



The cognitive and moral harms of platform decay

Michael J. Ardoline¹ · Edward Lenzo²

Published online: 26 July 2025
© The Author(s) 2025

Abstract

Platform decay is the phenomenon of major internet platforms, such as Google search, Facebook, and Amazon, systematically declining in quality in recent years. This decline in quality is attributed to the particular business model of these platforms and its harms are usually understood to be violations of principles of economic fairness and of inconveniencing users. In this article, we argue that the scope and nature of these harms are underappreciated. In particular, we establish that platform decay constitutes both a cognitive and moral harm to its users. We make this case by arguing that platforms function as cognitive scaffolds or extensions, as understood by the extended mind approach to cognition. It is then a straightforward implication that platform decay constitutes cognitive damage to a platform's users. This cognitive damage is a harm on its own; however, it can also undermine cognitive capacities that virtue ethicists argue are necessary for developing a virtuous character. We will focus on this claim in regards to the capacity to pay attention, a capacity that platform decay targets specifically. Platform decay therefore also constitutes both cognitive and moral harm, which simultaneously affects billions of people.

Keywords Platform decay · Virtue ethics · Deskilling · Enshittification · Extended cognition

Introduction

It is widely acknowledged that major internet platforms such as Google search, Facebook, and Amazon have been systematically declining in quality in recent years. This phenomenon is known as “platform decay” or, in Cory Doctorow’s coinage, “enshittification” (Doctorow, 2022). The reasons behind this decline in quality is attributed to the unique role these platforms have as ubiquitous intermediaries between users and other businesses, and its harms are usually understood to be violations of principles of economic fairness and of inconveniencing users (O’Reilly, et al., 2024, Doctorow, 2022, 2023, 2024). We will argue here that the harm to users caused by platform decay is much more extensive than currently appreciated. To do so, we first claim that many platforms should be understood as a form

of cognitive coupling such that some of them constitute cases of extended cognition. If one accepts this, then it is a straightforward implication that platform decay constitutes cognitive damage to a platform’s users. While a harm on its own, this cognitive damage can also undermine some cognitive capacities that virtue ethicists argue are necessary for developing a virtuous character, in particular, the capacity to pay attention. If so, then platform decay also constitutes a moral harm. Thereby, platform decay is not merely an inconvenience to users, nor is it only a concern of economic fairness. Platform decay is a cognitive and moral harm that simultaneously affects billions of people.

Platform decay

Platform decay became a mainstream topic of discussion following journalist and novelist Cory Doctorow’s work on what he calls “enshittification,” (Doctorow, 2022, 2023, 2024). “Platform decay” and “enshittification” are synonyms, though we adopt the former since it specifies the scope of the phenomenon. Platform decay describes the life-cycle of large platforms such as Facebook or Uber and the ways in which their services have declined in quality once

✉ Michael J. Ardoline
michaeldorfline@lsu.edu

Edward Lenzo
edwardlenzo@muhlenberg.edu

¹ Louisiana State University, Baton Rouge, United States

² Muhlenberg College, Allentown, United States

these platforms captured a commanding, often monopolistic market share (see the August 5th, 2024 United States court ruling that Google is an illegal monopoly). This lifecycle begins with offering useful, convenient, and usually free services. These services draw a large user base. Access to this user base then becomes a product that the platform sells to its business customers, usually advertisers. This continues until any meaningful competition disappears (for a more detailed account, see Srnicek, 2016, which analyzes this process of enclosure and rent-seeking in detail). At this point, platforms begin increasing rents to advertisers, cutting costs, going back on promises (e.g., Google no longer promises to “not be evil,” (Allyn, 2021), and re-engineering services from being useful and convenient to be primarily profit generating in ways that undercut their overall usefulness. For example, Google search increasingly returns advertisements instead of useful results, and Amazon’s internal search function ranks sponsored results over relevant ones (Doctorow, 2024). This leaves users with much worse services than they have grown used to, but without meaningful competition that they could use instead. This has “trapped billions of us on platforms that many of us do not like but feel we can’t leave” (Doctorow, 2023). This lack of meaningfully competing platforms also leaves business users, such as those who sell through Amazon, at the mercy of ever increasing rents in order to operate through the service and gain access to users’ attention (O’Reilly, et al., 2024).

This is certainly not a complete political or economic analysis of the situation, nor will we provide one here. There is a spectrum of attempts at such an analysis of the economics of platforms and the tech industry more generally that range from the claim that this situation represents a new evolution and perversion of capitalism (Zuboff, 2019) to the claim that it represents an overthrow of capitalism by a new form of feudalism (Varoufakis, 2024). Similarly, various remedies have been proposed to this situation, including greater anti-trust prosecution and regulation more generally (Shapiro et al., 2024), the formation of alternative coop platforms (Zuckerman, 2024) as well as economic protections for users, in particular end-to-end protections and the right of free exit (Doctorow, 2023). However, our aim here is not to engage with criticism of platforms in the terms of political economy, but to extend this critique’s scope to the cognitive and ethical dimensions of platform decay. Once our argument is established, we will return to these to show that they are narrowly focused on economics and inconvenience, and do not ameliorate the wider harms of platform decay. Therefore, additional remedies will be required.

Platforms as cognitive extensions

To understand the full range of platform decay’s effects on users, we must first establish that they constitute forms of extended cognition. There have been two waves of research in extended cognition. The first wave is marked by Clark and Chalmers’ *Parity Principle*. The second wave responds to critics of this principle by positing the *Complementarity Principle* instead. Both waves contribute to the understanding of platforms as extended cognition.

