 2018: p. 7). While Heersmink focuses on how being intellectually virtuous can help agents navigate the internet in an epistemically responsible fashion, others concentrate on the negative effects that the internet is already having on intellectual virtue cultivation. Nicholas Carr (2010) contends that the internet is turning us into shallow thinkers by diminishing our ability to concentrate, practice deep reading, and engage in appropriate reflection. Others have started analyzing how internet search engines like Google and Youtube can produce bad epistemic outcomes via algorithmic filtering mechanisms (Alfano et al. 2021) and autocompleted web search (Miller & Record, 2017). For instance, Simpson (2012) and Nguyen (2020) assess how ‘personalization’ algorithms undercut informational objectivity and stifle the virtues of open-mindedness and intellectual humility by facilitating confirmation bias and leading to the creation of online echo chambers and information bubbles.

While this is all vital work, it is also worth investigating the epistemological implications of emerging cognitive artifacts that are not yet pervasive in society but stand

¹ These traits can be contrasted with intellectual vices, which are traits that reflect poorly upon an agent in an epistemic sense or contribute to making one a lackluster thinker (e.g. close-mindedness, intellectual servility, intellectual arrogance, intellectual cowardice, and intellectual laziness).

² The origins of contemporary virtue epistemology can be traced back to Ernest Sosa’s work in the early 1980s (Sosa 1980). Sosa appeals to the concept of intellectual virtue when articulating his theory of ‘virtue perspectivism,’ which he contends can resolve the longstanding debate in epistemology between coherentism and foundationalism. The field of virtue epistemology is structurally analogous to the field of virtue ethics. Just as virtue ethicists ground the morality of actions in evaluations of moral virtues, so do virtue epistemologists ground the epistemic status of beliefs in evaluations of intellectual virtues.

a good chance of becoming ubiquitous. Accurately anticipating the epistemological problems posed by emerging technologies before these problems fully manifest themselves gives us a chance to formulate regulatory and design-based solutions in a timely manner. For example, Regina Rini (2020) discusses the impending social and political epistemic threat that machine learning technology poses as it pertains to ‘Deepfakes.’ Rini observes that visual and audio recordings currently function in society as what she calls ‘epistemic backdrops,’ meaning that they serve to regulate testimonial norms of sincerity and competence (as well as other forms of justificatory evidence like photographs). She then claims that the emergence of Deepfake technology may cause “a sudden collapse of the backstop to our testimonial practice, as we can no longer trust that recordings provide authoritative correction” (Rini, 2020: 13). This work of Rini’s is an example of what might be called *anticipatory social epistemology* because she is anticipating the *social* epistemological risks associated with an emerging cognitive artifact.

This paper engages in what might be called *anticipatory virtue epistemology*, as it anticipates some *virtue* epistemological risks related to an emerging cognitive artifact. The technology explored here is not machine learning systems and Deepfakes but rather a near-future version of brain-computer interface technology called *neuromedia* (Lynch, 2014, 2016, Carter, 2020, Pritchard 2018b). I analyze how neuromedia is poised to negatively affect the intellectual character of agents, focusing specifically on the virtue of intellectual perseverance, which involves a disposition to mentally persist in the face of challenges towards the realization of one’s epistemic ends (Battaly, 2017; King, 2014). There are two central aims of the paper. First, I present and motivate what I call ‘[the cognitive offloading argument](#)’, which holds that offloading cognitive tasks to technologies has the capacity to undermine intellectual virtue cultivation by rendering agents *cognitively passive*. I submit that high-tech cognitive enhancement devices like neuromedia incentivize excessive cognitive offloading, and so threaten to make agents especially cognitively passive. Then, I examine the cognitive offloading argument as it applies to the virtue of intellectual perseverance. This discussion illustrates that while neuromedia may increase cognitive efficiency, it threatens to do so at the cost of intellectual perseverance.

Anticipatory virtue epistemology of this sort might inspire various concrete policy proposals, such as government regulation requiring that producers of brain-enhancement devices disclose the relevant epistemic risks to consumers, just as medical risks for pharmaceutical drugs must be disclosed. Importantly, many contemporary cognitive artifacts like the smartphone and the internet are already functioning to erode the intellectual character of agents via excessive cognitive offloading. The worry expressed in this paper is not, therefore, that the emergence of neuromedia will bring with it a novel virtue epistemic problem involving intellectual perseverance, but rather that it will significantly worsen an already existing virtue epistemic concern.³ More than anything, it is the tightly coupled nature of neuromedia, or the fact that the

³ Analogously, Rini contends that Deepfakes are poised to significantly worsen an already existing social epistemic problem involving testimonial knowledge. Testimonial norms, according to Rini, are already being undermined via the widespread distribution of fake news on social media platforms. She argues that Deepfakes, as an especially pernicious form of fake news, will increase the severity of this social epistemic problem by facilitating a breakdown of the backstop of our testimonial norms.

technology is so deeply interwoven with our minds, that distinguishes it from contemporary cognitive artifacts and renders it especially concerning from the standpoint of virtue epistemology. As I explain, however, the advent of neuromedia does not condemn us to any particular epistemic fate. It is up to us to be aware of the virtue epistemic dangers associated with excessive cognitive offloading and to appropriately steer clear of these dangers. If used in an epistemically responsible manner, neuromedia may not undermine intellectual perseverance but instead allow us to persevere with respect to intellectual goals that we find more valuable by freeing us from different kinds of menial intellectual labor.

The structure of the rest of the paper proceeds as follows. Section II introduces brain-computer interface technology (or BCI for short), describes some recent and ongoing developments in the field of BCI research, and offers the concept of neuromedia as a cognitive enhancement BCI which could plausibly be created in the not-too-distant future. Section III presents and motivates the cognitive offloading argument, explains why the argument is especially concerning for a technology like neuromedia, and briefly considers an objection to the argument based on the extended mind thesis (Carter & Kallestrup, 2019; Clark & Chalmers, 1998; Pritchard, 2015). Section IV examines the cognitive offloading argument as it pertains to the virtue of intellectual perseverance before describing a few responsible forms of cognitive offloading that are compatible with the retention of intellectual perseverance. Section V contrasts cognitive offloading with what I call ‘cognitive inloading,’ maintaining that excessive cognitive inloading threatens to undermine intellectual perseverance by engendering a disposition of cognitive distraction (as opposed to a disposition of cognitive passivity). Section VI concludes.

2 Brain-Computer Interfaces and the Advent of Neuromedia

Brain-computer interfaces are hardware and software communication systems that directly connect the human nervous system to external computing devices via invasive or non-invasive electrodes. By establishing a direct connection to the human nervous system, BCIs allow subjects to execute technological actions using nothing but mental commands. There are three main functional components of BCIs: the electrodes, the signal processing unit, and the application device (Nam et al., 2018). The electrodes detect intentional changes in the brain patterns of the BCI user, a process that can be implemented in either an invasive or non-invasive manner. Non-invasive BCIs place the electrodes on the outer scalp, while invasive BCIs directly implant the electrodes under or over the brain’s cortical service.⁴ Once the electrodes analyze and extract the pertinent neural activity of the user, the signal processing unit then converts this neural activity into a command signal, which the application device subsequently carries out into a technological action.

⁴ Non-invasive BCIs typically use the techniques of electroencephalography (which measures the electrical activity of the brain) or magnetoencephalography (which measures magnetic fields generated by electrical currents in the brain), whereas invasive BCIs typically use a technique called electrocorticography (Abdulkader 2015).

BCIs have thus far been mainly developed for *therapeutic* purposes, most of which come in the form of motor and linguistic applications. Motor applications of BCIs like motorized wheelchairs and robotic arms help patients with central nervous system disabilities (e.g., Locked-in-Syndrome (LIS) and tetraplegia) to regain motor control.⁵ Linguistic applications of BCIs can also help these patients recover communicative abilities by enabling them to write sentences on a computer screen using nothing but their minds (Birbaumer et al., 2000).

While BCIs are great tools for different kinds of therapy, they can also *enhance* the cognitive abilities of non-disabled people. Many companies are already in the process of creating BCI enhancement devices for the consumer market. In 2007 the company NeuroSky introduced the first commercially available, noninvasive BCIs to the public. In 2017 Facebook announced that it is developing a non-invasive BCI device that would enable people to type at a fast rate (100 words per minute) using nothing but mental commands. Since making this announcement, Facebook acquired the startup company CTRL-labs, which has created a BCI wristband that lets users control computers by manipulating a virtual hand. In 2016 the billionaire entrepreneur Elon Musk launched the company Neuralink to create invasive BCIs which can be sold to consumers at scale. The devices under construction at Neuralink, called neural laces, include a chip that is placed outside of the head but connected to elastic threads containing electrodes implanted inside the skull. Other than neural laces (and neural implants more broadly), another promising avenue for developing BCI enhancement technologies is the prospect of *neuralnanorobots*, which are tiny robots that can record neuronal activity and operate on the brain at a molecular level. According to Martins et al. (2019), the field of neuralnanorobotics is poised to create the first Brain Cloud Interface, or BCI hooked up to the wireless cloud.

