



# How to Do Things with Information Online. A Conceptual Framework for Evaluating Social Networking Platforms as Epistemic Environments

Lavinia Marin<sup>1</sup> 

Received: 30 August 2021 / Accepted: 22 July 2022  
© The Author(s) 2022

## Abstract

This paper proposes a conceptual framework for evaluating how social networking platforms fare as epistemic environments for human users. I begin by proposing a situated concept of epistemic agency as fundamental for evaluating epistemic environments. Next, I show that algorithmic personalisation of information makes social networking platforms problematic for users' epistemic agency because these platforms do not allow users to adapt their behaviour sufficiently. Using the tracing principle inspired by the ethics of self-driving cars, I operationalise it here and identify three requirements that automated epistemic environments need to fulfil: (a) the users need to be afforded a range of skilled actions; (b) users need to be sensitive to the possibility to use their skills; (c) the habits built when adapting to the platform should not undermine the user's pre-existing skills. I then argue that these requirements are almost impossible to fulfil all at the same time on current SN platforms; yet nevertheless, we need to pay attention to these whenever we evaluate an epistemic environment with automatic features. Finally, as an illustration, I show how Twitter, a popular social networking platform, will fare regarding these requirements.

**Keywords** Skilled action · Epistemic agency · Social networking platforms · Epistemic environment · Situated cognition · Epistemic sensitivity

---

This article is part of the Topical Collection on *Information in Interactions between Humans and Machines*.

---

✉ Lavinia Marin  
l.marin@tudelft.nl

<sup>1</sup> Ethics and Philosophy of Technology Section, TU Delft, Delft, The Netherlands

## 1 Introduction. Epistemic Concerns with Social Networking Platforms

In recent years, mainstream social networking (SN) platforms such as Facebook, Twitter, Instagram or YouTube have come under heavy criticism for facilitating a variety of epistemic harms such as the spreading of misinformation and disinformation (Vosoughi et al., 2018), radicalising users through exposure to extremist content (Alfano et al., 2018), filtering the content to match only particular worldviews (Gertz, 2019) and entrapping users into epistemic bubbles (Nguyen, 2020). Epistemic harm is understood here as any process that hinders how epistemic agents form true beliefs, justify those beliefs, understand information or upgrade that information to knowledge. It is yet unclear to what extent the various epistemic harms occur as a systematic effect of design choices by SN platform owners or to what extent only certain types of users are affected—for example, those users who already were gullible and were displaying intellectual vices that the interactions will then amplify, or perhaps all users suffer to some extent (Fritts & Cabrera, 2022). Because of this unclarity, it is still an open question to what extent responsibility for epistemic harms on SN platforms belongs primarily with the platforms, with the users or with both parties but with different weights.

One major difficulty in pinpointing the share of responsibility for these epistemic harms stems from not knowing something more fundamental: whether these various examples of epistemic harms have something in common or if they are entirely unrelated. Thus far, approaches to the various epistemic failings on SN platforms have been analysed case by case by pointing out how particular design choices explain the harmful epistemic effects. For example, the spread of mis- and disinformation has been tied to the emotionally charged language in which misinformation is packed (Bakir & McStay, 2018) and to the tendency of personalisation algorithms to make more visible the content written with emotional language (Steinert, 2021); this was then shown to be a problem of designing emotional affordances on social media (Steinert & Dennis, 2022). Other examples of successful conceptual work that clarifies how design choices backfire into harmful epistemic effects are the analyses of the formation of epistemic bubbles on Twitter (Nguyen, 2020) or of ignorance niches online (Arfini et al., 2018), or the problem of networked virtues on social media (Alfano, 2021), etc. These philosophical approaches, while substantiated with empirical data and offering successful explanations, are limited in dealing with epistemic harms, as the analysis tackles one particular problem at a time happening on one platform. Design changes on SN platforms can happen very fast, often unseen to the millions of users who experience their effects unknowingly—see, for example, the Facebook emotional manipulation study (Gertz, 2016). Philosophical analysis will lag behind these fast changes as it can only analyse the harm after these experimental changes were carried out on a mass scale. This is a problem if we consider the undetected epistemic

harms currently happening on various SN platforms or the future epistemic harms still waiting to surprise us when software designers unleash the next innovation. How can we know whether a new platform is harmful to its users as epistemic agents if our conceptual tools are geared towards detecting only epistemic harms that have already happened?

We need an approach that evaluates the overall *epistemic environment* enacted by a SN platform, not only the various epistemic phenomena emerging in it. While the epistemic harms on SN platforms may be disconnected entirely from the environment within which these occur, we have already tried the piecemeal approach with limited results. Meanwhile, looking at the overall epistemic environment of SN networks has not been tried before for several reasons. To understand how an environment influences its agents, we need to look primarily at what agents do and then infer what affordances in the environment shaped that behaviour. Conceptually, it is difficult to single out epistemic actions on SN platforms because when users post, comment or share online, they perform speech acts that are not all truth assertions, for example they may express emotions, be ironic or signal that something is interesting (Arielli, 2018; Marsili, 2021). They may not mean to assert that “X is true”, but nevertheless, other users take it to mean that. Thus, actions become epistemically relevant to others while the initiator can claim no such intent (Rini, 2017). This conceptual difficulty should alert us to the relevance of the epistemic environment: individual actions and intentions are not enough to understand the overall effects except through the analytic lens of environment-user interactions.

To designate something as an epistemic *environment*, it needs to show a feedback dynamic: “the agents act on it and it acts back, i.e., the relational practices might alter the agents’ cognitive and behavioural dispositions” (Badino, 2022, p. 3). We are dealing with an environment whenever we can point to a mutual influence, a back and forth between agents creating something shared, and when the technical infrastructure facilitates and alters their interactions. Hence, it is worthwhile to explore an approach that looks at the epistemic environment enacted by a socio-technical system, in our case, SN platforms. Philosophers of technology already argued that, when a new socio-technical phenomenon emerges, human users and technology co-shape each other: the technology affords new abilities or changes old ones (Feenberg, 1991), while the technology is modified by different uses and emerging norms (Kudina, 2019). However, beyond this remark that technologies and humans co-shape and constitute each other, there remains the puzzle of how particular technologies achieve this.

Although SN platforms were not designed explicitly to be *epistemic* environments, they need to be analysed as such, given their influence on the users’ epistemic actions. Granted, SN platforms are not the first place where one would go to build knowledge. What we do with information found on SN platforms serves various purposes, such as communicating, maintaining social bonds and entertainment (Fuchs, 2014), while knowledge or understanding seems to be an afterthought. We may go on Facebook to enjoy some downtime, but we end up clicking on a link to an article others shared about the politics of climate change. We did not set out

to be informed when opening Facebook, but we ended up thinking about climate action, almost against our will.<sup>1</sup> Even when users do not aim to be informed by SM content, they end up being informed inadvertently about various topics (Lee & Ma, 2012), and some of this information ends up in their memory. In case we do not know anything else about the politics of climate change, at least we remember that Facebook post, and we may mention it in conversations with our friends when the topic comes up. Meanwhile, there are SN platforms such as Twitter or Reddit built entirely around the aim of being informed by others: we follow people and topics that we find interesting to find out what they think and to learn from their interpretation of the world. In choosing whom to follow, we already have some knowledge of the topic at hand, and we are not ignorant blank slates absorbing wisdom from Tweets.