The Parity Principle asserts that “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 and Chalmers, 1998, p. 8). In other words, the location of information processing is irrelevant to that process’ status as cognitive or not. While consciousness or perceptual awareness may be limited to the head, the active information processing that constitutes cognition is not. Instead, cognition extends to artifacts, the environment, other people, and so on. To illustrate this, they provide the now famous example of Otto and Inga. Both Otto and Inga plan to visit a museum. Inga uses her internal capacities of memory and navigation to do so. Otto, on the other hand, has a memory disorder and instead uses a notebook to remind him of his intention to visit the museum, as well as list directions for how to get there, what exhibits he wants to see, etc. Without the notebook, Otto is unable to perform this cognitive task, but with the notebook, he is able to successfully complete it in a way that is not meaningfully different from Inga. Insofar as the notebook performs the same cognitive function as Inga’s internal memory and navigation capacities, the Parity Principle would consider the notebook as much a part of cognition as the brain of either Otto or Inga. The notebook then extends cognition.

Alongside the Parity Principle, Clark and Chalmers propose three aspects of successful extension: availability, trust, and accessibility. In order to count as a cognitive extension on this view, the artifact must be readily available (Otto’s notebook is portable and on his person at nearly all times). It must be trusted (Otto presumes to be the only one writing in his notebook, and so pre-reflexively endorses what he reads there). That the extension be trusted does not require that the extension itself be trustworthy (Someone else may sneakily write in Otto’s notebook while he sleeps, but this would not be cognitively different than a case of ‘recovering’ false memories). And it must be accessible (Otto is able to open and read the notebook with ease as he goes about his day; it does not require a struggle or much effort to do so). Clark and Chalmers hesitantly offer a fourth criteria: past endorsement, meaning that the information accessed

has previously been accepted by the user. They are unsure of this (Clark and Chalmers, 1998, p. 17). We think past endorsement as a criteria should be rejected as it is either mistaken or redundant given trust of the artifact. If someone sneaks into Otto's room at night and modifies his notebook without his knowledge, in what way could this possibly decouple Otto from the notebook? He did not previously endorse the forged entry in the journal, but he cannot know this. Rather, he retains his trust in the artifact despite the tampering. The second wave of extended cognition that we take up below broadly avoids this issue.

Smart (2012, 2017) argue that on this framework, under certain conditions, the use of the web constitutes extended cognition. It is, however, cautioned that most of the web does not satisfy the accessibility criterion because information can be difficult to find, require scrolling through troves of irrelevant information to find what is useful, and so on. In the years since these articles, the ubiquity of internet use, as well as ease of access through devices like smartphones has increased notably. In addition, the process of enclosure that is central to platform decay leading to users' feelings of being trapped shows the degree to which these services have become a part of many everyday cognitive goings about the world. While we agree that much of the web does not count as extended cognition for the reason Smart (2012, 2017) cite, the ease and availability of platforms (especially via smartphones with mobile data) and the trust often placed in them (say that Google search results are reliable or that one's Twitter feed reflects common sentiment) mean that many of the largest platforms do meet the criteria for being counted as cognitive extensions.

Smart (2017) provides support for this by arguing that we should see not only the device (smart phone or computer), but the internet itself as a cognitive extension. This is because multiple different devices can be used to access the same service or information, and it is the service itself that makes certain cognitive tasks possible rather than the device alone. Without Google search or a similar service, even a smart device with a web browser would be unlikely to meet the availability and accessibility criteria for information on the internet to function as extended memory. We develop this example at the end of this section. We think that Smart's argument that we should not limit extension to devices is correct, and it straightforwardly supports our claim that platforms themselves count as cognitive extensions, not just the devices we use to access platforms. The reader may verify this by examining changes to their cognitive abilities the next time they find themselves with a smartphone without internet access of any kind.¹ This does not mean every

service, website, or even platform counts as an extension. For example, one may be skeptical of Wikipedia articles because of the ability for anyone to edit them, and so while generally reliable, one does not immediately endorse what they read there. If so, then this does not meet the threshold of trust required to count as a cognitive extension.² On the other hand, one is likely to rely on a GPS navigation service while driving in a way that does satisfy the trust criteria (in many cases, it is dangerous not to). We also note that if one does so while driving somewhere new, this is a case of extension without past endorsement.

While we make use of these criteria, they do risk falling into irresolvable debates over what counts as enough trust, just how accessible something needs to be, and the like. It is to avoid interminable debates such as these and those over just how alike to internal processes external processes need to be to satisfy the Parity Principle that a second wave of extended cognition embraced the complementarity approach instead. The complementarity approach eschews necessary and sufficient conditions in favor of a multidimensional analysis of the various ways in which humans and artifacts can be cognitively coupled (Menary, 2006; Sutton, 2010). Complementarity allows for there to be differences between internal and external sources of information processing, but for them to nevertheless count as cognitive because the interactions constitute a functionally coupled system of cognition such that the mind could not function as it does without such couplings. These couplings can range from cognitive scaffolding (which do not necessarily constitute cognition, but make possible forms of cognition that the uncoupled mind alone, if such a thing exists, is otherwise incapable of) to full extension of the sort we discussed above.

Sterelny (2010) argues that we can take "environmental fuels for cognition" (Sterelny, 2010, p. 473) seriously without undermining the inner/outer distinction, as the extended mind hypothesis would have it. He deploys an example of scaffolded digestion to suggest that, while it is undeniable we transform our digestive niche (e.g., through fire and cookware) and that it, in turn, transforms our digestive system (what and how we do and can eat; "we are obligatorily cooks" (Sterelny, 2010, p. 467)), it is implausible that this kind of scaffolding constitutes "extended digestion." What is crucial for Sterelny is the ongoing mutual transformation of internal (bodily) and external (environmental) resources, that is, coupling. He does not reject the extended

¹ Smart (2017) notes that the accessibility criteria may be violated in the smartphone case as the internet connection may not be reliable enough. We think the improvements to wireless internet access since

the article was published has made such reliable access a reality for enough people enough of the time to satisfy the criteria.