BCIs are so distinctive and appealing from a cognitive enhancement perspective because of how tightly coupled they allow our minds to be with cognitive artifacts. Cognitive systems are tightly coupled with artifacts to the extent that there is high-bandwidth information transfer between the artifacts and the mental processes they help support. Computing devices have gradually become more tightly coupled with mental processes as they have evolved throughout the years. The advent of smartphones in the mid-2000s enabled agents to achieve a kind of symbiosis with computers by granting them the ability to access the technology from anywhere in the world. However, contemporary smartphones still require that people use their hands (and often type onto a digital keyboard), and this kind of sensorimotor interaction functions to reduce the bandwidth of the system. To merge our minds more fully with information technologies, we need to create machines that allow for higher bandwidth or a much faster information transfer rate. BCIs promise to drastically improve the bandwidth problem by permitting users to execute technological actions (e.g. accessing the internet and retrieving data from the cloud) *via the power of thought alone*.

The emergence of BCIs like neural lace and neuralnanorobotics suggests that we may in the not-too-distant future create something approximating what the philosopher Michael Lynch (2014, 2016) calls *neuromedia*, which refers to a particularly

⁵ BCIs that target motor abilities are, strictly speaking, not cognitive artifacts. This paper focuses on BCIs which target cognitive capacities; specifically, cognitive enhancement BCIs.

sophisticated and cognitively integrated Brain Cloud Interface device. Lynch describes neuromedia thusly:

Imagine you had the functions of your smartphone miniaturized to a cellular level and accessible by your neural network. That is, imagine that you had the ability to download and upload information to the internet merely by thinking about it, and that you could access, similar to the way you access information by memory, all the various information you now normally store on your phone: photos you've taken, contact names and addresses, and of course all the information on Wikipedia (Lynch, 2014: p. 299).

The invention of neuromedia would be an unprecedented landmark in the history of technology, as it would permit humans to more or less fully merge their minds with information technologies. With a simple mental command, agents equipped with neuromedia can post onto social media platforms, perform complex calculations, navigate websites and online search engines, store personal information in the cloud, and retrieve information from the internet as easily as they retrieve information from biological memory. It may even be impossible to phenomenologically differentiate between technological actions generated by a neuromedia device and intellectual actions produced by the biological agent. Pritchard (2018) says that *phenomenological indistinguishability* is a defining feature of neuromedia: "By neuromedia I have in mind the development of information-processing technology that is so seamlessly integrated with our on-board cognitive processes that the subject is often unable to distinguish between her use of those on-board processes and the technology itself. The relationship of the subject to the technology is consequently no longer one of subject to instrument but rather 'feels' like a technological extension of her normal cognitive processes" (Pritchard, 2018: p. 328).

There are two key qualifications to make before proceeding. First, this paper is not an exercise in technological prophecy. Predicting the future of technology is a dangerous game. It is conceivable that neuromedia does not come to fruition in the way I have envisioned or that the advent of the technology is decades, if not centuries, away.⁶ Even if the advent of the technology is far-off, it is still worth considering the epistemic risks associated with neuromedia for roughly the same reasons that it is worth considering the ethical and existential risks related to AI superintelligence. We do not know for certain when and if we will create an AI superintelligence or whether we will merge ourselves with AI via brain-enhancement devices like neuromedia, but both of these technological prospects should be taken seriously given their potential to profoundly transform society and the individual human mind.

Second, this paper does not assume that people who equip themselves with neuromedia will automatically become better truth-trackers. At the dawn of the digital age, an idealistic vision of the internet predicted that the technology would produce wholly positive epistemic consequences for the world by promoting knowledge accumulation and the democratization of information. This idealistic vision failed to account for the fact that the internet can also be used for epistemically malicious purposes related to

⁶ One possibility is that the BCIs of the future will be noninvasive in nature instead of invasive. The neurotechnology company Kernel is focusing on manufacturing noninvasive BCIs that could be used for cognitive enhancement purposes.

technological manipulation, political extremism, conspiracy thinking, and the spread of fake news and computational propaganda. The emergence of neuromedia may face similar epistemic challenges and will be vulnerable to even more intrusive forms of technological manipulation related to brain-hacking (Ienca & Haselager, 2016). The cognitive offloading argument proclaims that even if these challenges are overcome, and neuromedia is used such that it reliably delivers truth, the technology will still be prone to undermine intellectual virtue insofar as it fosters a relationship of excessive epistemic dependence. I will now proceed to present and motivate the cognitive offloading argument. Then, in section IV, I examine the argument in the context of the virtue of intellectual perseverance, and the question of how neuromedia is poised to affect the cultivation and maintenance of intellectual perseverance.

3 The cognitive offloading argument

Cognitive offloading is a general intellectual strategy that involves outsourcing cognitive tasks to technological tools or other human agents. Many current technologies are used as cognitive offloading devices: we offload planning processes to our calendars, navigation capacities to our GPS systems, mathematical operations to our calculators, and memory storage to our computers and notebooks. Cognitive offloading poses no real risk to the cultivation of intellectual virtue *when done in moderation*. Moderate amounts of cognitive offloading may even be necessary for the development of some intellectual virtues (Carter, 2020). Cognitive offloading compromises intellectual virtue development only when it becomes excessive or when agents become overly reliant on artifacts for cognitive functioning. At least two factors determine whether dependence on a cognitive artifact qualifies as ‘excessive’: (i) how frequently the artifact is used and (ii) the range of cognitive tasks offloaded onto the artifact. There is no definitive point at which reliance on a cognitive artifact becomes ‘excessive.’ Rather, epistemic dependence can be conceptualized as existing on a bi-dimensional spectrum wherein such dependence becomes increasingly excessive as agents score higher with respect to variables (i) and (ii).

Typical use of a calculator, for example, does not involve excessive offloading because it is only used to perform a narrow cognitive task (i.e. execute mathematical operations) and is only invoked on occasion. By contrast, neuromedia is the ideal technology through which to examine the cognitive offloading argument because it is poised to foster a relationship of excessive epistemic dependence. As discussed, neuromedia is an all-purpose artifact that, due to its integration of the internet and artificial intelligence capabilities, can be deployed to offload a plethora of different cognitive functions, including (but not limited to) memory, learning, planning, decision-making, calculating, and communication.⁷ Moreover, neuromedia is prone to be used on a frequent basis because of how tightly coupled the technology is with our neural systems. As smart devices become more seamlessly coupled with our minds, it becomes easier

⁷ BCIs may render verbal communication unnecessary by making agents telepathic. Brain-to-brain interfaces allow for direct transmission of signals of ‘intended speech’ from one person to another. See DARPA’s ‘Silent Talk’ program: <https://www.wired.com/2009/05/pentagon-preps-soldier-telepathy-push/>.

and easier to offload cognitive tasks onto them, and people thereby become increasingly incentivized to do so. Finally, there may be strong social and economic pressures to hook oneself up to neuromedia, as it is plausible that the technology will become indispensable to personal and professional life in much the same way that the smartphone is today. People may feel as if they need to augment their brains with BCI technologies to stay competitive in the future economy. All of these factors lend credence to the idea that neuromedia will engender a relationship of excessive cognitive offloading.

The following argument illustrates how excessive cognitive offloading can have a detrimental impact on intellectual virtue cultivation:

3.1 The Cognitive Offloading Argument

- P1. The cultivation and maintenance of intellectual virtue requires sustained active cognition on the part of the agent.
- P2. Cognitive offloading implies less active cognition, and when done in excess, can render agents cognitively lazy or at least cognitively passive.
- C. Therefore, excessive cognitive offloading can impede the cultivation and maintenance of intellectual virtue by rendering agents cognitively lazy or at least cognitively passive.

‘The argument is cast at a high level of generality. It does not focus on how a particular cognitive artifact is poised to undermine intellectual virtue, and it does not single out a particular intellectual virtue that is poised to be undermined by cognitive artifacts. Instead, the argument demonstrates how excessive cognitive offloading can impede intellectual virtue development in general. The basic worry motivating the argument is that artifacts can become cognitive lynchpins (P2), allowing agents to avoid doing the difficult mental labor required for intellectual virtue possession (P1).

Of course, the argument can be applied to specific artifacts and intellectual virtues. There are already various applications of the cognitive offloading argument in the virtue epistemology literature. For example, in his book *The Internet of Us* (2016), Michael Lynch argues that the internet threatens to undermine the virtue of understanding by engendering a disposition of cognitive passivity. Understanding can be construed as either an epistemic end-state that is attainable via virtue or as itself a kind of virtue (Lynch, 2021). Understanding as an intellectual virtue involves the ability to grasp dependency relations or how different components of a system ‘hang together’ (Grimm, 2019).⁸ Lynch’s basic contention is that ‘Google-knowing’ (i.e. our increasing epistemic dependence on internet search engines like Google) is a fundamentally receptive form of knowledge, and as such, lacks the elements of active inquiry and information integration that are necessary for grasping dependency relations. While

⁸ Lynch provides the following helpful example of what the virtue of understanding entails: “Consider, for example, the lucky person who understands how her car works. She has this understanding in part, as we’ll discuss more fully in a moment, because she has certain skills, skills that give her the ability to see how various parts of a machine depend on one another: you can’t get the car to move without the battery and the battery won’t be charged without the alternator. You understand when you see not just the isolated bits, but how those bits hang together” (Lynch 2021: p. 4).

the internet allows agents to immediately accrue knowledge about any matter of fact, Lynch believes that this knowledge often does not amount to genuine understanding because it is shallow and passively received instead of actively apprehended. Carter (2018, 2020) discusses the cognitive offloading argument as it applies to the virtue of intellectual autonomy, or the capacity for epistemic self-direction. According to Carter, excessive offloading compromises intellectual autonomy by diminishing our freedom to achieve epistemic goods, where ‘achievements’ are understood as necessarily being the result of a difficult process that requires some degree of agential effort.⁹ The basic idea is that virtuous intellectual autonomy is weakened when subjects become overly dependent upon artifacts for cognitive functioning such that their process of belief-formation is no longer actively shaped by features of their intellectual agency but instead passively determined by the relevant technology. In this paper, I focus on the cognitive offloading argument not as it pertains to the virtues of intellectual autonomy or understanding but rather to the virtue of intellectual perseverance. Before getting to this task, though, the cognitive offloading argument must be properly motivated.