While most massive SN platforms were not constructed primarily as environments for acquiring knowledge, these give rise to spaces where information is exchanged intensively and, under certain conditions, these information exchanges may lead to knowledge formation or to ignorance sharing (Arfini, 2019). Even if many users are not affected directly by SN platforms as epistemic agents, they are affected indirectly when those around take their information from SN platforms and misuse it. If others' epistemic agency is affected by their online interactions on social media, we are feeling collateral effects as well, no matter how intellectually virtuous we may be. The epistemic environment enacted by SN platforms extends beyond the online realm into our offline lives and our relations with others because we are too connected to afford ignoring the epistemic effects of SN platforms. If we wanted to characterise various SN platforms as epistemic environments, in what terms could we do that? This paper provides a conceptual framework for evaluating how social networking platforms fare as epistemic environments for their human users. In the next section, I will explain why situated cognition offers a suitable approach for constructing this framework.

## 2 Operationalising the Epistemic Agency of Users Through a Situated Approach

In explaining the emergence of various epistemic harms on SN platforms, two approaches have been previously used: platform or user-centred. Platform-centred approaches look at design choices implemented and at how these affect the users (Plantin et al., 2018). User-focused perspectives point out how individual users fail to act as epistemic agents—either by giving in to specific intellectual vices, not using their intellectual virtues or lacking necessary skills (Heersmink, 2018). The two approaches differ in what they analyse as the source of the informational ills, while both tend to place more responsibility on the platforms and only secondarily on the users, thus acknowledging that users' actions are not happening in a vacuum. What

---

<sup>1</sup> This effect is called incidental exposure to online news—thanks to an anonymous reviewer for pointing this out.

users can do always depends on the affordances designed into the SN platforms and often on nudges and other subtle behavioural pushes to act in certain ways (Pennycook & Rand, 2021). There is agency on both sides, the users and the platforms, and there is an intertwining of how users affect platforms and vice-versa. The users can affect each other's beliefs; they can hijack the platforms and turn these into places of political activism and resistance (Cammaerts, 2015) for their own purposes while at the same succumbing to nudging and manipulation exerted by the SN platforms (Klenk, 2022). The user's epistemic agency is constrained and shaped by the environments in which they find themselves. But how exactly this constraint is happening needs further clarification, as this is specific to each informational environment.

If the challenge is how to analyse a particular environment in view of its epistemic effects on those interacting within, a situated approach seems well fitted as it considers both the environment and the agents interacting within as connected and affecting each other dynamically. The situated approach, as I am using it here, is inspired by situated cognition (or 4E cognition), a paradigm in cognitive sciences that endorses four main tenets, namely that cognition is embodied, embedded, extended and enacted (Robbins & Aydede, 2009, p. 3).<sup>2</sup> Whenever we know something, the cognitive processes are not happening merely in our brains, we also involve the world outside our skulls to shape tools for knowledge (a notebook, a simulation software), and we come to know things by acting in specific contexts. This approach does justice to the differences between knowers, even those raised in the same environment: what users are able to do with the information ultimately matters more than what they are exposed to.

Although some contextual approaches have emerged recently for evaluating epistemic environments, these are not situated approaches.<sup>3</sup> As a precursor to the situated approach in epistemology, Laura Candiotta has developed the term of epistemic cultures. Candiotta has pointed out that we are not solitary truth-seekers; the limits of what we can know are given by the limits of the environments in which we act (Candiotta, 2022). A situated approach looks at how epistemic dispositions are "context-dependent and practice-oriented" embedded in a particular environment or culture (Candiotta, 2022). Yet context is insufficient to explain the feedback mechanisms which are the mark of a situated cognition approach. We still needed to flesh out what is epistemic about the informational interactions between human agents and SN platforms. The situated approach will look particularly at how an environment promotes or hinders the agents' epistemic agency.

<sup>2</sup> "cognition depends not just on the brain but also on the body (the embodiment thesis) .... cognitive activity routinely exploits structure in the natural and social environment (the embedding thesis) .... [and that] the boundaries of cognition extend beyond the boundaries of individual organisms (the extension thesis)" (Robbins & Aydede, 2009, p. 3).

<sup>3</sup> For example, Blake-Turner had previously coined the concept of "epistemic environment" to designate "the circumstances, resources, and other factors of an epistemic community that determine whether one of its members is in a position to gain positive epistemic statuses" (Blake-Turner, 2020, p. 2). The term "epistemic environment" is context-sensitive but it does not touch upon the relations between epistemic agents acting in said environment; hence, it is not situated.

Agency is a way of acting in view of achieving a specific goal, whereby the behaviour adapts to the possibilities offered by an environment (Desmond & Huneman, 2022, pp. 22–23). Taking this generic definition of agency from situated cognition, I define epistemic agency as acting and fitting one's behaviour to an environment towards achieving some epistemic goals (knowledge, understanding, true belief, justification). This applies to the individual agent but also to groups and artificial agents. To have agency, agents need to be able to meaningfully modify their behaviour in response to the challenges that an environment poses and still achieve their goals. Yet doing things with information can take many forms which are not all conducive to epistemic agency. A Twitter user who never posts original content but systematically retweets whatever stirs one's interest is also managing information, but with no epistemic purposes implied. Exercising epistemic agency hinges on the kinds of information we have at our disposal and how we access it. We are constantly immersed in "informational flows" (Floridi, 2011, p. 32), of which we try to make sense based on our skills and capabilities. Since this immersion happens constantly, we do not process all information that touches us towards knowledge building. Turning information into knowledge is called "upgrading" by Floridi, who describes it as embedding the information "in a network of questions and answers that correctly accounts for it" (Floridi, 2011, p. 268). Briefly, we know that X when we are informed that X is the case and when we can say why X is the case by answering questions about how this piece of information is related to other pieces and how they rely on each other (Floridi, 2014, p. 71).

An epistemically productive environment will allow for users to easily upgrade information to knowledge or arrive at understanding. This model of information as upgradeable to knowledge by the epistemic subject emphasises that humans do not acquire knowledge passively from their environment; humans only have access to information that needs to be integrated with other pieces of information. Thus, two conditions need to be fulfilled to arrive at knowledge: firstly, from the subject's side, the epistemic agent should be able to integrate each piece of information within a justificatory network of account—this can happen consciously and through an effort, or it can be automatic (as is the case with expert knowers who pick up information and seamlessly upgrade it to knowledge due to their advanced epistemic skills); secondly, the epistemic environment needs to afford such an upgrade of information to knowledge by not actively hindering the epistemic agent's efforts or thwarting their skilled actions. The extent to which an environment affords epistemic agency is on a spectrum, ranging from hostile to friendly, and, furthermore, the same environment will be perceived differently by users based on their pre-existing skills. Thus, in order to evaluate the epistemic agency enabled by the environment, we need to evaluate the duo agent-environment taken together. This evaluation is complicated by the fact that not all processes of upgrading to knowledge are straightforward or predictable; oftentimes, the upgrading of information to knowledge happens through fortunate accidents, what is called serendipitous discovery (Copeland, 2019), in which the skills of the epistemic agent alight in a fortunate way with the ways in which the information is presented by the environment, making certain connections easier to make. Thus,

an epistemic productive environment will allow for epistemic agents to make use of their skills of upgrading information to knowledge, be this in an automatic or effortful manner, conscious or unconscious. But, given that designing for such environments can only anticipate a handful of use scenarios, and given that serendipitous discovery is mostly unpredictable, there are serious limitations to designing productive epistemic environments.