² While we agree with Smart here, Ludwig, (2015) argues that much of the web, including Wikipedia, do in fact count as extended cognition. If this is correct, then our claims in the following section are on even firmer ground.

mind hypothesis, but maintains that “extended mind cases are limiting cases of environmental scaffolding” (Sterelny, 2010, p. 465), i.e., those that exhibit especially intense coupling, and as such are of limited explanatory scope. Whether a particular platform counts as extension (we maintain that at least some do), it is certainly the case that platforms such as Google and Facebook at least function as cognitive scaffolding: they transform our cognitive capacities, shaping *how* we think through mechanisms such as offloading, and are in turn shaped by us through clicks, purchasing patterns, personalized algorithms, etc.³

While cognitive coupling is a matter of degree, and thus the borders between scaffolding and full extension are vague, the mutual transformation that characterizes coupling in general is sufficient for many of the claims we make moving forward. We note those cases where the difference makes a difference. We intend “extension” to denote a higher degree of cognitive coupling, whereas scaffolding is the more general case. Our claim that platform decay constitutes cognitive and moral harms applies in either case, though we take it that the degree of harm is proportional to the degree of cognitive coupling, and that at least some platform use meets the threshold for extension.

The claim that platform use constitutes extension is supported by Heersmink and Sutton (2020), which uses a complementarity analysis to evaluate the overall cognitive integration of various web-based services. The results of this analysis vary from service to service, with Wikipedia being judged as low to medium integration to Google search being judged as medium to high (Heersmink and Sutton, 2020, p. 157). We accept this analysis and use Google as one of our paradigm cases. In the use of Google search and similar services, users are more highly cognitively coupled than services that operate more like traditional websites. Google search is a ubiquitous platform, so the degree to which its users form a tightly coupled cognitive system with it is particularly meaningful for our case. Let’s look at this in detail. If cognition is extended, memory is not limited to what we remember of our own experience. Through services like Google search, we are able to “recall” the information stored by others. It is as if Otto and several friends with the same impairment (and similar handwriting) shared a notebook. We should then treat extended information access as analogous to memory. If one is skeptical of this, they may be reducing memory to semantic memory (the recall of specific information) and forgetting the role of transactive memory

(the recall of where and how to access information), that internal human memory is a reproduction rather than a process of accurate recall. Furthermore, empirical studies suggest that widespread internet access is altering the structure of memory (and its neural correlates) itself to rely less on semantic memory and more on transactive memory, effectively meaning we come to remember less specific information and more of how to access or retrieve that information, (see Firth et al., 2019, p. 122-3). If so, then one relies less on recalling specific information they have retained and more on trustworthy artifacts. To be successful, transactive memory requires that these sources and means of access remain reliable. As it becomes a trusted, accessible, and easy to use way of accessing information that we do not recall directly, Google search (in particular, its page ranking algorithm) becomes an extension of our transactive memory. Google search is also a paradigm case of platform decay, so we may now turn to the question of what platform decay means for the users whose cognition these platforms extend.

Extended cognitive decay

Aagaard (2020) raises an important issue relevant to exploring harms caused through cognitive extensions. Extended cognition has generally focused on cases of positive, cooperative, and successful extension. He calls this the dogma of harmony. That this dogma has implicitly structured extended cognition research means there is little work to build on for analyzing cognitive harm (though there are some hints, even in the classic extended mind papers). To address this, Aagaard advises researchers to not take extended cognition as a priori positive or cooperative interaction between humans and artifacts. He offers two examples: the formation of bad habits and cognitive deskilling, both of which are possible results of extending cognition, yet are not the desired ends of human users. These are both present in platform decay, and we now turn to the cognitive harms thereof. Timms and Spurrett 2023 build on this work and provide an account of what they call “hostile scaffolding.” Most work on cognitive extensions treat “benign” cases, or cases where the extension operates in line with the user’s presumed interests or goals. A cognitive extension is considered hostile when it “changes the cognitive demands of a task in ways that undermine the interests of the agent attempting the task, and in doing so serves those of another agent,” (Timms and Spurrett, 2023, p. 5).

In platform decay, we see a particularly interesting case, one where what was a harmonious extension (in Aagaard’s terms) becomes disharmonious and even damaging to the user. There are then two sources of cognitive harm in platform decay. First, there is the decay process wherein some

³ These platforms can be modeled using Sterelny’s multidimensional analysis of coupling as highly trusted, individualized/individualizing, and as shaping both individual and collective activity. Some such platforms, at least, would count as highly coupled on his account, but in line with complementarity approaches we aim to avoid questions of thresholds as far as possible.

features that users have come to rely on, that is, are cognitively coupled with, have notably diminished in quality or disappear altogether. Second, there is the resultant product (the decayed platform) which, because of the specific rent-seeking aspects of platform decay, increasingly takes the shape of hostile scaffolding. These moments are formally separable in that a service could merely get worse, becoming an inefficient or unusable copy of its former self without necessarily becoming hostile in Timms and Spurrett's sense. This would still constitute a cognitive harm on our account, though a lesser one. However, given the economic realities of platform decay, its usual trajectory is decay into hostility. In the case of Google search making advertisements harder to distinguish from search results, the user's action to search is hijacked from their goal of seeking true and reliable information to the platform's goal of increasing profits (we assume that advertising dollars spent does not have a necessary correlation with truth and reliability).

That platform decay constitutes a cognitive harm is a direct implication of platforms being either cognitive scaffolds or extensions. As we saw in the previous section, platforms function as cognitive scaffolding for many if not most of their users. As cognitive scaffolds, if these platforms are worsened, if they become worse at performing the information processing for which they are depended on, then the cognition that they support or extend is impaired. Therefore, platform decay constitutes cognitive damage to the users of that platform. We should not take this lightly. "If we remove the external component the system's behavioral competence will drop, just as it would if we removed part of its brain," (Clark and Chalmers, 1998, p. 8–9). Similarly, in an interview, Clark claims that altering external supports to scaffolding may be equivalent to causing brain damage (Getzels, 2018). Levy, 2007 formulates this concern as the Ethical Parity Principle: "Since the mind extends into the external environment, alterations of external props used for thinking are (*ceteris paribus*) ethically on a par with alterations of the brain," (Levy, 2007, p. 61).⁴ Similarly, Vold, 2018 claims that Otto's "notebook has a cognitive status, and moral status, equivalent to Otto's brain. For this reason, stealing Otto's notebook would be more on par with kicking Otto in the head, causing him serious cognitive damage, than it would be to stealing some property, like his gym bag," (Vold, 2018, p. 497). It is important to note that part of the severity results from the fact that while the artifact may

be replaced (the notebook) if destroyed, the information it contained (Otto's notes) cannot (Levy, 2007, p. 62).