Consider first P1. This paper upholds a *virtue responsibilist* conception of intellectual virtue (Montmarquet, 1993; Zagzebski, 1996), according to which virtues are acquired character traits that are developed via habituation and that agents are at least partially responsible, and therefore praiseworthy, for obtaining and preserving. Virtue responsibilism contrasts with a *virtue reliabilist* conception of intellectual virtue, which holds that virtues are any traits that function to reliably generate true beliefs (Greco, 2010; Sosa, 2007). The virtue reliabilist considers both acquired character traits like intellectual perseverance and innate sub-personal traits like sense perception to be intellectual virtues insofar as these traits are reliable truth-trackers.¹⁰ Unlike virtue reliabilism, virtue responsibilism does not take innate sub-personal traits, or what might be called ‘cognitive abilities’, to be intellectual virtues.

Cognitive abilities are inherent dispositions or dispositions whose maintenance and acquisition do not require continuous mental effort on the agent’s part. For example, having good eyesight is a mere cognitive ability because it is innately given and passively exercised. Responsibilist intellectual virtues, on the other hand, are fostered dispositions that are formed and maintained by habitually executing a set of actions. Since intellectual virtues are excellences of the mind, the primary type of action required for their acquisition are *intellectual actions*, or voluntary actions that can be performed internally (i.e. within the confines of one’s head) and that do not necessarily involve the manifestation of any observable behavior. Examples of intellectual actions include (but are not limited to) believing, wondering, entertaining, assessing, hypothesizing, inferring, attending, imagining, recognizing, questioning, recalling, ignoring, affirming, grasping, and doubting. The phrase ‘active cognition’ in the context of the cognitive offloading argument refers to the performance of intellectual actions. To say

⁹ See Bradford (2015) for more on this difficulty-based view of achievement.

¹⁰ Virtue-responsibilism and virtue-reliabilism do not exhaust the possible positions one can take with respect to the nature of intellectual virtue. Heather Battaly has developed a third view that combines elements of both responsibilism and reliabilism. The view, which she labels *personalism*, concurs with responsibilism that intellectual virtues must be cognitive character traits but rejects the responsibilist idea that intellectual virtues are necessarily praiseworthy or that we are necessarily (partly) responsible for possessing such virtues (Battaly 2018a).

that intellectual virtue possession requires sustained active cognition (P1) is just to say that these virtues are acquired intellectual character traits formed via the habitual performance of said actions.

The relevant set of intellectual actions necessary for virtue possession depends upon the particular virtue in question. For instance, to develop the virtue of open-mindedness, agents must do things like consistently entertain alternative points of view, question their belief systems, and recognize when their ideological preconceptions are shown to be mistaken (Baehr, 2011b; Riggs, 2010). By contrast, to cultivate the virtue of intellectual humility, agents must habitually execute a different set of intellectual actions, such as attending to and owning their intellectual limitations (e.g. gaps in knowledge, intellectual character flaws), as well as accurately assessing and potentially even downplaying their intellectual strengths (Whitcomb et al., 2017). *Virtue-conducive intellectual actions* can be understood as any intellectual actions that promote intellectual virtue development. *Virtue-inconducive intellectual actions*, by contrast, can be understood as any intellectual actions which impede the development and maintenance of intellectual virtue.¹¹

Notably, P1 outlines a necessary but not sufficient condition for intellectual virtue possession. The relevant intellectual actions must not only be habitually performed but also (a) motivated by epistemically respectable reasons and (b) skillfully executed in the sense that they are guided by practical wisdom or what Aristotle calls ‘phronesis.’ To say that virtuous agents are motivated by epistemically respectable reasons is to say that their performance of virtue-conducive intellectual actions is rooted in “an intrinsic concern for or ‘love’ of epistemic goods” (Baehr, 2011a). Phronesis, by contrast, is what enables agents to ‘hit the mean’ with respect to each virtue such that they avoid corresponding vices of excess and deficiency. Agents who harbor practical wisdom exercise virtue-conducive actions “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).¹²

The veracity of P1 alone entails that intellectual virtues, understood along responsibilist lines, can never be *directly engineered* into agents via technological augmentation. Cognitive artifacts can directly engineer cognitive abilities into agents (e.g. glasses instantaneously give agents with poor vision the ability to have good eyesight), but such artifacts can never automatically confer cognitive character traits given that these traits require sustained active cognition on the part of agents.¹³ Pritchard (2018) makes this point in the context of discussing the educational implications of

¹¹ Notably, there may be cases in which some set of intellectual actions Z functions to promote the development of a virtue Y while simultaneously impeding the development of a different virtue X. Cases like this denote the possibility that a set of intellectual actions may be both virtue-conducive and virtue-inconducive depending upon the virtue in question.

¹² A deficiency of practical wisdom can be found in the agent who, in an effort to cultivate open-mindedness, habitually engages with and seriously entertains all alternative perspectives, even the perspective of a Nazi who denies the existence of the Holocaust or that of a conspiracy theorist who alleges that the earth is flat. Any agent who seriously considers these kinds of alternative perspectives plausibly possesses the character trait of open-mindedness to an excess degree such that it has become a vice (Battaly 2018b).

¹³ Some artifacts even have the power to directly engineer novel cognitive abilities which transcend the capacities of our nonaugmented neural systems. For instance, a neuromedia device equipped with object recognition artificial intelligence technology might allow agents to automatically identify different places

neuromedia, arguing that while neuromedia enables agents to instantaneously access any fact from the internet, thereby rendering the educational practice of fact retention unnecessary, the technology will never allow agents to automatically acquire traits of intellectual virtue: “Ones intellectual virtues, and ones intellectual character more generally, are constituted by being reflective, managerial traits that guide ones employment of ones cognitive abilities (extended or otherwise) and ones knowledge (again, extended or otherwise). This means that it is built into these traits that they are manifestations of ones [technologically] unextended cognitive character” (Pritchard, 2018: p. 341). Pritchard concludes on this basis that educators will still be tasked with advancing the intellectual character (if not the cognitive abilities) of students in a world where neuromedia is ubiquitous.

It may be helpful to make a distinction between direct engineering and indirect engineering in this context. Intellectual virtues can be *indirectly engineered* into agents by different technologies and educational institutions. Indeed, people will often say that their education made them more virtuous. What they really mean by this is that the educational institution provided them with ample opportunities and resources to internalize the requisite cognitive practices of habituation which undergird intellectual virtue development. However, it is ultimately a subject’s own agency, not their educational institution, that is the proximate cause of their newfound virtue. To say that virtues can be directly engineered into agents by cognitive artifacts is to say that these artifacts can supersede one’s cognitive agency as the proximate cause of virtue possession. This direct engineering of virtue is impossible from the standpoint of virtue responsibilism, which mandates that one’s cognitive agency plays an active role in virtue development.

When P1 is combined with P2, it becomes clear that in addition to cognitive artifacts not being able to directly engineer virtues into agents, overreliance on some artifacts can actually impede intellectual virtue development. This insight emerges from the fact that (P1) the development and maintenance of intellectual virtue requires the sustained and skillful execution of intellectual actions, and (P2) excessive cognitive offloading fosters a disposition of cognitive passivity wherein agents have their artifacts perform technological actions in place of virtue-conducive intellectual actions that they would otherwise perform to fulfill cognitive tasks.

One of the main functions of a cognitive artifact is to help carry the ‘cognitive load’ for any given mental operation. The offloading of a mental operation to an artifact means that the biological agent is no longer entirely responsible for completing the operation in question, which implies that the agent does not have to engage in as much active cognition to execute said operation. There exists a negative correlation between cognitive offloading and active cognition: as offloading grows increasingly excessive, the level of active cognition required for daily cognitive functioning decreases (Barr et al., 2015). Tightly coupled, all-purpose cognitive artifacts like neuromedia are the most concerning from the perspective of intellectual virtue development precisely because they foster a relationship of excessive cognitive offloading, and therefore, a disposition of cognitive passivity. The worry, to repeat, is that neuromedia users may

Footnote 13 continued

and people in the world. Similarly, natural language processing technologies can give agents the ability to translate foreign languages in real-time.

not have sufficient opportunities to ingrain the cognitive habits and skills required for intellectual virtue possession, given that the technology significantly depletes their mental lives of active cognition.