However, we should notice that some environments are more hostile than others and some are outright detrimental to any epistemic agents, thus making the matter of epistemic environment design a question of what needs to be avoided, rather than what can be fostered, for example, a censorship system in a dictatorial country: essential information is hidden from the regular citizens, and what gets broadcast through official channels is propaganda. This makes for an information-deprived environment that is also polluted with misleading and false propositions. Another way that the environment can hinder epistemic agency is when agents are showered with too much information and when the criteria between falsity and truth are unclear. Another way of making an environment hostile to knowledge is by having users trapped in epistemic bubbles (Nguyen, 2020). In an epistemic bubble, the epistemic leaders usually reinterpret facts according to their new criteria of what counts as knowledge or evidence, criteria which they do not share with other members in order to promote dependence and discourage critical thinking of the users (Marin & Copeland, 2022). To construct knowledge in such a bubble, one needs to take after the leaders and repeat their methods without understanding why these methods are reliable.<sup>4</sup>

To narrow the concept of epistemic agency for a socio-technical environment, we need to look at the relevant features of the environment that foster and shape the user's agency, also called affordances. Affordances are "possibilities for action provided [...] by the environment" (Gibson, 1977). How well a user is attuned to the environment depends on how well users can detect the affordances "out there" and integrate them into habits.<sup>5</sup> Affordances are not merely binary choices for action; they present to agents the possibilities for various levels of skilled actions (Rietveld & Kiverstein, 2014, p. 326). Affordances are the answer to how an environment allows, encourages or hinders our skilled actions, given that most of our skilled actions are automatic to an extent: "[a]ctions are not necessarily the product of deliberation upon desires and intentions ... as agents we are often motivated by a perceptual grasp of what a given situation 'affords'"

---

<sup>4</sup> Such an epistemic bubble emerged for Russians since their army invaded Ukraine in February 2022. The general Russian public has access to information about war crimes going on in Ukraine or testimonies of Ukrainians that they do not want to "be saved" through the internet. Yet this information is discarded and not taken seriously by many because the environment of propaganda and ideology in which they grew up discourages any objective epistemic norms and epistemic autonomy.

<sup>5</sup> Affordances have both an objective and subjective character since these are concrete features in the environment but can be used only when the agent has the skills and abilities of the agent to perceive them (Fayard & Weeks, 2014). To perceive a door handle as an affordance for opening the door, one must first have learned how to open the door through repeated interactions. A person that never opened a door in her life will not see the door handle as an affordance for opening, merely as an unremarkable thing in the room.

(van Grunsven, 2018, p. 133). As an affordance, a basketball is an object to be thrown in the basket during a game, yet, for the skilled player, the ball can be made to take specific trajectories that a novice could not envision. The expert can achieve more actions with the ball than the novice, and these actions are not accidental. The ball is not exceptional by any standard; it simply affords sophisticated actions because it remains constant and because the expert player developed one's skills through time. In this case, the player's agency lies solely in the agent's ability rather than on the ball, which must be only predictable. Yet imagine a ball that changes shape with every game. This would effectively hinder the players' game since it would be impossible to develop skills with such a ball. Thus, conditions for having epistemic agency lie both in the users' developing skills and in the environment that affords skilled actions.

This paper's main conceptual challenge is to distinguish between affordances that are relevant for epistemic agency and those that are merely incidental. Here, situated cognition approaches can provide a helpful angle. The difference between skilled action and amateurish action lies in how appropriately the agent uses existing information for one's goals and how adaptable the agent is to novelty, such as new uses or competent usage in new contexts (Pavese, 2021), which means that a skilled agent will pick up on the affordances of the environment faster and will be able to act in more skilful ways with these affordances than a novice. For the epistemic environments enacted on SN platforms, we need to look at what range of actions with information users have at their disposal is determined by affordances, whether their skills are available to draw upon for these actions and finally, what actions become habits through repetition. I will explain next why each of these is necessary for exercising epistemic agency.

### 3 Informationally Skilled Action and Epistemic Environments: a Conceptual Framework

Skilled actions<sup>6</sup> with information are central for exercising epistemic agency in a particular environment because these require less effort to make. It is undoubtedly possible to act as an epistemic agent without being skilled, e.g. by having luck, in the process of learning, or by sheer effort. Still, these isolated unskilled actions cannot explain the overall workings of the epistemic environment. From a situated cognition perspective, an environment is successful in fostering specific actions when it taps into the already existing skills of its agents because it makes action easier to take, and agents look for the most effortless ways of acting (Tyler et al., 1979); concerning the informational behaviours, avoiding effort implies a reliance on heuristics

---

<sup>6</sup> I chose skills and not capabilities or virtues because skills are more granular and thus can be analysed easier, and also because skills do not carry with them normative implications. A skills-focused approach is not excluding capabilities or virtues; these can be developed alongside skills. Skills are more basic than virtues and contribute to the possibility of habitual action that leads to building the virtues themselves (Pavese, 2021).



as cognitive shortcuts (Broncano & Carter, 2021) or on signals from others in our community of trusted peers. Having skills in a particular area makes the actions less effortful than amateurish actions. Hence, we need to pay attention, particularly to the features of the environment that will trigger or inhibit the exercise of these skills. When epistemic agents upgrade information to knowledge, they rely on pre-existing skills in dealing with information, skills acquired through formal education and later refined through life-long practice. Skills are “purposeful and goal-directed activities ... learnable and improvable through practice” (Pavese, 2021). Generic informational skills are abilities for capturing, manipulating and evaluating information, for example, “the ability to search, select, and evaluate information in digital media” (van Deursen, 2016, p. 6).

Information scholars distinguish between four broad categories of information skills “information seeking, information retrieval, information management and communication of information” (Barry, 1997, p. 225). All are important for SN users, undoubtedly, but of interest here is how SN platforms as epistemic environments hinder or foster the exercise of all these skills. The kinds of informational skills that an agent can use hinge on the environment where one interacts with information, specifically by the possibilities for action (also known as affordances) designed in the environment and on the social norms around information (Gallagher, 2017, p. 174). It is not solely the agent’s own decision to use or forego one’s skills; this choice is also shaped by the affordances designed or found in the environment. One may have all the skills needed for successful action, but if the environment is hostile or rigid, no actions can be performed. The next question to answer is as follows: what kind of environmental features (affordances) should be in place so that users can meaningfully upgrade information to knowledge? From a situated cognition perspective, affordances are not about offering binary choices for action; they also present to agents the possibilities for various levels of skilled actions (Rietveld & Kiverstein, 2014, p. 326); in other words, one can perform an action in a better or worse manner.

However, with SN platforms, we are dealing with a particularly hyper-dynamic environment, an environment that adapts its content by predicting users’ preferences and this is problematic from a situated perspective. Most if not all SN platforms are algorithmically curated, whereby what information we see and how it is presented to us is decided by algorithms which try to tailor the information to users’ liking so that engagement is maximised in terms of time spent on the platform and activity (the personalisation algorithms). The choice to personalise the information that reaches the users’ feeds has been justified as a necessary feature since, without it, it would be difficult to “preserve individuals’ time, and attention” (Reviglio & Agosti, 2020, p. 1) under the constant state of informational overwhelm online. From this perspective, algorithms are presented as the only guides in the informational storms. No environment in which agents act is ever inert; agents change it through their actions even if the change may be visible only at a geological scale. However, with SN platforms, the environment is hyper-dynamic, and it adapts to the users’ actions immediately: when a Facebook user clicks on a link to a video with cats, within minutes, as one scrolls down the news feed, cats have colonised their news feed. This unusual

reactivity of the environment is problematic.<sup>7</sup> To have epistemic agency, the user needs to act in response to an environment and adapt their behaviour according to their goals—which are irrespective of the epistemic goals of the system. However, to adapt one’s behaviour to an environment, the environment needs to be predictable and react in stable ways. If the goal is for users to have agency, then a predominantly rigid environment is just as bad as a predominantly hyper-dynamic one: the first one would not allow users to adapt behaviours, while the latter would offer no point of reference to adapt to.