When Google search is altered from providing the most relevant search results to a given query to providing a mixture of advertisements and results (with ads now appearing where top ranked relevant results were previously placed), it is a case of hostile scaffolding. They aim to direct a user away from what may be most relevant (the user's presumed interest) towards whoever has paid Google to appear linked to a given search term (an interest that conflicts with that of the user's). In doing so, a cognitive harm is done. The information processing that the user has come to rely on is made worse. In more extreme cases of platform decay, what once counted as an extension may, when the service degrades, fail to count as one. For example, the ease and reliability Google search once provided satisfied the criteria for a cognitive extension. However, once that service declines in quality to the point that it is no longer clear what is a paid advertisement and what is a relevant and useful search result, the user may no longer seamlessly integrate search results into their cognitive processes nor immediately endorse them, thereby violating both the criteria of trust and accessibility. Such a user then loses an extension and the cognitive capacities that came along with it. This again harms the user's cognition in a way equivalent to damaging or removing parts of their brain against their will.

The intentional worsening of these platforms is the ethical equivalent of tearing out pages of Otto's notebook, or, applying the Ethical Parity Principle, giving its users brain lesions. One could perhaps argue that such worsening is not intentional. In such cases, platform decay would still then be an example of harm through negligence (say, having been entrusted with Otto's notebook and leaving it outside in the rain). However, given that the changes to these services are design choices mandated by a corporation's deliberative decision making structure, these changes are intentional. We do not mean these corporations intend to cause brain damage, but that platform decay itself is intentional and such decay constitutes cognitive harm. Therefore, such harm cannot be written off as negligence. Furthermore, Timms and Spurrett note that hostile scaffolding flourishes when one party has a large amount of control of an environment. It is then unsurprising that hostility increases in platforms after they reach monopolistic or near monopolistic market share. They make the connection noting that the design principles used to keep people gambling in casinos are applied to platform design (to, say, keep Instagram users scrolling through content), turning users' smartphones into "instruments of their [own] possible bespoke exploitation" (Timms and Spurrett, 2023, p. 14).

Now that we have established that in principle platform decay constitutes cognitive harm, the question arises of

⁴ Levy, (2007) proposes two versions of the principle, strong and weak. The strong variant follows from accepting full extension. We do and so make use of that version. The weak version follows from accepting only that the mind is embedded in its environment. If one accepts the weak version and that cognition is scaffolded by platforms, the following arguments still hold, but the severity of the harm is reduced.

what sorts of cognitive harms are being done. We won't undertake here an encyclopedic account of such harms. Instead, we will focus on a paradigm example, and one that allows us to make the connection between cognitive harms and moral harms: the deskilling of attention.

Cognitive deskilling

In the previous section, we made the claim that platform decay constitutes cognitive harm in general. Here an opponent might raise the challenge that this harm is overblown. They could accept both that platform decay is happening and accept the ethical parity principle, so that yes, decay does imply a level of cognitive harm. However, they could make the challenge that this harm is actually offset if we compare the user in their coupled-with-a-decayed-platform state not with their coupled-with-a-not-yet-decayed-platform state, but with their pre-coupled state. If the service is still worth using, then the user must be benefiting from it somehow (let's grant them that we will ignore here the monopolistic effects, economic realities, and the various social pressures to use certain platforms). That is, the decayed coupling must be better than the user's pre-coupled state, and so, taking a long view, there is still a net gain for the user. If so, then the charge that decay causes the ethical equivalent of brain damage would be misplaced because overall, the user has still received a cognitive augmentation, even if it is a lesser one than they are used to. We reject this on the grounds that this challenge misunderstands the nature of cognitive adaptation. In response, we will argue that platform decay constitutes cognitive deskilling. We will then extend this in the next section to argue that because some of these skills are necessary for acquiring virtues, platform decay results in moral deskilling as well.

Generally, concerns over cognitive deskilling arise from the offloading of tasks which humans do internally onto some external artifacts. This idea arises out of the more general discussion of deskilling, coming from economics, which is the process by which the skills traditionally had by workers are offloaded onto machines (usually due to choices made by management), leading to workers no longer needing those skills to perform their labor. These skills then often all but disappear in workers, (Braverman, 1998). Cognitive deskilling is the loss of cognitive abilities, capacities, etc. by the offloading of those activities onto artifacts and other scaffolds for cognition, leading to the loss of these skills in users when they are not coupled to that artifact. For example, GPS use is shown to lead to worse navigational knowledge acquisition compared to paper map use (Münzer et al., 2006). Worries of cognitive deskilling are perhaps as old as technology itself, with the discussion of

writing potentially atrophying memory in Plato's *Phaedrus* being the most well-known ancient example. Deskilling of all varieties became a wide-scale economic concern during the Industrial Revolution, and Marx gives perhaps the earliest account of the political economy of cognitive deskilling in the *Grundrisse*. (Marx 1993). Like with platform decay, we leave the fundamentally important work on the economic aspects of cognitive deskilling to others.