I will now raise two noteworthy objections to the cognitive offloading argument, which, when fully considered and unpacked, really serve as qualifications to the argument. The first objection targets P2, alleging that cognitive offloading, even when done excessively, does not necessarily render agents cognitively lazy or cognitively passive because the process of mastering and responsibly overseeing our offloading devices itself mandates a substantial degree of sustained active cognition. Following Carter (2018), one might distinguish between *information gathering labor* and *labor involved in managing our offloading gadgets* and contend that cognitive artifacts only automate the former kind of intellectual labor. The basic idea is that “we cannot have cognitive offloading ‘all the way down’; at some point, there is intellectual labour—labour associated specifically with the responsible management of offloading technologies—which cannot be offloaded to further technologies” (Carter, 2018: p. 15). For example, use of a personal notebook as a memory storage device demands continuous agential effort in the form of writing, reading, monitoring, and updating. P2, then, arguably offers an overly simplistic picture of cognitive offloading. Cognitive offloading does not fully replace a brain task with an equivalent technological task but rather involves transformed cognitive tasks requiring some degree of agential effort. So, while cognitive artifacts may deprive agents of epistemic opportunities to habituate some intellectual virtues; namely, those loosely associated with information gathering labor, they provide ample opportunities to garner other virtues associated with artifact management labor or what might be called cognitive transformation labor.¹⁴

This objection is correct as far as it goes. However, the intellectual labor involved in mastering and managing our cognitive artifacts is often trivial compared to the intellectual labor that is offloaded onto these artifacts. Learning how to use and navigate the internet, for instance, necessitates a certain degree of effort, but this effort pales in comparison to the effort that it would take to independently (without any technological aids) accumulate all of the information presently stored on the internet. Moreover, the cognitive transformation labor related to cognitive offloading will itself be increasingly offloaded as artifacts become more technologically autonomous and tightly coupled with our minds. Unlike notebooks, neuromedia largely removes humans from ‘the loop’ of technological operation, as the technology mandates only a minimal level of cognitive-intentional input on the part of agents. The upshot is that the more high-tech and tightly coupled with our minds a cognitive artifact is, the more disposed it will be to foster an epistemically consequential level of cognitive passivity within agents.

A second objection to the cognitive offloading argument targets the argument as a whole, alleging that it presupposes an overly internalist picture of cognition by neglecting the possibility that cognitive artifacts can become extensions of the mind. For a cognitive artifact to qualify as a cognitive extender, according to Clark and Chalmers (1998), the artifact must play the same functional role as an internal brain

¹⁴ I am indebted to an anonymous reviewer for raising this possible objection to P2.

process, and in particular, must be as trustworthy, reliably available, and easily accessible as internal cognitive resources like biological memory.¹⁵ Brain-enhancement devices are extremely plausible candidates for satisfying these traditional criteria for cognitive extension precisely because of how tightly coupled they are with our neural systems. If devices like neuromedia *metaphysically* extend our minds and do not merely *phenomenologically* feel like cognitive extenders, then reliance on them arguably no longer even takes the conceptual form of cognitive offloading. Cognitive offloading involves outsourcing a cognitive task to something external to one's mind, but if an artifact is a cognitive extender, then it is no longer properly conceived of as outside one's mind but is instead viewed as a constitutive component of one's mental apparatus.

Conceptualizing cognitive artifacts in this manner suggests that instead of "emptying our minds of their riches" (Carr, 2010: 192), our smart devices may be epistemically enriching our minds by extending them into our technologies.¹⁶ As Clowes (2013) remarks, "From the internalist (and embedded) vantage-point it is as if our minds are being steadily off-loaded and dissipated into our tools, potentially with the result that we become less autonomous, less knowledgeable and arguably less interesting... For HEC [hypothesis of extended cognition] theorists this does not follow, for rather than being dissipated into our tools, we are incorporating them into us" (Clowes, 2013: p. 127). Contra the cognitive offloading argument, then, it may be the case that cognitive artifacts function to *extend* intellectual virtues (instead of *undermining* them) by becoming constitutive components of an agent's mind.

This second objection faces two complications. First, many scholars reject the extended mind thesis as being too radical and opt instead for the embedded mind thesis, which is the more metaphysically conservative view that technologies and other environmental props at best function as cognitive scaffolds that causally support (as opposed to partially constitute) mental processes (Rupert, 2004, Adams & Aizawa 2008). Second, and more pressingly, even if artifacts like neuromedia count as constitutive components of an agent's *cognitive architecture*, this does not entail that they are constitutive components of the agent's *intellectual character*. Technologies like neuromedia may extend cognitive processes while simultaneously undermining intellectual virtue development.

The concept of extended intellectual virtue is highly controversial, although there has recently been a wave of research in the virtue epistemology literature promoting the idea (Battaly, 2018c; Carter, 2020; Cash, 2010; Howell, 2016; Pritchard, 2010, 2015). The consensus within this literature is that the conditions for extended virtue (as well as extended knowledge) are significantly more demanding than the conditions for mere

¹⁵ This is what's known as *the parity principle* for cognitive extension: "If, as we confront some task, a part of the world functions as a process which, *were it done in the head*, we would have no hesitation in recognizing as part of the cognitive process, then that part of the world is (so we claim) part of the cognitive process" (Clark & Chalmers 1998: 8).

¹⁶ One particularly intriguing possibility is that neuromedia facilitates what philosophers have called *internet-extended knowledge* (Smart 2012, Bjerring & Perdersen 2014, Ludwig 2014).

extended cognition because intellectual virtue represents a distinctive kind of cognitive achievement (Carter & Kallestrup, 2019).¹⁷ Whereas extended cognition can be facilitated by passively hooking oneself up to technology, integrating an artifact into a subject's intellectual character such that it extends their agency tenably requires active cognitive engagement with the artifact, and in particular, the kind of intellectual labor associated with mastering and responsibly managing artifacts that was discussed in the context of the previous objection. For example, Pritchard (2015) says that students can actively incorporate educational tools into their intellectual characters provided that they are taught to deploy and manage these tools in an epistemically diligent and conscientious manner: "it is critical to the process of cognitive integration that the subject displays due diligence in making epistemic use of the device, and this will involve the manifestation of intellectual virtues like conscientiousness. This manifestation of intellectual virtue is not incidental to cognitive integration, but rather a key element, in that without it we would lose a grip on the idea that the subject has actively integrated this device into her cognitive practices" (Pritchard, 2015: p. 123).¹⁸ Similarly, Carter (2020) suggests that devices like neuromedia, even when relied upon excessively for offloading purposes, can extend virtuous intellectual autonomy, granted that agents take appropriate 'cognitive ownership' over the device's role in their process of belief formation and direction of epistemic inquiry.

For the sake of brevity, I will put the controversial concept of extended virtue epistemology aside and choose to restrict the cognitive offloading argument to cases where cognitive artifacts are not epistemically integrated into the subject's intellectual character. However, it is worth noting that even if the concept of extended virtue epistemology is coherent, it may be that the level of cognitive passivity engendered by a device like neuromedia prevents or at least dissuades agents from sufficiently integrating the technology into their intellectual characters. If Pritchard (2015) and others are correct in claiming that the facilitation of extended intellectual virtue itself requires the manifestation of certain virtues like epistemic diligence and conscientiousness, then excessive cognitive offloading may be antithetical to extended virtue to the extent that such offloading undermines or discourages the cultivation of the relevant character traits. The upshot is that the present objection risks getting things backward: the cognitive offloading argument may serve as a challenge to the realization of extended intellectual character and not vice versa. Having presented and motivated the cognitive

¹⁷ Pritchard (2010) is an exception to this consensus, as he argues that the conditions for *metaphysical* cognitive integration are the same as the conditions for *epistemic* cognitive integration.

¹⁸ Pritchard (2015) invokes the concept of extended virtue epistemology as a response to something resembling the cognitive offloading argument. He observes that an ostensible tension exists between a virtue-theoretic account of the epistemic goal of education and the fact that students tend to excessively depend upon cognitive artifacts for educational purposes. Namely, if the primary epistemic goal of education is to promote the intellectual character development of students, but excessive dependence upon cognitive artifacts (especially high-tech ones like neuromedia) impedes intellectual character development, then educators should arguably discourage the use of such artifacts in the classroom. Pritchard avers that this ostensible tension is illusory and that a virtue-theoretic account of the epistemic goal of education is entirely compatible with the use of cognitive artifacts for educational purposes. He says that as long as educators teach students to use artifacts like smartphones in an intellectually virtuous manner, they can actively integrate these artifacts into their intellectual characters.

offloading argument, I will now proceed to examine the argument as it applies to the virtue of intellectual perseverance.

4 Cognitive Efficiency and Intellectual Perseverance

In order to discern whether, according to the cognitive offloading argument, a specific cognitive artifact *Y* is disposed to undermine a specific intellectual virtue *X*, one must know (i) which set of intellectual actions *Z* an agent has to habitually perform in order to cultivate and sustain *X* and (ii) whether *Y* allows and incentivizes agents to offload *Z*, or substitute *Z* for functionally isomorphic technological actions. Put differently, the key question is: what kinds of actions *Z* are necessary for the development and maintenance of intellectual virtue *X*, and does the perpetual use of artifact *Y* lead to a diminishment of *Z* in an agent's cognitive economy?

This section answers this question in the case where *Y* is neuromedia and *X* is the virtue of intellectual perseverance. To understand how, according to the logic of the cognitive offloading argument, neuromedia is disposed to influence intellectual perseverance, it must first be determined what *Z* is in this case. In other words,

Q1 Which set of intellectual actions must an agent habitually perform to cultivate the virtue of intellectual perseverance?