When users engage with SN platforms, many do so under the impression that these are merely entertainment media and thus fail to see their epistemic responsibility arising in that context.<sup>8</sup> There is an epistemic responsibility in allowing oneself to be exposed to various pieces of information that one cannot control nor assess, and then for not taking into account how this information becomes a belief without justification. Being a user on a SN platform is like driving on auto-pilot on a highway of information: we keep seeing things thrown at us by personalisation algorithms; often, we cannot distinguish between meaningful and meaningless information, yet we are inadvertently informed (Lee & Ma, 2012) as some of this information stays ingrained in our memory. We have taken our hands off the steering wheel without realising that we are epistemically responsible for where the road takes us: we could be good drivers on the highway of information if only we had an awareness of how dangerous the situation at hand is and used our skills to process it meaningfully.<sup>9</sup> A primary demand for any informational environment fostering epistemic agency is to

---

<sup>7</sup> It is also possible that users may choose to thwart the personalisation features of SN platforms, by deliberately clicking on items they dislike, or by using private browsing each time. Personalisation is not (yet) unavoidable, and users can develop skills to counteract it, especially if they are concerned about their privacy and being behaviorally profiled. I am grateful to a reviewer for this observation linking privacy concerns of users with personalisation algorithms.

<sup>8</sup> An audience’s tendency to consume news content for entertainment purposes has been already analysed in media and communication studies under the term of *infotainment* (Baym, 2008) showing that some of the epistemic problems of SN platforms were already emerging with the classic mass media platforms.

<sup>9</sup> To understand how quasi-automatic systems pose problems for human agency, a parallel with the problem of skilled actions for drivers of self-driving cars may be useful here, since the philosophical discussions in that area have had to grapple with similar problems of integrating humans in an environment with a “mind of its own”. In 2016, a lethal accident happened with a Tesla driver using the “auto-pilot” feature: the car’s sensors failed to detect a white truck against a very bright sky and crashed into it. In the aftermath of the accident, Tesla representatives repeatedly pointed out how it was the driver’s fault for having taken his hands off the steering wheel despite hearing the acoustic signal warning him to put the hands back on the wheel. However, philosophers of technology have pointed out that such a dual-mode driving system (allowing for human control during the auto-pilot phases) does not allow for meaningful human control (Mecacci & Santoni de Sio, 2020). The entire system (driver, car, and environment such as infrastructure) needs to afford two main conditions for meaningful human control: tracking and tracing (Santoni de Sio & Van den Hoven, 2018). For our purposes here, the tracing condition is relevant for users on SN platforms. The tracing condition states that a system under meaningful human control needs at least one agent to understand and assume moral responsibility for the situation in which one is found (Santoni de Sio & van den Hoven, 2018). In the case of the Tesla accident, the driver thought that the self-driving system was more autonomous than he believed, so he relegated his moral responsibility as a driver to the car without adequately understanding the system’s limits and, therefore, the limits of his responsibility.

allow its users to become aware when their actions are epistemically relevant, such that they can choose whether to use their informational skills or not. An additional condition would be that users should have more epistemic agency than the artificial agents in the environment because they get to use their skills in adapting to the environment and not, instead, have the environment adapt to them. These demands can be operationalised as follows into three conditions that need to be fulfilled by the environment: (a) the users need to be afforded a range of skilled actions; (b) users need to be sensitive to the possibility of using their skills; (c) the habits built when adapting to the platform should not undermine the user's skills. I will explain briefly each one.

### 3.1 Ability to Engage in a Range of Skilled Action

Users should be able to effectively use their pre-existing skills, which entails having a range of actions at one's disposal, not merely binary choices. For example, user Bob sees a misleading post on Facebook about vaccines and their side effects. Bob has the skills to evaluate the claims in that post and thinks this is a false claim. Here are the actions at Bob's disposal: to flag the post (which is a binary action) or challenge the post's author through a comment. The latter is more granular and allows Bob to deploy one's skills with more or less effort. Bob may decide it is not worth it, or he could think it is worth it and write a long comment with links to relevant sources. From this ability to use perspective, it seems that users do have some affordances at their disposal to engage in sophisticated actions. Any designed interaction that allows for users' granular input is preferable to the binary yes/no choices. Binary choices are those achievable by one click of a button (share, like, dislike, flag), while granular choices are those with free input (text boxes usually, but also uploading images). Most mainstream SN platforms currently allow for some granular user actions.

### 3.2 Epistemic Sensitivity

Epistemic sensitivity allows the user to discriminate between meaningful situations, when one should act, and meaningless ones. Successful epistemic actions presuppose a certain degree of sensitivity, which I will call epistemic sensitivity for information. Moral sensitivity was the inspiration for this concept. Moral sensitivity is the ability to "interpreting the situation, role taking how various actions would affect the parties concerned, imagining cause-effect chains of events, and being aware that there is a moral problem when it exists" (Rest et al., 1999). When we enter a new situation, we do not know right from the start if it is morally relevant or not, and we enter it with no expectation that we should act as moral agents right at that moment, but we need to pick up on the features of the situation. To illustrate, imagine that you are casually discussing yesterday's game with a co-worker when suddenly she turns sombre and asks your advice about what to do with her senile grandmother, whether she should send her to an elderly home or keep her at home. The situation has become morally relevant because she is

asking for your advice on a moral issue, yet some people lacking moral sensitivity will not notice that, and they will offer the most cost-effective advice without any moral considerations. Moral sensitivity is a capacity we need to train in ourselves, as we are not born with it. Similarly, *epistemic sensitivity* is about picking up when a random piece of information presented to us is relevant to our beliefs, i.e. it may alter them, challenge them or re-enforce them. Not all information surrounding us is equally meaningful. We infuse it with meaning by attempting to connect it with our background knowledge and figuring out whether the information at hand can challenge or enrich our knowledge. But, to do this work, we need to be aware of which pieces of information are worthwhile to investigate and which are simply noise (for us). This is also an intuition, a gut-level reaction, that gets trained by repeated interactions with various pieces of information in different contexts and by seeing how others act with information.

When talking about epistemic sensitivity, it may seem that the entire responsibility for having it falls on the user alone. We use the term sensitive to designate humans, not the environment in which they interact. And yet, there are desensitising environments. A worker on a construction site is desensitised to noises around her, as there is constant drilling and hammering and shouting. She hears the noises, but she filters them out and is even capable of having a normal tone of conversation among the incredible noise going on around her. Meanwhile, if taken outside this loud environment, one can only hope that she will react appropriately to loud noises such as a car honking or a fire alarm. Thus, as a minimal condition, we want an environment not to desensitise the user to the salience of information. The environment cannot do much to sharpen this sensitivity, but at least it should not blunt it. Sensitivity can be blunted by overexposure, for example, moral desensitisation happens through repeated exposure to disturbing information. Desensitisation can occur by overexposure to a specific kind of information, but also to any kind of information, especially on SN platforms which are informationally abundant on any topic. Unfortunately, anything can become a source of boredom or disinterest when consumed in too much quantity, from pictures of genocide to cat videos.