Returning to the lifecycle of platform decay, the first stage is to offer a useful, cheap, and easily accessible service. This service is presumably something that the user cannot do as well on their own. The user then integrates this service into their everyday going about the world, and so become cognitively coupled to the service. In doing so, they are offloading some of their abilities because the service does them better and with greater ease. This leads to atrophy of those skills. Therefore, some of the harms caused by platform decay are cases of cognitive deskilling. Take the shift from semantic memory to transactive memory discussed above. Information recall is a cognitive skill. One becomes worse at recalling specific information, but better at recalling how to find it. However, when the systems that extend transactive memory by facilitating finding that information (such as Google search) are made worse, one is not only cognitively harmed in that their transactive memory is made worse; their semantic memory has also been partially deskilled. This adds severity in that not only is there a harm akin to brain damage in degrading the service from its useful to less useful, decayed forms, the user has also been deskilled. Because of this, their coupling with a decaying platform has made them worse off than they were in their pre-coupled state. Overcoming this either requires the effort of retraining the lost skills or the use of some cognitive extension. Given the monopolistic conditions of platform decay, the latter case is one where the user has then been made more dependent on a worse product!

While deskilling has a negative connotation, not all deskilling is bad. First, not all extension and offloading is deskilling. Whether or not offloading cognitive tasks results in deskilling and the atrophying of important capacities differs from case to case. Writing by hand can often improve memory, *contra* the *Phaedrus*. Also, some skills are worth offloading to free up time and make possible upskilling. Even if writing down equations led to not developing mental math skills, the gain in complexity of the mathematics one is able to do via offloading is worth the trade-off even if we value the ability to do mental math. Furthermore, some skills are simply too context dependent to be worth preserving when their instrumental value vanishes. Moreover, deskilling in general results from external artifacts doing a task at least reasonably well and often better than humans can (think again of GPS use versus unaided navigation or

even paper maps), which is a benefit. This is not to say that these cases do not raise problems that need to be addressed (following the GPS example, see Gillett and Heersmink, 2019). However, the deskilling which results from platform decay is not mitigated by any of these concerns. Platform decay does involve the loss of valuable skills, the skills lost are fundamental to the use and acquisition of other important skills (and do not provide a trade-off in the form of upskilling), and, as we saw above, decaying platforms do not even perform their services reasonably well, yet they nurture user dependence on them. The trade-off may have been worth it when the platform performed its function well, but once they reach a certain point of decay, their users are left worse off than before they started using the platform. It is this lack of mitigating factors and the ethical aspect of the skills in question that make platform decay a clear case of moral harm. We turn to this now.

Moral deskilling

Moral deskilling, like cognitive deskilling, is a wide-ranging and complex topic that extends to technology use in general. However, cases of moral deskilling are much more straightforward in the case of platform decay and more directly make the argument that platform decay harms its users' ethical capacities. That is, there is an inherent ethical harm to platform decay in that the degradation of platforms that have become central to billions of people's cognitive scaffolding damages their virtuous capacities. Vallor, (2015) gives the canonical account of moral deskilling via technological offloading and the relevant concerns for virtue ethics, and we follow her account here. In short, in order to form virtues *qua* positive habits, we must be able to exercise certain skills. The repeated exercise of skills is the process of habit formation. Because certain skills are necessary for certain virtues, these skills can be thought of as having an ethical dimension. In moving from cognitive to moral skills and harms, we want to be clear that this position does not imply that those with cognitive damage or disability are somehow inherently less moral than those without. It is a common position in disability studies that being disabled is morally neutral, whereas disabling someone is a harm (see Barnes, 2016, especially Chap. 5). Our focus here is entirely on such latter cases.

Vallor gives several examples of potential cases of moral deskilling. Consider the case of where a military makes use of autonomous weapons to the extent that their human soldiers no longer face combat directly. In such cases, Vallor argues, these soldiers would have no opportunities to repeatedly exercise the skills that make a good soldier. In particular, without combat experience, they would be

unable to form the virtue of courage proper to a soldier. In an important sense, these would be soldiers in name only. And insofar as making good decisions about combat and war in general requires the virtue of courage, this military would be left with no human beings capable of properly judging the orders given to and actions carried out by their autonomous weapons.

Vallor's examples are instructive, but disanalogous to our case in a few ways. We are less concerned here with the offloading aspect of deskilling and more with how the decline in quality of service and the increase of hostile scaffolding affects virtue formation (we assume most users are interested in being or becoming good people or at least are not interested in using a platform to become a worse person). In other words, we are concerned with ethical deskilling whether or not the skill is actually offloaded onto a platform. In particular, platform decay involves increasing distractions which harm attention (a claim we defend below). It does not matter for our purposes here whether the ability to pay attention is offloaded onto the platform or not (it isn't), but with the fact that the repeated use of such platforms and their role as cognitive scaffolding atrophies the user's skill of paying attention.

Shuster and Lazar (2025) build on Vallor's work to argue that algorithmic recommendation systems pose a danger to attention. Algorithmic recommendation systems are systems for ranking and choosing what and in what order a user encounters information on a platform. They note, as we have, the possibility of moral deskilling through the use of such systems. Furthermore, they give a possible mechanism by which the deskilling of attention in particular happens in terms of classical conditioning. Algorithmic recommendations are designed not just to present information, but to produce engagement (usually measured through clicks or interactions). The goal of increasing engagement is then to get users to interact with more content,⁵ and this content is either enjoyable, anger inducing, or provokes another strong emotional response. This response rewards the user, and so is a form of reinforcement learning that favors multitasking, distraction, impulsive behavior, and other forms of engagement that disrupt attention. This reinforcement learning has been shown to make people more easily distracted, that

⁵ There are some exceptions here, such as Google's "long click," which tracks whether a user returns to Google in a short time after accessing a website. If not, this is a long click. Whether a website regularly leads to long clicks can then be fed back into a recommendation algorithm and used as data to rank that website higher than others that lead to less long clicks. Alternatively, long clicks can be dispreferred so as to increase the number of returns to search, leading to more ad engagement, as Google has been accused of doing recently. Algorithmic recommendations then do not inherently prefer distracting or hijacking content, but the economics of platforms and platform decay heavily favor producing algorithms that do.