To answer Q1, one must first provide a definition of intellectual perseverance. The trait of intellectual perseverance involves being mentally persistent and resilient with respect to one's epistemic projects. To intellectually persevere is not to give up and finish the cognitive task at hand when confronted with challenges and inconveniences. Nathan King (2014) offers the following paradigmatic examples of the virtue:

Thomas Edison endured years of work and thousands of failures in his quest to develop the incandescent light bulb. Booker T. Washington overcame slavery, racism, and poverty in order to gain an education and disseminate his political views. Helen Keller overcame blindness and deafness in order to do the same. Isaac Newton labored for years to develop the calculus needed for his system of physics (King, 2014: p. 3502).

Heather Battaly (2017) formally defines the virtue thusly: "Intellectual perseverance is a disposition to overcome obstacles, so as to continue to perform intellectual actions, in pursuit of one's intellectual goals" (Battaly, 2017: p. 671). An *intellectual goal* is just any epistemic end that a subject possesses. Intellectual goals are relativized to subjects, and so by extension, the disposition of intellectual perseverance is as well. For example, my current intellectual goal might be to finish a philosophy article, whereas someone else might have the intellectual goal of learning how to speak a different language, or getting a bachelor's degree, or reading a novel, or solving a math problem. All of these goals count as 'intellectual' because they all constitute some kind of epistemic success. Agents harbor the trait of intellectual perseverance insofar as they are disposed to consistently execute intellectual actions which further their intellectual goals despite any challenges or obstacles that may arise. The trait of intellectual perseverance can be contrasted with the trait of physical perseverance, which involves remaining committed to attaining one's physical or practical goals in

the face of obstacles. For instance, a marathon runner who starts to cramp up during a race but finishes the race anyways exhibits the trait of physical perseverance.

Notably, one can distinguish between the *trait* of intellectual perseverance and the *virtue* of intellectual perseverance. As Battaly (2017, 2020) clarifies, intellectual perseverance does not always manifest itself as an intellectual virtue. As mentioned in section III, to qualify as an intellectual virtue, a character trait arguably must have the appropriate motivational structure, which is to say, it must be motivated out of a desire to attain epistemic goods like truth and knowledge. However, an agent does not need to be motivated out of a desire for truth in order to instantiate the trait of intellectual perseverance. Someone might be motivated to achieve their intellectual goals and persist in the face of obstacles because realizing those goals brings a monetary reward. Alternatively, an agent might be motivated out of spite or a sense of competitiveness to see a project through until it is finished. Such a person would possess the trait of intellectual perseverance, but the trait would not be an intellectual virtue in this case. Relatedly, Battaly observes that intellectual perseverance also fails to be a virtue when the intellectual goals in question do not meet a certain threshold of epistemic value. Consider, for example, an agent whose intellectual goal consists entirely in stoking political division on the internet. Such an agent might be motivated to accomplish this goal in the face of difficulties and therefore possess the trait of intellectual perseverance, but the agent would not be engaging in intellectually virtuous behavior. For the sake of simplicity, I will assume in this paper that agents who are interested in cultivating intellectual perseverance possess intellectual goals which meet the required threshold of epistemic value and are motivated to pursue these goals for epistemically respectable reasons.

Given this characterization of intellectual perseverance, an initial answer to Q1 which seems reasonable is ‘the set of intellectual actions conducive to an agent’s epistemic ends.’ According to this line of thought, brain-enhancement devices like neuromedia threaten to undermine intellectual perseverance by accomplishing intellectual goals for agents, thereby robbing them of many internal opportunities for action that are necessary for the cultivation and maintenance of the virtue. At this point one might interject: ‘If neuromedia can undermine intellectual perseverance by robbing agents of opportunities for intellectual action which further their epistemic ends, does it follow that the technology might also advance the virtue by preventing agents from performing intellectual actions that amass epistemic opportunity costs? That is, just as neuromedia can hinder intellectual perseverance by ridding an agent’s mind of *epistemically productive cognition* (i.e. active cognition which furthers the agent’s intellectual goals), is it not also possible that the technology can facilitate the habituation of intellectual perseverance by ridding an agent’s mind of epistemically unproductive cognition, or what might be called *cognitive noise* (i.e. active cognition which is unconducive to, or distracts from, the agent’s intellectual goals)?’

Consider the following thought experiment, which illustrates the point that the interjector has in mind. Suppose that John is prone to mind wandering such that whenever he encounters some cognitive challenge, he tends to ignore the challenge and think about other things. Put simply, John uses mind wandering as a kind of cognitive avoidance mechanism. Suppose further that John’s current intellectual goal is to complete an academic paper for his class. As soon as he sits down to write the paper,

John is confronted by an immediate obstacle: the mountain of preliminary research that he must complete before he is ready to start the writing process. John knows that this a daunting task that involves looking through his class notes, searching the internet for relevant academic sources and quotations, integrating different arguments and ideas, and devising a detailed outline for the paper. When confronted with this obstacle to his intellectual goal, John is disposed to engage in mind wandering or start to have thoughts that have nothing to do with the task at hand. These thoughts solicit further irrelevant thoughts, ultimately leading John down a cognitive rabbit hole that prevents him from making any intellectual progress.

If John were equipped with a neuromedia device, he would be confronted with significantly fewer obstacles because the device would carry out many cognitive tasks for him. Moreover, assume that John equips himself with what might be called *passive neuromedia*. Whereas an *active* BCI requires intentional mental effort to operate (i.e. the issuing of a mental command), a *passive* BCI does not mandate any mental effort or voluntary psychological control, but instead monitors the user's brain activity in real-time and executes calculated actions based upon the data it collects via this monitoring. Steinert et al., 2019 discuss a recently developed passive BCI device that serves to "detect affective states and use this information to improve the interaction of humans and computers. For example, the system might adapt the computer game or task to be performed when it detects that the user is frustrated or bored, by decreasing the level of difficulty or by introducing more engaging elements" (Steinert et al., 2019). A passive neuromedia device may, without the user knowing, integrate and reorganize stored memories, engage in various predictive processing functions, and perhaps even directly cause epistemic and emotional changes to one's mental states via deep brain stimulation (Wardrope, 2014). The aforementioned passive BCI device monitors the affective mental states of users with the preprogrammed goal of helping them optimize cognitive performance.

Passive neuromedia might be conceived of as the intellectual equivalent of a Fitbit device. A Fitbit is a health-promoting wearable technology that advances one's physical fitness goals by collecting and recording a user's biomedical data, such as exercise, sleep quality, and heart rate (Owens & Cribb, 2019). Both Fitbit and passive neuromedia are examples of '*quantified self*' technologies, or self-tracking tools that are deployed in the service of self-optimization (Kunze et al., 2013). Elon Musk has even described the brain-computer interface devices they are creating at Neuralink as a 'Fitbit in your skull' (Koetsier 2020). The principal difference between the two technologies is that neuromedia is implantable instead of wearable and functions to help agents more easily accomplish their intellectual goals as opposed to their physical fitness goals.¹⁹ We might imagine that John is hooked up to a passive version of neuromedia which actively tracks his brain states and thoughts and then uses this neural data to facilitate the fulfillment of his intellectual goals, which in this case

¹⁹ Smart, Clowes, & Heersmink elaborate on how self-tracking technologies can be used to optimize cognitive performance: "Inasmuch as future devices are able to track physiological signals that index a variety of epistemic feelings (such as the feeling of knowing or the feeling of difficulty), it is possible that future technologies may be able to adaptively modify their modes of operation to support human end-users with regard to a variety of cognitive- and epistemically-relevant performances" (Smart, Clowes, & Heersmink 2017: p. 42).

means automatically scanning the internet and generating a list of relevant academic sources, appropriately integrating and condensing these sources, and devising an outline of his research paper. One might think that this device would help John to become more intellectually perseverant, not less. With fewer cognitive obstacles to confront, neuromedia-equipped John would arguably engage in less mind wandering as a cognitive avoidance mechanism as there would no longer be significant obstacles to avoid.

In response to this suggestion, it is imperative to distinguish between *cognitive efficiency* and *intellectual perseverance*. Cognitive efficiency involves being able to accomplish intellectual goals at an optimal rate using minimal wasted mental resources. Intellectual perseverance, on the other hand, consists in being able to overcome obstacles in the pursuit of one's intellectual goals. The problem is that John's neuromedia device does not foster the disposition to overcome obstacles; rather, as the example illustrates, the technology functions to *remove* obstacles from his cognitive economy. The upshot is that while neuromedia might help John become more cognitively efficient by ridding his mind of cognitive noise (in this case, cognitive noise associated with mind-wandering), this does not mean that the device supports the development of intellectual perseverance. Quite to the contrary, *neuromedia may increase cognitive efficiency at the cost of intellectual perseverance*. In her conceptual analysis of the virtue Battaly emphasizes that accomplishing an intellectual goal is not sufficient for persevering with respect to that goal and claims that goals which are accomplished too easily do not (and cannot) involve the exercise of intellectual perseverance: "the trait of IP does not allow intellectual goals to be too easily achieved. It excludes goals that are so easily achieved that they don't permit any obstacles (the goal of solving ' $1 + 1 = x$ ') and don't require any effort on the part of the agent" (Battaly forthcoming: 169). The problem in this context is that a device like neuromedia allows agents to attain intellectual goals in the absence of nearly any effort or obstacles.