### 3.3 Habit Building

Our informational habits shape the kinds of epistemic agents we are more than the propositional attitudes that we hold. We may have a strong belief about the sentence “climate change is caused by human actions”, and our belief may be justified and count as knowledge, but if we never do anything when we are confronted with climate deniers or with misleading formulations of this sentence, our knowledge does not count for much in the social world. We often allow false or misleading statements to circulate around us, even though we know better. However, is there a university duty to act when misleading sentences are uttered nearby? I cannot settle here this issue which is a broader problem for ethics and social epistemology. Yet an interesting insight from situated cognition is that the decision to act itself is not something that we are aware of most of the time; it is an embodied and automatic

reaction, intuitive and quite often unconscious<sup>10</sup> (Dijksterhuis & Nordgren, 2006, p. 105). The more skilled we are in acting, the more the action becomes unconscious as it falls into the realm of habits. Thus, the decision to act or not when epistemic harms are happening around us is triggered by habit. This means that our primary duty, as epistemic agents inhabiting shared epistemic environments, is to build the right kind of informational habits.

Thus far, I have outlined informational skills as important for epistemic agents. Skills become effective only when turned into habitual actions as a skill we never call upon is almost non-existent. Yet habits go both ways. One can become more skilled in performing towards a goal through repeated interactions, but repeated interactions can also build habits that undermine one's previous skills. We want to build the right informational habits that foster our epistemic agency, while, at the same time, to avoid the harmful epistemic habits that would undermine our skills. Research about habit building for online users has mostly outlined the negative effects of SN platforms on their users, for example digital addiction (Allcott et al., 2021) and loss of attentional control (Williams, 2018) as tokens of negative habits acquired on SN platforms. We know much less about how users build informational habits on SN platforms, as this topic is still underdeveloped. By contrast, the contribution of attentional control to moral agency<sup>11</sup> has been explored widely and should serve as inspiration for social epistemology. Users' attention has been seen as a limited resource that needs safeguarding from digital distractions (Newport, 2019), a way of expressing one's freedom by choosing what to pay attention to (Williams, 2018). The ability to focus one's attention on one thing at a time (monotasking) seems to be important for dealing with information in meaningful ways.

When we act skilfully, we need to bracket out certain perceptions so that we can focus on the action at hand, the equivalent of blocking out the noise around us (van Grunsven, 2021, p. 2). On SN platforms, users' attention can be hijacked by the notifications on a page, flashing with red colours, but also by the users' habit of scrolling further. Imagine that we just saw a post that deserves our attention, we need to enter into the inquiry mode, yet our hand automatically scrolls down the page for the next post. Our gaze follows. The habits of scrolling further and jumping from a snippet of information to another are problematic here for the user's epistemic agency. How did users get to these habits, to twitchy fingers and diagonal reading of posts, always jumping to the next one? User actions were supported by specific affordances that made some paths easier to take than others. In developing these detrimental habits, the platforms were complicit by designing features such as infinite scrolling and other so-called dark patterns designed in the user interaction (Gray et al., 2018). The SN platforms do not hinder users from evaluating the information nor from reflecting on it, but rather their detrimental effect lies in fostering habits of taking in information at a rapid speed that makes evaluation and reflection difficult for users.

<sup>10</sup> "intuition as a gut feeling based on unconscious past experience. intuition, in other words, involves feeling that something is right or wrong, or that A is better than B, while being largely unaware where that feeling came from, or what it is based on" (Dijksterhuis & Nordgren, 2006, p. 105)

<sup>11</sup> In the realm of moral agency, we need to control our attention so that we can relate to others meaningfully as we owe others our attention as precondition to moral perception (Murdoch, 1985).

Thus far, I have proposed that three dimensions are fundamental for evaluating whether a particular environment allows for users as epistemic agents: what range of actions are users afforded; can users be sensitive to the information at hand; are users developing new habits, and, if yes, are these detrimental to their skills. Together, these three dimensions make a space upon which every SN platform can be mapped. These three dimensions are almost impossible to maximise for simultaneously with current SN platforms, and this is a problem that future designers of personalised informational environments will need to tackle. A hyper-dynamic environment that constantly changes its content delivery to match the users' perceived desires will not allow them to adapt to the relevant epistemic features of the environment. At the same time, users develop new habits in their interaction with information, such as scrolling endlessly through information or consuming multiple information streams. These habits are not informational skills as such, but habits formed around information: speedy consumption, distraction, multitasking. Habits yield automaticity and unconscious responses, which would not be a problem if the habits were built on existing skills. However, unskilled habits with information contribute to users' loss of epistemic sensitivity because these new habits render automatic actions that users should pay attention to and evaluate. Thus, a perfect epistemic storm emerges when users are unaware that they are responsible for the information they see and circulate, as their informational habits become more unconscious. In this situation, nobody will assume responsibility for the epistemic harms emerging in this environment since nobody intended for these to happen (except for the creators of bots and the fake users injecting propaganda deliberately). This void of responsibility is what makes epistemic harms so hard to deal with systematically on current mainstream SN platforms.

#### 4 A Case Study: Twitter as an Epistemic Environment

To operationalise the proposed framework, I will take Twitter as an example and show how it would fare on the dimensions mentioned above. Twitter is a social networking platform that became famous because it allows only short posts called tweets of maximum 280 characters. This makes it that most messages are short and often deprived of context, unless the Twitter user chooses to create a thread of tweets instead of packing the content into one tweet. Twitter affords it users' binary as well as skilled actions. Binary actions are retweeting without comments and clicking on the heart-shaped button. Skilled actions with information are the more creative ones: commenting, composing a tweet with craft and uploading images. A skilled Twitter user will formulate pithy and concise statements, easily understandable by the targeted audience, giving the gist of one's opinion with maximum economy of words. This is a compositional skill, similar to that of writing haikus or micro-stories. A user may choose to develop the compositional skills for Twitter, but one would probably be influenced by the feedback received from other users. Twitter's news feed design does not seem to particularly reward this skill. Twitter news feeds are made up of rapid successions of tweets, where one tweet gets visibility for other users in only a few seconds. A shared tweet usually dies in a day after it has been shared and

commented on by many. An uncommented tweet becomes invisible in a matter of minutes. In this economy of competing tweets, the content matters more than the form—the more controversial topic, the more viral the tweet gets—and many users will not see it as worthwhile to work on their compositional skills. Because what gets more visibility is the mostly the controversial information, there is hardly an incentive for regular users to craft pithy tweets or to try and document elaborate ones. There are, of course, exceptions to this. Some experts use Twitter successfully to inform a wider audience about their take on the hot topics, and they offer timely and insightful analyses that would otherwise be lost if these were published in magazines or academic journals. To offer these analyses, the experts will circumvent the Twitter standard format by posting threaded tweets. Thus, Twitter fares neutrally concerning the range of skilled informational actions, allowing for both binary and sophisticated actions. Yet, because user engagement is pursued by exposing users to as many tweets as possibly can fit on a screen, the rational choice for a user is to not invest too much time in crafting tweets, and instead to invest the time in making as many tweets as possible. For most users, unskilled actions seem to be the easiest route to getting as much engagement and visibility as possible.