is, they become even worse at multitasking than someone who does not regularly multi-task, (see Firth et al., 2019, p. 120-2 for an overview of studies on the effects of multi-tasking on attention). Because attention is a certain component of cognition, this constitutes a cognitive harm. The moral consequences are equally clear: if one is constantly distracted by a precisely calibrated alternation of rage-bait about their political opponents and soothingly cute videos of cats in whimsical situations that they repeatedly scroll past a close friend's cry for help posts, their ability to be a caring friend has been harmed. Furthermore, it has harmed the fundamental skill of paying attention. Shuster and Lazar analyze algorithmic recommendation systems on platforms working relatively well. Yet these systems are all made worse, and so presumably is the damage they may do, under conditions of platform decay.⁶

An incredibly common feature of platform decay is the ever increasing number of distractions built into them. This has spread to all corners of the internet and there are competing theories as to why. On the surveillance capitalism account, this is mainly the result of the economic dependence of most web-based series on advertisements (Zuboff, 2019). This dependence on advertising is particularly pronounced in decaying platforms. Facebook, for example, has continually increased the amount of advertisements and sponsored content shown to users, at the expense of content generated by friends, that is, the people the users actually have social relations with (Doctorow, 2023, though this is also obvious to, and frequently remarked on, by its current users). This increase in advertising means an increase in various actors vying for users' attention through the platform. Each ad is an attempt to distract a user from their current task and redirect them elsewhere, and therefore is directly aimed at disrupting attention. An opposing account to Zuboff is the Algorithmic Attention Rent account found in O'Rielly et al. 2024. This account claims that the control of users' attention is the source of a platform's ability to charge rents to their business customers, whether that attention is directed towards advertisements or not. Because of their algorithmic shaping of information access, these platforms are able to charge access to the scarce resource of user attention. This ability to funnel users' attention to different business customers is essentially the diverting of users away from their intended goals of using a platform (e.g., they are shown the product of the business customer who pays the most despite the user searching for the most relevant product). That is, it is a form of hijacking users' attention in Timms and Spurrett's sense, and this hijacking is necessary to their business model. If so, then platforms, insofar as they remain capable

of extracting rents from business customers, are inherently hostile scaffolding.

Given either of these accounts of the economics of platforms, their functioning is predicated on manipulating user attention in a harmful way, either through distraction or hijacking. As Vallor points out, attention is fundamental to ethical comportment.

"A person who cannot be counted on to pay attention when you tell her about the recent death of your closest friend, or who is unable to stay focused on the grave and imminent danger to which you're trying to alert her... is not someone who can be said to be virtuous. This is true even of a person who makes a sincere effort to pay attention to her social environment but who has unwittingly lost the cognitive ability to succeed in this task," (Vallor, 2015, p. 117).

Similarly, Gardiner (2022) argues that the habit of paying attention to the right things at the right time in the right way is a virtue (which she calls "attunement"). This virtue is a particularly important one in that it plays a role in most virtuous actions. "Attention determines which possibilities a person takes seriously and which environmental features they are sensitive to, monitor for, and neglect," (Gardiner, 2022, p. 50). Insofar as properly exercising a virtue means being aware of contextual features of the situation and recognizing possibilities for good action, attention is central to virtuous action in general. Similarly, forming virtues, and by extension a good character, requires acting in the right ways repeatedly until one forms good habits. Attention is then necessary for virtue cultivation as well. Therefore, harm to one's ability to pay attention constitutes a harm to one's character.

While the harm done to capacities for paying attention is much wider than platform decay, given the role of advertising and the rent-seeking model that is central to platform decay as we currently see it, this harm is of fundamental importance here (that said, we believe these arguments would generalize to all forms of design focused on hijacking and disrupting users' attention). Such distracting hostile extensions increase the likelihood of what (Bruineberg and Fabry 2022) calls extended mind wandering, or mind wandering via an artifact, such as "mindless" scrolling on a smartphone. They argue that extended mind wandering has the costs but not the benefits of unextended mind wandering (Bruineberg and Fabry, 2022, p. 18). These costs include the general decrease in cognitive ability that comes with being distracted, whereas non-extended mind wandering has positive effects on creativity, self-insight, and disinhibition during learning tasks that its extended variant misses out on. Therefore, increasing the likelihood of extended mind

⁶ Understanding platforms as cognitive extensions, which Shuster and Lazar do not explore, is also a reason to consider their harm as more pernicious.

wandering is a moral harm that instills habits which undercut the formation of virtues.

Even useful and well-designed technologies can produce ethical deskilling. Platform decay, however, is a particularly clear-cut case in that common forms of decay not only do not provide opportunities to acquire virtues; they actively undercut the conditions of several important virtues. In virtue ethics terms, the repeated, habitual use of decaying platforms will result either in forming vices, in preventing the formation of core capacities required for the formation and use of virtues, or in undermining cognitive skills required for those capacities. Following the Ethical Parity Principle then, the design of decayed platforms should not be seen as the moral equivalent of offering users a choice in which they chose poorly (and so placing much of the responsibility on the user). Instead, such design changes are the moral equivalent of direct damage to whatever parts of the user's brain allows them to pay attention.

Damage to the capacity to pay attention is the most straightforward and perhaps most fundamental case of moral deskilling, but it is not the only one. For example, consider the recent rush to install LLM-based AI in various platforms despite often limited usefulness and various downsides. While this is not fundamental to platforms in the way advertising or attention control seem to be, the situation is made possible by platform's monopolies (or near monopolies). This AI rush has potentially wide-ranging effects (Heersmink, 2024). Hicks et al., (2024) argue that such AI's must, by definition, be understood to be Frankfortian bullshitters because they are incapable of caring about the truth or falsity of what they say. What this means for the cognitive aspect of virtues in general and specific virtues such as practical wisdom is an open question, but if these services become ubiquitous as cognitive extensions, there is a high risk of deskilling.