One might ask: who cares if neuromedia has a detrimental impact on intellectual perseverance so long as it increases cognitive efficiency? After all, the primary function of intellectual perseverance is to help agents accomplish their intellectual goals. If neuromedia is a better means for attaining one's intellectual goals than intellectual perseverance ever could be, then perhaps it is worth implanting neuromedia at the cost of intellectual perseverance. The natural reply to this line of reasoning by the virtue epistemologist is to observe that intellectual virtues are not merely *instrumentally* valuable but are also *intrinsically* valuable. That is, intellectual virtues do not simply serve as a means to the end of acquiring truth but are also ends in and of themselves. This means that intellectual perseverance is worth possessing for its own sake and is not just valuable insofar as it helps agents accomplish their intellectual goals. Thus, the virtue epistemologist might concede that neuromedia is a better mechanism than intellectual perseverance for meeting one's epistemic ends but deny that the tradeoff is worth it because of the intrinsic value that would be lost in sacrificing intellectual virtue.

The 'mind-wandering John' thought experiment reveals that the initial answer given to Q1 is incomplete. There is a difference between being disposed to perform intellectual actions that further one's epistemic ends *in the absence of any obstacles* and being disposed to do so *in spite of any obstacles*. While the latter disposition describes the disposition of intellectual perseverance, the former refers to what might be called

the disposition of *intellectual engagement*.²⁰ It is (at least conceptually) possible to possess the disposition of intellectual engagement while lacking the disposition of intellectual perseverance. Suppose, for example, that an agent habitually performs intellectual actions which further their epistemic ends but does so without ever having to confront any significant obstacles along the way. While this agent may develop the disposition to perform said actions (i.e. the disposition of intellectual engagement), they will not develop the disposition to overcome obstacles and, consequently, fail to cultivate intellectual perseverance. *Thus, to cultivate intellectual perseverance, agents must not only habitually engage in epistemically productive cognition but must do so in spite of the presence of at least some obstacles.*

It is an open question as to exactly how many obstacles are necessary to meet the minimal threshold for developing the virtue of intellectual perseverance. The required obstacles can vary in nature and may be either internally manufactured by one's own mind (in the form of anxiety, boredom, the desire to do something else, etc.) or externally imposed by one's environment (in the form of social oppression, financial instability, lack of health, etc.). Some agents might have more opportunities to cultivate intellectual perseverance precisely because they encounter more obstacles in life than others. Thus, while a technologically nonaugmented agent might be less cognitively efficient than a neuromedia-equipped agent, the former, all else being equal, has more opportunities to cultivate intellectual perseverance in virtue of being confronted with a greater number of obstacles in the course of pursuing their epistemic ends.

This is not to say that neuromedia users will have zero, or even few, opportunities to cultivate intellectual perseverance, however. It is crucial to emphasize that *everything is contingent on how the technology is used*. Neuromedia need not rid our cognitive economies of all obstacles but could instead just alter the kinds of obstacles we face by shifting our intellectual goals. For example, people may deploy the technology solely to automate different types of menial intellectual labor so that they can focus on more worthwhile epistemic ends. Smartphone users often offload particularly dull cognitive tasks like navigating or remembering addresses as a way to reduce their epistemic opportunity costs and ensure that they have the mental bandwidth to attend to the intellectual goals that they care about the most. This type of cognitive offloading does not necessarily undermine intellectual perseverance; it just allows agents to persevere with respect to intellectual goals that they find more valuable. Relatedly, in their book *Deep Thinking: Where Machine Intelligence Ends and Human Creativity Begins* (2017), Garry Kasparov and Mig Greengard aver that artificial intelligence systems may serve to amplify human creativity and other intellectual virtues by freeing us from the drudgery of various kinds of tedious unwanted intellectual labor. In a recent interview on the book, Kasparov remarks:

Let's look at this historical process. Machines that replace manual labor, they have allowed us to focus on developing our minds. More intelligent machines, I understand

²⁰ Notably, intellectual engagement is not synonymous with cognitive efficiency, as the former does not imply intellectual accomplishment in the way that the latter does.

that they will take over more menial aspects of cognition and will elevate our lives towards curiosity, creativity, beauty, joy, you can continue this line.²¹

There is nothing in principle preventing neuromedia from being used in a way that aligns with Kasparov's techno-optimistic vision. 'Mind-wandering John' may choose only to offload the relatively dull intellectual labor associated with writing an academic paper and save the more stimulating, creative intellectual work for himself to accomplish. Reliance on neuromedia starts to become excessive from the standpoint of virtue epistemology when agents use it to outsource not only menial intellectual labor, but also more stimulating, distinguished cognitive tasks. The central worry, to reiterate, is that we will enter a future characterized by widespread cognitive atrophy in which people have their neuromedia devices accomplish all manner of intellectual goals for them, including those related to their most important daily work. Consider by way of analogy the futuristic society depicted in the Pixar movie WALL-E. The human inhabitants of this fictional universe are confined to a generation starship and have become extremely obese and physically lazy due to their overreliance upon autonomous artificial intelligence systems, which take care of their every need. The physical health predicament of humanity in this movie serves as a counterpart to the kind of epistemic health concern expressed by the cognitive offloading argument. Excessive reliance on cognitive artifacts like neuromedia threatens to render agents cognitively lazy and cause the degeneration of intellectual character traits in much the same way that excessive reliance on autonomous AIs renders the humans in WALL-E physically lazy and obese.

It bears repeating, however, that none of this is inevitable. Cars, planes, motorcycles, and other vehicles already largely eliminate the need for physical exercise, but people still find a way to remain physically fit. Agents typically offload essential physical tasks like getting to and from work to their cars and then run on a treadmill as a way to stay in shape. Similarly, neuromedia users may decide to outsource important intellectual labor to the device but still find time to remain 'mentally fit' by incorporating different brain games or mental gymnastics into their daily schedules. Brain games like sudoku and luminosity generate artificial obstacles and intellectual goals which grant agents the opportunity to cultivate and exercise intellectual perseverance in a consequence-free environment. Perhaps staged intellectual fitness activities of this sort will increase in popularity with the societal uptake of neuromedia, and neuromedia-equipped agents will still be motivated to actively develop virtues like intellectual perseverance despite having the technological option of resigning to cognitive passivity.²²

I do not know what the epistemic future of humanity holds, and in particular, whether the advent of neuromedia will lead to some kind of epistemic dystopia or perhaps be more in line with Kasparov's vision of techno-optimism. The cognitive offloading argument simply alerts us to the virtue epistemic dangers of excessive use of neuromedia for offloading purposes. To summarize: the pertinent question in discerning whether, according to the cognitive offloading argument, a cognitive artifact

²¹ The interview can be found here: <https://medium.com/conversations-with-tyler/garry-kasparov-tyler-cowen-chess-iq-ai-putin-3bf28baf4dba>.

²² I am indebted to an anonymous reviewer for bringing this possibility involving staged intellectual fitness activities to my attention. The analogy between using a treadmill to remain physically fit and playing brain games to remain intellectually fit derives from said reviewer.

Y is disposed to undermine an intellectual virtue X is: what kinds of intellectual actions Z are necessary for the development and maintenance of X, and does the perpetual use of Y lead to a diminishment of Z in the agent's long term cognitive economy? This section has shown that for the virtue of intellectual perseverance, 'Z' is identified with intellectual actions which (a) further an agent's intellectual goals and which (b) are executed in the presence of at least some obstacles. Neuromedia compromises IP to the extent that it accomplishes intellectual goals for agents and removes obstacles from their cognitive economies.

5 Cognitive Inloading and Intellectual Perseverance

Thus far, it has been established that, when used as a cognitive offloading device, a technology like neuromedia has the capacity to undermine intellectual virtue by *stripping* an agent's cognitive economy of *virtue-conducive intellectual actions*. However, neuromedia will also serve as a 'cognitive inloading' device, meaning that it will be used not only to outsource cognitive tasks from our minds to the technology, but also to insource information from the technology to our minds. Excessive cognitive inloading threatens to undermine intellectual virtue by *adding* to an agent's cognitive economy the presence of *virtue-inconducive intellectual actions*. The worry in this case is that smart devices will be used too frequently to insource informational content that is epistemically harmful or that leads agents epistemically astray. Reflection upon this converse epistemic risk suggests that there are at least two different ways in which artifacts like neuromedia can undermine intellectual perseverance: (i) by rendering agents cognitively lazy or cognitively passive due to excessive cognitive offloading and (ii) by rendering agents *cognitively distracted* due to excessive cognitive inloading. The best way to illustrate the latter concern is to focus on how the so-called online attention economy is eroding the attentional capacities of contemporary internet users.

We now live in the era of what Shoshana Zuboff (2019) has dubbed *surveillance capitalism*, which is a new form of capitalism in which companies, and in particular Big Tech companies like Google, Facebook and Amazon, use the surveillance of human behavior to gather personal data about digital consumers with the intention of monetizing this data to make a profit. From the perspective of surveillance capitalism, attention is a precious commodity. Web designers, app creators, online advertisers, and tech companies are constantly competing for the limited attention of digital consumers by deploying many subtle and not-so-subtle technological mechanisms.²³ Machine learning algorithms, for example, monitor the online behavior of internet users in real-time and generate increasingly personalized content and advertisements which are unprecedented in their ability to capture user attention (Fogg, 2009). Even when people are not consuming new media content, they are constantly receiving push notifications from apps that light up the screens of their smartphones, drawing them back into the digital zeitgeist. James Williams (2018) explains and analyzes how this 'online

²³ As Hanin (2021) puts it, we are confronted by "an asymmetric matchup between our fallible mechanisms of self-restraint and armies of engineers, programmers, designers, and executives working to extract ever-smaller "slivers" of our focus in a highly competitive attention economy (Wu 2016, 268; Williams 2018, xi–xii, 33)" (Hanin 2021: p. 3).

attention economy' is fostering a disposition of cognitive distraction within digital consumers, thereby inhibiting them from autonomously governing their mental lives and achieving sustained focus on a single task: "the main risk information abundance poses is not that one's attention will be *occupied* or *used up* by information, as though it were some finite, quantifiable resource, but rather that one will *lose control over* one's attentional processes" (Williams, 2018: p. 37).