Meanwhile, the epistemic desensitisation of users is a genuine concern on Twitter. When former president Trump started tweeting in his early days as a politician, most people were outraged by his choice to disregard specific facts or interpretations while favouring others. However, as time passed, people got used to it, and ceased to react, merely noticing that “Trump is at it again”. Desensitisation of Twitter users does not happen simply because they are exposed to too much information but because on Twitter, as on many other massive SN platforms, there emerged a relativisation of epistemic norms and criteria for justification. On Twitter, anyone can say whatever they like, and if it is not factually false or reported as hate speech, it will stay on the platform. Epistemic norms for truthfulness, soundness or relevance are circumvented and seem not to matter anymore. This may be construed as a matter of freedom of speech, and I am certainly not proposing to censor such users, but something like awarding less visibility to the users that repeatedly post unreliable information is feasible and desirable in view of making some users less harmful to their peers.<sup>12</sup> However, we need to be wary of the cumulative effect of random Tweets about anything expressing worldviews that nobody bothers to justify when challenged. These Tweets are public, so it may seem that these are acts of speech contributing to the social sphere. Yet, often, SN users utter things in public yet act as if they were simply speaking in their living room to a bunch of friends (Marin, 2021, p. 369). Twitter users will defend their right to say anything by stating that it is their personal opinion and whoever does not want to hear it should unsubscribe. Imagine if a pupil were to say this to their geography teacher: “I think the world is flat, and it is my opinion. If you do not like it, you can unsubscribe”. This would not fly because, like other institutions of knowledge, the school is responsive to specific standards of truthfulness and can be held accountable for those. But meanwhile, on Twitter (and most other mainstream SN platforms out there), there

<sup>12</sup> I am grateful to one of the reviewers for pointing out this alternative.

is no accountability for what epistemic standards users follow. Users can ignore any requests to justify their speech acts, no matter how outrageous, and can enclose themselves in epistemic bubbles to be shielded from such requests. Trying to be an epistemically sensitive user on Twitter may seem ultimately useless since it is not rewarded socially by one's peers. Ultimately, most users will perceive that there are other institutions for truth-telling out there and that whatever is said on Twitter makes no difference in the "real world".

Concerning the habits that Twitter users have developed, there are many possibilities to explore here, and the list cannot be exhaustive by any means. Everyday habits are the impulse to tweet whatever looks interesting or whatever thought crosses one's mind, scrolling down to refresh the page, or liking and retweeting without reading the content. These habits, however, are not quite relevant epistemically, although the twitches we develop to check our phone for new tweets or the notifications we get on the phone are certainly distracting. One epistemically relevant aspect is whether we can refocus our attention on issues that matter when needed. It seems that yes. On Twitter, users are afforded at least one option to direct their attention at will. If a tweet jumps at them as worthy of engaging with, users have the option to open that tweet in a new window, thus focusing their attention on it. This provides an option to exit the main page loaded with different tweets and to arrive on a page for a tweet alone. By contrast, Facebook does not afford this option: posts are read on page of the news feed, usually triggering the users' reflex to keep scrolling down. However, even with this feature, on Twitter, the user's focus is not entirely controlled at will, the various page elements remain in the user's corner of the eye, and comments and notifications still can fight for their attention. An affordance geared explicitly towards facilitating users' control of attention would be introducing a button that darkens the entire page except for the reading window.

To sum up this brief evaluation, Twitter fares neutrally on the skilled action dimension, negatively on the epistemic sensitivity dimension and neutrally on the habit building. This is because negative habits of scattering attention as compulsive checking are built through the technical device (the smartphone) than by the actual platform. A user with strong informational skills (and other skills which I could not get into here, such as critical thinking, social skills digital literacy skills) and who cares about epistemic norms for justification in everyday life is perfectly safe to join Twitter or any other similar SN platform that fares neutrally on one dimension. At the same time, the user needs to stay alert to how the SN platform reshapes one's informational habits by re-evaluating one's experience every few months. Meanwhile, a user with unformed epistemic standards should probably not join such a platform because one will be dependent on luck to not get entangled in communities that act as epistemic echo chambers (Nguyen, 2020), and that may shape the norms one follows [ref hidden for review]. Table shows the summary of this evaluation on the three dimensions. The table also outlines some examples of design proposals for tackling the epistemic environment and allowing Twitter score positively on all three dimensions.



	Range of skilled action	Epistemic sensitivity	Habit building
<i>Twitter evaluation</i>	0	-	0
<i>Design features of Twitter</i>	Tweeting (280 characters), threading several tweets, commenting, retweets with comments	Overexposure to outrageous tweets and moral emotions expressed; visibility of the controversial tweets; no consequences for overriding epistemic norms	Retweeting or linking without reading, infinite scrolling, notifications, focusing on one tweet at a time
<i>Redesign proposals</i>	Compositional windows that obscure the rest of the page; adding different metrics other than likes, such as rating a tweet in terms of its insight, truthfulness, unconventional thinking, concern for others, evidence-based, criticality or reflection	Ranking of Twitter users by reliability of information tweeted or shared (Rini, 2017), or by how likely they are to engage in inflammatory content. Serendipity button—showing users less visible tweets from people that they do not follow	Reading window that allow a user to see one tweet at a time and focus on it (darkening the rest of the page); retweeting with comments by default (not allowing simple shares); asking the user to take a break after seeing 10 tweets; showing one tweet on a page, without any infinite scrolling possible

Designing for an epistemic environment fostering users' informational skills is not a straightforward matter, as I previously mentioned the design depends on the pre-existing skills of the user. One would need to design for different classes of users with different levels of informational skills, from barely literate to the digitally fluent, thus putting in place classes of affordances allowing for ranges of actions. The design intervention would need to first identify the major pitfalls in every category and then proactively design ways of mitigation. For Twitter, the main issue seems to be epistemic insensitivity which is generated by an accumulation of interactions with users that infringe epistemic norms without any consequences. Tackling this would require some way of qualifying users into some sort of epistemic categories such as trustworthy, careful and fact-checker and signalling this with user badges. This is inspired by a design feature already proposed by Regina Rini (2017), whereby she proposed to classify users with badges based on how trustworthy is the information that they usually share. Using badges to signal the user's level of skills and combining these badges with the ranking of tweets that a user sees on the homepage could ensure some level of epistemic control from the user side. The proposal here is not to automatically rank the tweets with the most trustworthy tweeters first, although this would be conducive to the higher quality of the epistemic environment, but to allow users to select which epistemic feature they want to prioritise when they see tweets. Users would be asked to select from a range of informational skills such as insight, truthfulness, unconventional thinking, concern for others, evidence-based, criticality or reflection, and rank them. Then, the tweets they see would be ranked according to these skills, starting with the most important skill for that user. To ensure this complex system of ranking users, one would also

need user feedback by rating users on the same dimensions. Thus, skilled users in critical thinking that would rate other users on the critical thinking dimension would carry more weight than an anonymous user's evaluation. This system of ranking and rating users and showcasing their skills would ensure that informational skills and expertise are rewarded and seen by others, ultimately generating a desire for other users to get those badges. Finally, concerning the issue of habit building, this is the trickiest one since most of the detrimental habits formed on Twitter are not unique to Twitter and are related mostly with scattering attention and distraction. The question of habits would need to be addressed by designers in a more positive light: what kinds of habits does a user want to build in accordance with their epistemic goals? And how can the platform help one achieve these habits by setting for oneself the goals? The answer to habit building questions would be, again, highly tailored to individuals and their skill levels.