Conclusion

We have argued that due to the extended nature of cognition, platform decay constitutes cognitive harm. This harm is both in the worsening of the functions of artifacts that users' cognition has come to depend on and in the cognitive deskilling of users through this dependence on a worse product. Because some of the cognitive skills harmed, in particular that of paying attention, are also moral skills whose repeated exercise is necessary for forming virtues, this deskilling also constitutes a moral harm. Insofar as these services are ubiquitous, they are a part of the cognitive scaffolding of billions of human beings (What did you Google while reading this? How many times did you check Twitter?). Therefore, the degrading of these services is not merely an inconvenience

to its users nor should we only be concerned with economic fairness in regards to monopolistic platforms. We must also reckon with the fact that platform decay constitutes a cognitive and moral harm which is equivalent to low-level brain damage to billions of users. This may sound extreme, but these platforms do indeed serve billions of users and if the ethical parity principle holds, the harms we note are at least analogous to a minor concussion to the majority of those users.

If so, then the proposed ameliorations on offer, such as Doctorow's End-to-end and Right to Exit Principles, Zuboff's Right to the Future Tense and Right of Sanctuary, or O'Rielly et al.'s call for mandatory disclosure of operation metrics, are inadequate. We think these principles are a fine start, but if we are correct about the scale and nature of the harms involved, they do not treat the whole problem or offer suitable demands for redress. Nor do we think, given both the ubiquity of platform usage and their monopolistic conditions, that simply not using various decayed platforms is a reasonable response. This would be to give up on legitimate goods. Before succumbing to decay, Google search was a legitimate marvel on a par with Wikipedia for the good it provided regarding access of information (this is not to absolve pre-decay Google of any corporate misgivings, only to say that the service itself was a good one that greatly benefitted users). Unless the economic situation is radically changed, these platforms are unlikely to have meaningful competitors. And so, to say simply that one should avoid using platforms is to give up on a better world, a world we know is possible because, until a few years ago, we lived in it.

Perhaps responses to moral deskilling in general could help? Wong, (2019) suggests a Confucian response to the problem of moral deskilling in that we can maintain and cultivate the right skills and habits by ritualizing our use of technology in beneficial ways. Wong claims that the problem is not faulty design or user choice, but that these technologies disrupt ritual practices and social relations in a way that causes people to neglect rituals and tradition (Wong, 2019, p. 7–8). This is a promising approach to many forms of ethical deskilling, but not platform decay. Due to the ubiquity of these platforms and their ability to break down social norms (norms are often the unnamed things in the motto "move fast and break things"), it is unlikely that norms in the form of ritual practice alone can ameliorate the harms done by platform decay. Similarly, one might suggest that we change our comportment to platforms such that they no longer function as cognitive extensions and remain only loosely coupled to them. This suggestion ignores (1) the social and economic pressures that make certain platforms ubiquitous, and (2) that it may not be feasible or efficient to use these services in such a way (imagine what it would

require to verify every Google search result such that you can honestly be said to be using it with low trust). That said, some direct approach to the problem of moral deskilling must be a part of the general response to platform decay.

The insufficiency above shows that platform decay must be addressed by and through platforms themselves in both their ethical and economic dimensions. In their ethical dimension, greater ethical concern must be taken in the design of platforms. Here we echo Vallor's own response to the issue of moral deskilling. She counsels that "moral skills are essential prerequisites for the effective development of practical wisdom and virtuous character, and since market and cultural forces are not presently aligned to bring about the more ethical of the ambiguous potentials presented here, the future shape of these developments calls for our closer attention—and perhaps active intervention," (Vallor, 2015, p. 109). This active intervention takes the shape of forefronting ethical design such that platforms are remade to not be hostile, exploitative, or morally deskilling (or at least have a reasonable trade-off of upskilling). However, this alone does not address the economic drivers at the heart of platform decay. As we have left the economic work to others who are more qualified than us, we will also not offer a specific economic prescription. We instead note that there are a variety of proposals that potentially address this. These include Lina Kahn's argument that the monopolies behind these platforms be broken up (Khan, 2019). Alternatively, platforms can be nationalized or run as public utilities, given their relation to the public good. These responses would reshape both the technology at the heart of platforms and the economic incentives by which they are run. In doing so, we would be building artifacts and environments worth extending our minds into.

Author contributions M.A. and E.L. wrote, edited, reviewed, and revised the manuscript.

Data availability No datasets were generated or analysed during the current study.

Declarations

Competing interests The authors declare no competing interests.

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/>.