Importantly, being cognitively distracted is not synonymous with being confronted by cognitive distractions. One must be careful not to commit the fallacy of equivocation here. Cognitive distractions are types of obstacles, the presence of which can help agents inculcate the virtue of intellectual perseverance given that they are able to overcome the relevant obstacles and remain focused on their intellectual goals. Being cognitively distracted, by contrast, presupposes that an agent has already succumb to obstacles such that their attention is preoccupied with what I previously called 'cognitive noise', or intellectual actions that are inconducive to their intellectual goals. Cognitive noise is antithetical to the development and exercise of intellectual perseverance, and so by extension, cognitive distraction is as well. The online attention economy serves to undermine intellectual perseverance by rendering digital consumers cognitively distracted. As Williams indicates, the disposition of cognitive distraction entails a diminished capacity for attentional control, which makes it particularly difficult to remain focused on one's intellectual goals and not succumb to obstacles when they present themselves in the form of cognitive distractions. Just consider the lack of attentional control on display when digital consumers perpetually succumb to the temptation to check their smartphones and mindlessly scroll through social media. We are constantly inloading information from the internet that distracts from our epistemic ends and inspires the performance of virtue-inconducive intellectual actions.

Neuromedia may exacerbate the epistemic harms of the online attention economy by bringing agents into closer cognitive contact with Web-based information. One worry here is that the technological forces of surveillance capitalism will escape the confines of the smartphone and play a more directly intrusive role in the mental lives of neuromedia users, given how tightly coupled their cognitive systems will be to the Web. Surveillance capitalism currently operates by collecting 'browsing data' from agents, which is to say, data extracted from the digital footprint agents leave behind on the internet. Neuromedia could lead to a new, more invasive form of surveillance capitalism, which Lesaja and Xavier-Lewis (2020) call *neurocapitalism*, in which tech companies directly extract 'neural data' from agents, or data concerning the brain states, emotions, and thoughts of consumers. The possibility of neurocapitalism suggests that tech companies could produce even more fine-grained consumer personality profiles and essentially understand the minds of consumers at the most intimate level imaginable, thereby augmenting their capacity for technological persuasion and ability to capture user attention in the name of profit.²⁴

The upshot of this discussion is that there are at least two vices of deficiency associated with the virtue of intellectual perseverance: cognitive laziness and cognitive distractedness. Nathan King (2014) argues that intellectual perseverance can be

²⁴ This kind of brain hacking scenario potentially has negative implications for other intellectual virtues as well, such as intellectual autonomy and open-mindedness.

viewed in Aristotelian terms as a mean between a vice of excess and a vice of deficiency. The vice of excess, according to King, is ‘intransigence’ whereas the vice of deficiency is ‘irresolution’. Intransigent agents do not know when it is appropriate to call it quits on a cognitive endeavor and tend to continue working on a project even when it is clear that the project is not going anywhere meaningful or is not epistemically worthwhile to pursue. Irresolute agents, by contrast, are disposed to prematurely abandon their intellectual goals as soon as they confront obstacles, or at least postpone these goals in the face of obstacles. The dispositions of cognitive laziness and cognitive distractedness might be conceptualized as two different forms of what King calls ‘irresolution.’²⁵ Neuromedia is especially well-suited to inculcate both of these vices of deficiency given how tightly coupled the technology is with our minds in comparison to contemporary smart devices.

6 Conclusion

To conclude, it is worth reiterating that the uptake of cognitive enhancement technologies does not inevitably cause the degradation of intellectual character. Cognitive artifacts can also promote intellectual virtue development via the kind of indirect engineering discussed in section III. Pritchard (2018) gives two examples of how cognitive artifacts might play a virtue-conducive role for an agent. The first is a possible app that he describes which helps and incentivizes internet users to become better inquirers in the context of online search engines. This envisioned app is in line with other so-called *digital wellness technologies* that have recently gained traction, such as the ‘Moment’ app and Apple’s ‘Downtime’ feature (Sullivan & Reiner, 2021). Digital wellness technologies like these might be construed as virtuous manifestations of surveillance capitalism because they collect and harvest personal data, not as a means to monopolize consumer attention at all costs but rather to nudge users towards behaviors and habits that promote their mental well-being and stated intellectual goals. The second example that Pritchard discusses concerns a real-life project that he was a part of which involved bringing an educational technology called MOOC (Massive Open Online Course) to prisons in an effort to spur the intellectual character development of inmates. The MOOC technology was focused on teaching the inmates critical thinking skills and introducing them to basic philosophical topics, and according to Pritchard, the project was at least partially successful. Neuromedia might also be used to promote intellectual virtue development. For instance, imagine a neuromedia device that supports the cultivation of open-mindedness by suggesting alternative perspectives to a topic that the agent may be overlooking or by pointing out cognitive biases that may be present in the agent’s thinking about an issue. In all of these examples, the relevant cognitive artifacts encourage the cultivation of intellectual virtue by motivating the agential performance of virtue-conducive intellectual actions.

It also bears repeating that the value of this paper is not contingent on a technology like neuromedia becoming ubiquitous in the near future. The cognitive offloading

²⁵ Battaly (2020) distinguishes between three different vices of deficiency for intellectual perseverance: slacking off, quitting, and procrastinating. It would be interesting to discern how her analysis both differs and overlaps with mine.

argument is still relevant to our current technological predicament as it applies to the internet, smartphones, and other AI-based technologies. The focus on neuromedia and cognitive enhancement BCIs is justified by the fact that these technologies are poised to take cognitive offloading to the extreme and so pose an especially serious risk to the cultivation of intellectual virtue. In addition, as described in section II, there are numerous recent and ongoing developments in the field of BCI research which suggest that the emergence of a cognitive enhancement device resembling neuromedia may be on the horizon. The goal of anticipatory virtue epistemology is to understand, analyze, and mitigate the epistemic risks associated with emerging technologies like neuromedia, sort of like the goal of AI safety research is to understand, analyze, and mitigate the ethical and existential risks related to artificial general intelligence and superintelligence (Bostrom, 2014). As cognitive artifacts continue to become more powerful and tightly coupled with cognition as the twenty-first century unfolds, it will be increasingly paramount for epistemologists to come to terms with how these technologies are transforming and are prone to transform the human mind.

Acknowledgements I deeply grateful to Susan Schneider, Heather Battaly, Michael Lynch, and two anonymous referees for offering helpful comments on previous versions of this paper. I would also like to thank William Lycan and Lewis Gordon for their mentorship as well as the Philosophy Department and Computer Science and Engineering Department at the University of Connecticut for their support over the years.

Open Access This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence, and indicate if changes were made. The images or other third party material in this article are included in the article's Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the article's Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this licence, visit <http://creativecommons.org/licenses/by/4.0/>.

References

- Abdulkader, S. N., et al. (2015). Brain Computer Interfacing: Applications and Challenges. *Egyptian Informatics Journal*, 16(2), 213–230.
- Adams, F., & Aizawa, K. (2008). *The Bounds of Cognition*. Blackwell: Wiley.
- Alfano, M., Fard, A. E., Carter, J. A., Clutton, P., & Klein, C. (2021). Technologically scaffolded atypical cognition: The case of YouTube's recommender system. *Synthese*, 199(1), 835–858.
- Baehr, J. (2011a). *The Inquiring Mind: On Intellectual Virtues and Virtue Epistemology*. Oxford University Press.
- Baehr, J. (2011b). The structure of open-mindedness. *Canadian Journal of Philosophy*, 41(2), 191–213.
- Barr, N., Pennycook, G., Stolz, J. A., & Fugelsang, J. A. (2015). The brain in your pocket: Evidence that smartphones are used to supplant thinking. *Computers in Human Behavior*, 48, 473–480.
- Battaly, H. (2020). Quitting, Procrastinating, and Slacking Off. In: I. Kidd, H. Battaly & C. Quassim (Eds.), *Vice Epistemology* (pp. 167–188). NY: Routledge.
- Battaly, H. (2017). Intellectual Perseverance. *Journal of Moral Philosophy*, 14(6), 669–697.
- Battaly, H. (2018a). A Third Kind of Intellectual Virtue: Personalism. In H. Battaly (Ed.), *The Routledge Handbook of Virtue Epistemology*. Routledge.
- Battaly, H. (2018b). Closed-mindedness and dogmatism. *Episteme*, 15(3), 261–282.
- Battaly, H. (2018c). Extending epistemic virtue: Extended cognition meets virtue-responsibilism. In P. Duncan, K. Jesper, P. Orestis, & J. A. Carter (Eds.), *Extended Epistemology*. Oxford University Press.