## 5 Discussion: Limitations and Extensions of This Framework

Previous discussions of epistemic harms and failures on SN and SM platforms focused on the content of the information circulated by users assuming that it is problematic to have so much misleading and blatantly false information floating around unchallenged; this was framed as a problem of informational quality (Illari & Floridi, 2014). The situated approach sketched here points out a problem emerging before even considering the quality of the information, namely the problem with the ways in which users are hindered in their capacity to deal with the information that matters to them and their possibilities for meaningful actions with said information. Epistemic harms occur not only when we are confronted with misleading or false information, but in a variety of other situations: when we fail to upgrade meaningful information to knowledge, when we get detrimental habits for handling information, when we become insensitive to what matters or when the epistemic norms become relative. An epistemic friendly environment is one that fosters a user's effective skill deployment or, at least, does not create new habits that induce a user to lose one's skills. If an environment fulfils these minimal criteria, it is conducive to having epistemic agency with information. Users are shaped by their repeated interactions with platforms and, through habituation, become skilled at these informational actions.<sup>13</sup> Depending on how various SN platforms fare on the three dimensions, we can place any informational environment on a spectrum ranging from epistemically fruitful to detrimental for the user's agency. However, a caveat needs to be mentioned: we cannot evaluate a SN platform before it is launched and actually used. In a truly situated and ecological approach, the environment is not out there, but it is constituted by the user's actions and their choices to engage with certain affordances.

---

<sup>13</sup> Many SN users have a variety of skills that are developed to a larger or lesser extent. Designing affordances for the multiplicity of skills levels of users is a challenge that I did not tackle in this paper, but it remains a significant open question that needs to be settled in future research.

In this paper, I have identified three features of the socio-technical environment as necessary for promoting the epistemic agency of users: allowing users to have epistemic sensitivity to information, affording granular actions with information and not building detrimental habits with information. The framework proposed in this paper allows a mapping of the potential for epistemically harmful behaviours as afforded by each platform and thus allows users to select the epistemic environment in which they want to interact. However, the choice of a platform should not be a matter of the users' responsibility alone; collective and coordinated action is still needed to make platforms accountable for their design choices. Many of the current mainstream SN platforms score low on these three dimensions because of various problems such as information overwhelm, engagement optimisation for users and the systematic attention scattering. SN platforms are not environments that we encounter out there; rather, these are designed environments with conscious choices behind each designed interaction. SN platforms are commercial enterprises which try to optimise user engagement usually by sophisticated techniques of capturing the user's attention, through the design of persuasive technologies (Atkinson, 2006). The software companies' reasons for choosing to design for a kind of user experience have to do more with financial profit than with the users' well-being, while also trying to fend off the public pressures from activists<sup>14</sup> and existing regulations. To what extent we should expect these platforms to design purposefully towards maximising users' epistemic agency requires entering into a broader discussion about the politics of SN platforms which has been going on already for some time (Fuchs, 2014; Gilroy-Ware, 2017; Zuboff, 2019). There are strong incentives—financial and political—for keeping the social media platforms working as addiction machines and captivators of attention through persuasive design. Meanwhile, my purpose in this paper was confined to highlight one dimension of user agency which has been previously neglected, epistemic agency through skilled action, and to point out the ways in which designed environments can foster this kind of agency.

**Author Contribution** Not applicable. Single-authored paper.

**Funding** This project has received funding from the European Union's Horizon 2020 research and innovation programme under the Marie Skłodowska-Curie grant agreement no 707404. The opinions expressed in this document reflect only the author's view. The European Commission is not responsible for any use that may be made of the information it contains.

## Declarations

**Ethics Approval** Not applicable. No participants or experiments were used in this paper which is fully theoretical.

**Consent to Participate** Not applicable

---

<sup>14</sup> See for example the Centre for Humane Technology as a visible initiative to make accountable both the digital companies and the legislators: <https://www.humanetech.com/>.

**Consent for Publication** Not applicable.

**Competing Interests** The author declares 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/>.

## References

- Allcott, H., Gentzkow, M., & Song, L. (2021). *Digital addiction* (No. w28936). National Bureau of Economic Research.
- Alfano, M., Carter, J. A., & Cheong, M. (2018). Technological seduction and self-radicalisation. *Journal of the American Philosophical Association*, 4(3), 298–322.
- Alfano, M. (2021). Virtues for agents in directed social networks. *Synthese*, 199(3–4), 8423–8442. <https://doi.org/10.1007/s11229-021-03169-6>
- Arielli, E. (2018). *Sharing as speech act. Versus*, 47(2), 243–258.
- Arfini, S., Bertolotti, T., & Magnani, L. (2018). The diffusion of ignorance in on-line communities. *International Journal of Technoethics (IJT)*, 9(1), 37–50.
- Arfini, S. (2019). *Ignorant cognition*. Springer International Publishing.
- Atkinson, B. M. C. (2006). Captology: A critical review. In D. Hutchison, T. Kanade, J. Kittler, J. M. Kleinberg, F. Mattern, J. C. Mitchell, M. Naor, O. Nierstrasz, C. Pandu Rangan, B. Steffen, M. Sudan, D. Terzopoulos, D. Tygar, M. Y. Vardi, G. Weikum, W. A. IJsselsteijn, Y. A. W. de Kort, C. Midden, B. Eggen, & E. van den Hoven (Eds.), *Lecture Notes in Computer Science. Persuasive Technology* (Vol. 3962, pp. 171–182). Springer Berlin Heidelberg. 10.1007/11755494\_25
- Badino, M. (2022). *Bubbles and chambers: Post-truth and belief formation in digital social-epistemic environments*. <http://philsci-archive.pitt.edu/20235/>
- Bakir, V., & McStay, A. (2018). Fake news and the economy of emotions. *Digital Journalism*, 6(2), 154–175. <https://doi.org/10.1080/21670811.2017.1345645>
- Barry, C. A. (1997). Information skills for an electronic world: Training doctoral research students. *Journal of Information Science*, 23(3), 225–238. <https://doi.org/10.1177/016555159702300306>
- Baym, G. (2008). Infotainment. In W. Donsbach (Ed.), *The international encyclopedia of communication*. John Wiley & Sons, Ltd. <https://doi.org/10.1002/9781405186407.wbieci031>
- Blake-Turner, C. (2020). Fake news, relevant alternatives, and the degradation of our epistemic environment. *Inquiry*, 1–21. <https://doi.org/10.1080/0020174X.2020.1725623>
- Broncano, F., & Carter, J. A. (2021). *The philosophy of group polarization: Epistemology, metaphysics, psychology* (1st). *Routledge studies in epistemology*. Routledge.
- Cammaerts, B. (2015). Social media and activism. *Journalism*, 1027–1034.
- Candiotta, L. (2022). Epistemic emotions and co-inquiry: A situated approach. *Topoi*. Advance online publication. <https://doi.org/10.1007/s11245-021-09789-4>
- Copeland, S. (2019). On serendipity in science: Discovery at the intersection of chance and wisdom. *Synthese*, 196(6), 2385–2406. <https://doi.org/10.1007/s11229-017-1544-3>
- Desmond, H., & Huneman, P. (2022). The integrated information theory of agency. *The Behavioral and Brain Sciences*, 45, e45. <https://doi.org/10.1017/S0140525X21002004>
- Dijksterhuis, A., & Nordgren, L. F. (2006). A theory of unconscious thought. *Perspectives on Psychological Science : A Journal of the Association for Psychological Science*, 1(2), 95–109. <https://doi.org/10.1111/j.1745-6916.2006.00007.x>
- Fayard, A.-L., & Weeks, J. (2014). Affordances for practice. *Information and Organization*, 24(4), 236–249. <https://doi.org/10.1016/j.infoandorg.2014.10.001>