Works cited

- Aagaard, J. (2020). 4E cognition and the dogma of harmony. *Philosophical Psychology*, 34(2), 165–181. <https://doi.org/10.1080/09515089.2020.1845640>
- Allyn, B. (2021). Ex-Google Workers Sue Company, Saying It Betrayed 'Don't Be Evil' Motto. *NPR*, November 29, 2021, sec. Technology. <https://www.npr.org/2021/11/29/1059821677/google-dont-be-evil-lawsuit>
- Barnes, E. (2016). *The minority body: A theory of disability*. Oxford University Press.
- Braverman, H. (1998). *Labor and Monopoly Capital: The Degradation of Work in the Twentieth Century*. 25th Anniversary ed. edition. New York: Monthly Review Press.
- Bruineberg, J., & Regina Fabry (2022). Extended Mind-Wandering. *Philosophy and the Mind Sciences*, 3, 1–30. <https://doi.org/10.33735/phimisci.2022.9190>
- Clark, A., & David Chalmers (1998). The extended Mind. *Analysis*, 58(1), 7–19.
- Doctorow, C. (2022). Social Quitting. *Medium* (blog). <https://doctorow.medium.com/social-quitting-1ce85b67b456>
- Doctorow, C. (2023). As Platforms Decay, Let's Put Users First. Electronic Frontier Foundation. <https://www.eff.org/deeplinks/2023/04/platforms-decay-lets-put-users-first>
- Doctorow, C. (2024). The 'enshittification' of TikTok. *Wired*. <https://www.wired.com/story/tiktok-platforms-cory-doctorow/>
- Firth, J., Torous, J., Stubbs, B., Firth, J. A., Genevieve, Z., Steiner, L. S., Alvarez-Jimenez, M., et al. (2019). The 'online brain': How the internet May be changing our cognition. *World Psychiatry*, 18(2), 119–129. <https://doi.org/10.1002/wps.20617>
- Gardiner, G. (2022). Attunement: On the cognitive virtues of attention. In M. Alfano (Ed.), *Social virtue epistemology*. Colin Klein, and Jeroen de Ridder. Routledge. 48–83.
- Getzels, P. (2018). Andy Clark - What is extended mind?? *Closer To Truth* <https://www.youtube.com/watch?v=kc-TdMjuJRU>
- Gillett, A. J., & Richard Heersmink (2019). How navigation systems transform epistemic virtues: Knowledge, issues and solutions. *Cognitive Systems Research*, 56, 36–49. <https://doi.org/10.1016/j.cogsys.2019.03.004>
- Heersmink, R., & John Sutton (2020). Cognition and the web: Extended, transactive, or scaffolded?? *Erkenntnis*, 85(1), 139–164. <https://doi.org/10.1007/s10670-018-0022-8>
- Heersmink, R. (2024). Use of large Language models might affect our cognitive skills. *Nature Human Behaviour*, 8(5), 805–806. <https://doi.org/10.1038/s41562-024-01859-y>
- Hicks, M., Townsen, J. H., & Slater, J. (2024). ChatGPT is bullshit. *Ethics and Information Technology*, 26(2), 38. <https://doi.org/10.1007/s10676-024-09775-5>
- Khan, L. M. (2019). The separation of platforms and commerce. *Columbia Law Review*, 119(4), 973–1098.
- Levy, N. (2007). *Neuroethics: Challenges for the 21st Century*. 1st edition. Cambridge, UK; New York: Cambridge University Press, 2007.
- Ludwig, D. (2015). Extended cognition and the explosion of knowledge. *Philosophical Psychology*, 28(3), 355–368. <https://doi.org/10.1080/09515089.2013.867319>
- Marx, K. (1993). *Grundrisse: Foundations of the Critique of Political Economy*. Translated by Martin Nicolaus. Reprint edition. London: Penguin Classics.

- Menary, R. (2006). Attacking the bounds of cognition. *Philosophical Psychology*, 19(3), 329–344. <https://doi.org/10.1080/09515080600690557>
- Münzer, S., Zimmer, H. D., Schwalm, M., Baus, J., & İlhan Aslan. (2006). Computer-Assisted navigation and the acquisition of route and survey knowledge. *Journal of Environmental Psychology*, 26(4), 300–308. <https://doi.org/10.1016/j.jenvp.2006.08.001>
- O'Reilly, T., & Strauss, I., & Mariana Mazzucato (2024). Algorithmic attention rents: A theory of digital platform market power. *Data & Policy*, 6. <https://doi.org/10.1017/dap.2024.1>
- Plato (1997). Phaedrus. In *Plato: Complete Works*, edited by John M. Cooper and D. S. Hutchinson, 506–57. Indianapolis, Ind: Hackett Publishing Co.
- Schuster, N., & Seth Lazar (2025). Attention, moral skill, and algorithmic recommendation. *Philosophical Studies*. <https://doi.org/10.1007/s11098-023-02083-6>
- Shapiro, A., Levitt, M. & Patrick Jarenwattananon (2024). From Amazon to Facebook and Google, Here's How Platforms Can 'Decay.' *NPR*, sec. Business. <https://www.npr.org/2024/01/04/1222040458/amazon-google-facebook-twitter-social-media-online>
- Smart, P. (2012). The Web-Extended Mind. *Metaphilosophy*, 43(4), 446–463. <https://doi.org/10.1111/j.1467-9973.2012.01756.x>
- Smart, P. (2017). Extended cognition and the internet. *Philosophy & Technology*, 30(3), 357–390. <https://doi.org/10.1007/s13347-016-0250-2>
- Srnicek, N. (2016). *Platform Capitalism*. 1st edition. Cambridge, UK; Malden, MA: Polity.
- Sterelny, Kim. (2010) Minds: Extended or Scaffolded? *Phenomenology and the Cognitive Sciences*, 9(4), 465–81. <https://doi.org/10.1007/s11097-010-9174-y>
- Sutton, J. (2010). Exograms and interdisciplinarity: History, the extended mind, and the civilizing process. In R. Menary (Ed.), *The extended Mind* (pp. 189–225). MIT Press.
- Timms, R., & David Spurrett (2023). Hostile scaffolding. *Philosophical Papers*, 52(1), 1–30. <https://doi.org/10.1080/05568641.2023.2231652>
- Vallor, S. (2015). Moral deskilling and upskilling in a new machine age: Reflections on the ambiguous future of character. *Philosophy & Technology*, 28(1), 107–124. <https://doi.org/10.1007/s13347-014-0156-9>
- Varoufakis, Y. (2024). *Technofeudalism: What killed capitalism*. Melville House.
- Vold, K. (2018). Overcoming Deadlock: Scientific and Ethical Reasons to Embrace the Extended Mind Thesis. *Filozofija I Društvo* 29, no. 4: 489–504. <https://doi.org/10.2298/fid1804489v>
- Wong, P. H. (2019). Rituals and machines: A Confucian response to Technology-Driven moral deskilling. *Philosophies*, 4(4), 59. <https://doi.org/10.3390/philosophies4040059>
- Zuboff, S. (2019). *The Age of Surveillance Capitalism: The Fight for a Human Future at the New Frontier of Power*. 1st edition. New York, NY: PublicAffairs.
- Zuckerman, E. (2024). How We've Enshittified the Tech Economy. <https://www.prospectmagazine.co.uk/ideas/technology/63324/how-weve-enshittified-the-tech-economy>

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