- Birbaumer, N., Kubler, A., Ghanayim, N., Hinterberger, T., Perelmouter, J., Kaiser, J., & Flor, H. (2000). The thought translation device (TTD) for completely paralyzed patients. *IEEE Transactions on Rehabilitation Engineering*, 8(2), 190–193.
- Bjerring, J. C., Pedersen, N. J., & Linding, L. (2014). All the (many, many) things we know: Extended knowledge. *Philosophical Issues*, 24(1), 24–38.
- Bostrom, N. (2014). *Superintelligence: Paths, dangers, strategies*. Oxford University Press.
- Bradford, G. (2015). *Achievement*. Oxford University Press.
- Carr, N. G. (2010). *The Shallows: What The Internet Is Doing To Our Brains*. W.W. Norton.
- Carter, J. A. (2020). Intellectual autonomy, epistemic dependence and cognitive enhancement. *Synthese*, 197, 2937–2961.
- Carter, J. A. (2018). Autonomy, Cognitive Offloading, and Education. *Educational Theory*, 68(6), 657–673.
- Carter, J. A., & Kallestrup, J. (2019). Varieties of Cognitive Integration. *Noûs*, 4, 867–890.
- Cash, M. (2010). Extended cognition, personal responsibility, and relational autonomy. *Phenomenology and the Cognitive Sciences*, 9(4), 645–671.
- Chisholm, R. M. (1989). *The Theory of Knowledge* (3rd ed.). Prentice Hall.
- Clark, A., & Chalmers, D. J. (1998). The extended mind. *Analysis*, 58(1), 7–19.
- Clowes, R. W. (2013). The Cognitive Integration of E-Memory. *Review of Philosophy and Psychology*, 4(1), 107–133.
- Fogg, B. J. (2009). Creating persuasive technologies: an eight-step design process. Persuasive '09. April 26–29, Claremont, California, USA.
- Gillett, A., & Heersmink, R. (2019). How navigation systems transform epistemic virtues: Knowledge, issues and solutions. *Cognitive Systems Research*, 56(56), 36–49.
- Goldman, A. I. (1967). A causal theory of knowing. *Journal of Philosophy*, 64(12), 357–372.
- Greco, J. (2010). *Achieving knowledge: A virtue-theoretic account of epistemic normativity*. Cambridge University Press.
- Greenfield, S. (2014). *Mind change: How digital technologies are leaving their mark on our brains*. Rider Books.
- Grimm, S. (2019). Understanding as an Intellectual Virtue. In H. Battaly (Ed.), *Routledge Companion to Virtue Epistemology*. New York, USA: Routledge.
- Hanin, M. L. (2021). Theorizing Digital Distraction. *Philosophy and Technology*, 34(2), 395–406.
- Heersmink, R. (2016). The internet, cognitive enhancement, and the values of cognition. *Minds and Machines*, 26(4), 389–407.
- Heersmink, R. (2018). A virtue epistemology of the Internet: Search engines, intellectual virtues and education. *Social Epistemology*, 32(1), 1–12.
- Howell, R. J. (2016). Extended Virtues and the Boundaries of Persons. *Journal of the American Philosophical Association*, 2(1), 146–163.
- Ienca, M., & Haselager, P. (2016). Hacking the brain: Brain–computer interfacing technology and the ethics of neurosecurity. *Ethics and Information Technology*, 18(2), 117–129.
- Kasparov, G., & Greengard, M. (2017). *Deep Thinking: Where Machine Intelligence Ends and Human Creativity Begins* (1st ed.). Public Affairs.
- Kelp, C. (2013). Extended Cognition and Robust Virtue Epistemology: Response to Vaesen. *Erkenntnis*, 78(2), 245–252.
- King, N. L. (2014). Perseverance as an intellectual virtue. *Synthese*, 191(15), 3501–3523.
- Kunze, K., Iwamura, M., Kise, K., Uchida, S., & Omachi, S. (2013). Activity Recognition for the Mind: Toward a Cognitive Quantified Self. *Computer*, 46(10), 105–108.
- Lesaja, S., & Palmer, X.-L. (2020). Brain-Computer Interfaces and the Dangers of Neurocapitalism. *ArXiv*. [2009.07951](https://arxiv.org/abs/2009.07951).
- Ludwig, D. (2014). Extended cognition and the explosion of knowledge. *Philosophical Psychology*, 3, 1–14.
- Lynch, M. P. (2014). Neuromedia, extended knowledge and understanding. *Philosophical Issues*, 24(1), 299–313.
- Lynch, M. P. (2016). *Internet of Us: Knowing More and Understanding Less in the Age of Big Data*. NY, USA: WW Norton.
- Lynch, M. (2021). Understanding and coming to understand. In S. Grimm (Ed.), *Making Sense of the World: New Essays on the Philosophy of Understanding*. Oxford University Press.
- Martins, N., Angelica, A., Chakravarthy, K., Svidinenko, Y., Boehm, F. J., Opris, I., Lebedev, M. A., Swan, M., Garan, S. A., Rosenfeld, J. V., Hogg, T., & Freitas, R. A., Jr. (2019). Human Brain/Cloud Interface. *Frontiers in Neuroscience*, 13, 112.

- Miller, B., & Record, I. (2017). Responsible epistemic technologies: A social-epistemological analysis of autocompleted web search. *New Media and Society*, 19(12), 1945–1963.
- Montmarquet, J. A. (1993). *Epistemic Virtue and Doxastic Responsibility*. Rowman & Littlefield.
- Nam, C., Nijholt, A., & Lotte, F. (2018). *Brain-Computer Interfaces Handbook: Technological and Theoretical Advances*. CRC Press.
- Nguyen C. T. (2020). Echo chambers and epistemic bubbles. *Episteme*, 17(2), 141–161.
- Nozick, R. (1981). *Philosophical Explanations*. Harvard UP.
- Owens, J., & Cribb, A. (2019). ‘My Fitbit Thinks I Can Do Better!’ Do Health Promoting Wearable Technologies Support Personal Autonomy? *Philosophy and Technology*, 32(1), 23–38.
- Pritchard, D. (2015). Intellectual virtue, extended cognition, and the epistemology of education. In J. Baehr (Ed.), *Intellectual virtues and education: essays in applied virtue epistemology* (pp. 113–127). Routledge, Taylor and Francis Group.
- Pritchard, D. (2010). Cognitive ability and the extended cognition thesis. *Synthese*, 175(1), 133–151.
- Pritchard, D. (2018). Neuromedia and the Epistemology of Education. *Metaphilosophy*, 49(3), 328–349.
- Pritchard, D. (2018a). ‘Extended virtue epistemology’. *Inquiry: An Interdisciplinary Journal of Philosophy*, 61(5–6), 632–647.
- Ranisch, R. (2015). Cognitive Enhancement to Overcome Laziness: Ethically Wrong or Just Stupid? *American Journal of Bioethics Neuroscience*, 6(1), 42–43.
- Riggs, W. (2010). Open-mindedness. *Metaphilosophy*, 41(1–2), 172–188.
- Rini, R. (2020). Deepfakes and the Epistemic Backstop. *Philosophers’ Imprint*, 20(24), 1–16.
- Rupert, R. D. (2004). Challenges to the hypothesis of extended cognition. *Journal of Philosophy*, 101(8), 389–428.
- Schwengerer, L. (2021). Online Intellectual Virtues and the Extended Mind. *Social Epistemology*, 35(3), 312–322.
- Simpson, T. W. (2012). Evaluating Google as an Epistemic Tool. *Metaphilosophy*, 43(4), 426–445.
- Smart, P. R. (2012). The Web-Extended Mind. *Metaphilosophy*, 43(4), 446–463.
- Smart, P. (2018). (Fake?) News Alert: Intellectual Virtues Required for Online Knowledge! *Social Epistemology Review and Reply Collective*, 7(2), 45–55.
- Smart, P., Clowes, R. W., & Heersmink, R. (2017). Minds Online: The Interface between Web Science, Cognitive Science, and the Philosophy of Mind. *Foundations and Trends in Web Science*, 6(1–2), 1–234.
- Sosa, E. (2007). *A Virtue Epistemology*. Oxford University Press.
- Steinert, S., Bublitz, C., Jox, R., & Friedrich, O. (2019). Doing Things with Thoughts: Brain-Computer Interfaces and Disembodied Agency. *Philosophy and Technology*, 32(3), 457–482.
- Sullivan Specker, L., & Reiner, P. (2021). Digital Wellness and Persuasive Technologies. *Philosophy and Technology*, 34, 413–424.
- Sunstein, C. R. (2006). *Infotopia: How many minds produce knowledge*. Oxford University Press.
- Wardrope, A. (2014). Authenticity and autonomy in deep-brain stimulation. *Journal of Medical Ethics*, 40(8), 563–566.
- Whitcomb, D., Battaly, H., Baehr, J., & Howard-Snyder, D. (2017). Intellectual Humility: Owning Our Limitations. *Philosophy and Phenomenological Research*, 94(3), 509–539.
- Williams, J. (2018). *Stand out of our light: Freedom and resistance in the attention economy*. Cambridge University Press.
- Zagzebski, L. T. (1996). *Virtues of the Mind: An Inquiry Into the Nature of Virtue and the Ethical Foundations of Knowledge*. Cambridge University Press.
- Zuboff, S. (2019). *The age of surveillance capitalism: The fight for a human future at the new frontier of power*. Public Affairs.

Publisher’s Note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.