- Feenberg, A. (1991). *Critical theory of technology* (Vol. 5). Oxford University Press.
- Floridi, L. (2011). *The philosophy of information*. Oxford University Press.
- Floridi, L. (2014). Perception and testimony as data providers. In F. Ibekwe-SanJuan & T. M. Dousa (Eds.), *Theories of Information, Communication and Knowledge* (pp. 71–95). Springer.
- Floridi, L., & Illari, P. (Eds.). (2014). *The philosophy of information quality*. Springer International Publishing. <https://doi.org/10.1007/978-3-319-07121-3>
- Fritts, M., & Cabrera, F. (2022). Fake news and epistemic vice: Combating a uniquely noxious market. *Journal of the American Philosophical Association*, 1–22. <https://doi.org/10.1017/apa.2021.11>
- Fuchs, C. (2014). *Social media: A critical introduction*. SAGE.
- Gallagher, S. (2017). *Enactivist interventions: Rethinking the mind*. Oxford University Press.
- Gertz, N. (2016). Autonomy online: Jacques Ellul and the Facebook emotional manipulation study. *Research Ethics*, 12(1), 55–61.
- Gertz, N. (2019). The four Facebooks. *The New Atlantis*, 58, 65–70.
- Gilroy-Ware, M. (2017). *Filling the void: Emotion, capitalism and social media*. Duncan Baird Publishers.
- Gray, C. M., Kou, Y., Battles, B., Hoggatt, J., & Toombs, A. L. (2018, April). The dark (patterns) side of UX design. In *Proceedings of the 2018 CHI conference on human factors in computing systems* (pp. 1–14).
- Heersmink, R. (2018). A virtue epistemology of the Internet: Search engines, intellectual virtues and education. *Social Epistemology*, 32(1), 1–12.
- Illari, P., & Floridi, L. (2014). Information quality, data and philosophy. In L. Floridi & P. Illari (Eds.), *The Philosophy of Information Quality* (pp. 5–23). Springer International Publishing.
- Klenk, M. (2022). (Online) manipulation: Sometimes hidden, always careless. *Review of Social Economy*, 80(1), 85–105.
- Kudina, O. (2019). *The technological mediation of morality: Value dynamism, and the complex interaction between ethics and technology*. <https://doi.org/10.3990/1.9789036547444>
- Lee, C. S., & Ma, L. (2012). News sharing in social media: The effect of gratifications and prior experience. *Computers in Human Behavior*, 28(2), 331–339. <https://doi.org/10.1016/j.chb.2011.10.002>
- Marin, L. (2021). Sharing (mis) information on social networking sites. An exploration of the norms for distributing content authored by others. *Ethics and Information Technology*, 1–10. <https://doi.org/10.1007/s10676-021-09578-y>
- Marin, L., & Copeland, S. (2022). Self-trust and critical thinking online: A relational account. *Social Epistemology* (in press)
- Marsili, N. (2021). Retweeting: its linguistic and epistemic value. *Synthese*, 198(11), 10457–10483. <https://doi.org/10.1007/s11229-020-02731-y>
- Mecacci, G., & Santoni de Sio, F. (2020). Meaningful human control as reason-responsiveness: The case of dual-mode vehicles. *Ethics and Information Technology*, 22(2), 103–115. <https://doi.org/10.1007/s10676-019-09519-w>
- Newport, C. (2019). *Digital minimalism: On living better with less technology*. Penguin.
- Nguyen, C. T. (2020). Cognitive islands and runaway echo chambers: Problems for epistemic dependence on experts. *Synthese*, 197(7), 2803–2821. <https://doi.org/10.1007/s11229-018-1692-0>
- Pavese, C. (2021). Knowledge How. In Edward N. Zalta (Ed.), *The Stanford encyclopedia of philosophy* (2021st ed.). Metaphysics Research Lab, Stanford University.
- Pennycook, G., & Rand, D. (2021). Nudging social media sharing towards accuracy.
- Plantin, J. C., Lagoze, C., Edwards, P. N., & Sandvig, C. (2018). Infrastructure studies meet platform studies in the age of Google and Facebook. *New media & society*, 20(1), 293–310.
- Rest, J. R., Narvaez, D., Bebeau, M. J., & Thoma, S. J. (1999). *Postconventional moral thinking: A Neo-Kohlbergian approach*. L. Erlbaum Associates.
- Reviglio, U., & Agosti, C. (2020). Thinking outside the black-box: The case for “Algorithmic Sovereignty” in social media. *Social Media + Society*, 6(2). <https://doi.org/10.1177/2056305120915613>
- Rietveld, E., & Kiverstein, J. (2014). A rich landscape of affordances. *Ecological Psychology*, 26(4), 325–352. <https://doi.org/10.1080/10407413.2014.958035>
- Rini, R. (2017). Fake news and partisan epistemology. *Kennedy Institute of Ethics Journal*, 27(2S), E-43-E-64. <https://doi.org/10.1353/ken.2017.0025>
- Robbins, P., & Aydede, M. (2009). A short primer on situated cognition. In P. Robbins & M. Aydede (Eds.), *The Cambridge handbook of situated cognition* (pp. 3–10). Cambridge University Press.
- Santoni de Sio, F., & van den Hoven, J. (2018). Meaningful human control over autonomous systems: A philosophical account. *Frontiers in Robotics and AI*, 5, 15. <https://doi.org/10.3389/frobt.2018.00015>
- Steinert, S. (2021). Corona and value change. The role of social media and emotional contagion. *Ethics and Information Technology*, 23(1), 59–68.

- Steinert, S., & Dennis, M. J. (2022). Emotions and digital well-being: On social media's emotional affordances. *Philos. Technol.*, *35*, 36. <https://doi.org/10.1007/s13347-022-00530-6>
- Tyler, S. W., Hertel, P. T., McCallum, M. C., & Ellis, H. C. (1979). Cognitive effort and memory. *Journal of Experimental Psychology: Human Learning and Memory*, *5*(6), 607.
- van Deursen, A. J. A. M. (2016). *Digital skills: Unlocking the information society*. Palgrave Macmillan.
- Van Grunsven, J. (2018). Enactivism, second-person engagement and personal responsibility. *Phenomenology and the Cognitive Sciences*, *17*(1), 131–156. <https://doi.org/10.1007/s11097-017-9500-8>
- van Grunsven, J. (2021). Perceptual breakdown during a global pandemic: Introducing phenomenological insights for digital mental health purposes. *Ethics Inf Technol*, *23*, 91–98. <https://doi.org/10.1007/s10676-020-09554-y>
- Vosoughi, S., Roy, D., & Aral, S. (2018). The spread of true and false news online. *Science (New York, N.Y.)*, *359*(6380), 1146–1151. <https://doi.org/10.1126/science.aap9559>
- Williams, J. (2018). *Stand out of our Light*. Cambridge University Press. <https://doi.org/10.1017/9781108453004>
- Zuboff, S. (2019). *The age of surveillance capitalism: The fight for a human future at the new frontier of power: Barack Obama's books of 2019*. Profile books.

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