

Knowledge and Disinformation

Ideally, we want to resist mis/disinformation but not evidence. If this is so, we need accounts of misinformation and disinformation to match the epistemic normative picture developed so far. This chapter develops a full account of the nature of disinformation. The view, if correct, carries high-stakes upshots, both theoretically and practically. First, it challenges several widely spread theoretical assumptions about disinformation – such as that it is a species of information, a species of misinformation, essentially false or misleading, or essentially intended/aimed/having the function of generating false beliefs in/misleading hearers. Second, it shows that the challenges faced by disinformation tracking in practice go well beyond mere fact checking. I begin with an interdisciplinary scoping of the literature in information science, communication studies, computer science, and philosophy of information to identify several claims constituting disinformation orthodoxy. I then present counterexamples to these claims and motivate my alternative account. Finally, I put forth and develop my account: disinformation as ignorance-generating content.

12.1 Information and Disinformation

Philosophers of information, as well as information and communication scientists, have traditionally focused their efforts in three main directions: offering an analysis of information, a way to measure it, and investigating prospects for analysing epistemic states – such as knowledge and justified belief – in terms of information. Misinformation and disinformation have traditionally occupied the backseat of these research efforts.¹ The

¹ While fully-fledged accounts of the nature of disinformation are still thin on the ground, a number of information scientists and philosophers of information have begun to address the problem of disinformation (Hernon 1995, Skinner and Martin 2000, Calvert 2001, Lynch 2001, Piper 2002, Fallis 2009, Walsh 2010, Rubin and Conroy 2012, Whitty et al., 2012, Karlova and Fisher 2013).

assumption has mostly been that a unified account of the three is going to become readily available as soon as we figure out what information is. As a result, for the most part, misinformation and disinformation have received dictionary treatment: for whatever the correct analysis of information was taken to be, misinformation and disinformation have either been taken to constitute the false variety thereof (misinformation) or the intentionally false/misleading variety thereof (disinformation) – by theorists endorsing non-factive accounts of disinformation – or, alternatively, something like information minus truth (misinformation) or information minus truth spread with an intention to mislead (disinformation) in the case of theorists endorsing factive accounts of information.

This is surprising in more than one way: first, it is surprising that philosophers of any brand would readily and unreflectively endorse dictionary definitions of pretty much anything – not to mention entities with such high practical stakes associated with them – such as mis/disinformation. Second, it is surprising that not more effort on the side of information, communication, and computer scientists has been spent on identifying a correct account of the nature of disinformation given the increasingly high stakes of issues having to do with the spread of disinformation that threaten our democracies, our trust in expertise, our uptake of health provision, and our social cohesion. We are highly social creatures, dependent on each other for flourishing in all walks of life. Our epistemic endeavours make no exception: due to our physical, geographical, and psychological limitations, most of the information we have is sourced in social interactions. We must inescapably rely on the intellectual labour of others, from those we know and trust well, to those whose epistemic credentials we take for granted online. Given the staggering extent of our epistemic dependence – one that recent technologies have only served to amplify – having a correct account of the nature of mis/disinformation, in order to be able to reliably identify it and escape it, is crucial.

Disinformation is widespread and harmful, epistemically and practically. We are currently facing a global information crisis that the Secretary-General of the World Health Organization (WHO) has declared an ‘infodemic’. Furthermore, crucially, there are two key faces to this crisis: two ways in which disinformation spreads societal ignorance: One concerns the widespread *sharing of disinformation* (e.g. fake cures, health superstitions, conspiracy theories, political propaganda, etc.), especially online and via social media, which contribute to dangerous and risky political and social behaviour. Separately, though at least as critical to the wider infodemic we face, is the prevalence of disinformation-generated

resistance to evidence: even when the relevant information available is reliably sourced and accurate, many information consumers fail to take it on board or otherwise resist or discredit it (Klintman 2019) due to the rising lack of trust and scepticism generated by the ubiquity of disinformation. An important pay-off, then, of a correct analysis of the nature of disinformation is an understanding of how to help build and sustain more resilient trust networks. It is urgent that we gain such answers and insights: according to the 2018 Edelman Trust Barometer, UK public trust in social media and online news has plummeted to below 25 per cent, and trust in government is at a low of 36 per cent. This present crisis in trust of corresponds with a related crisis of distrust, in that the dissemination and uptake of disinformation, particularly on social media, have risen dramatically over the past few years (Lynch 2001, Levinson 2017, Barclay 2022).

12.2 Against Disinformation Orthodoxy

In what follows, I will scope the scientific and philosophical literature, identify three very widely spread – and rarely defended – assumptions about the nature of disinformation, and argue against their credentials.

- (1) **Assumption §1:** Disinformation is a species of information (e.g. Shannon 1948, Carnap and Bar-Hillel 1952, Frické 1997, Fallis 2009, 2015, Cevolani 2011, D’Alfonso 2011, Dinneen and Brauner 2015).

These theorists take information to be non-factive and disinformation to be the false and intentionally misleading variety thereof. On accounts like these, information is something like meaning: ‘the cat is on the mat’, on this view, carries the information that *the cat is on the mat* in virtue of the fact that it means that the cat is on the mat. Disinformation, on this view, consists in spreading ‘the cat is on the mat’ in spite of knowing it to be false and with the intention to mislead.

Why think in this way? Two rationales can be identified in the literature, one practical and one theoretical.

12.2.1 *The Practical Rationale*

Factivity doesn’t matter for the information scientist. In the early days of information science, the thought behind this went roughly as follows: for the information scientist, the stakes associated with the factivity/non-factivity of information are null – after all, what the computer scientist/

communication theorist cares about is the quantity of information that can be packed into a particular signal/channel. Whether the relevant content will be true or not makes little difference to the prospects of answering this question.

It is true that, when it comes to how much data one can pack into a particular channel, factivity doesn't make much difference. However, times have changed, and so have the questions the information scientist needs to answer: the 'infodemic' has brought with it concerted efforts to fight the spread of disinformation online and through traditional media. We have lately witnessed an increased interest in researching and developing automatic algorithmic detection of misinformation and disinformation, such as PHEME (2014), Kumar and Geethakumari's (2014) 'Twitter algorithm', Karlova and Fisher's (2013) diffusion model, and the Hoaxy platform (Shao et al. 2016), to name a few. Interest from developers has also been matched by interest from policymakers: the European Commission has brought together major online platforms, emerging and specialised platforms, players in the advertising industry, fact-checkers, and research and civil society organisations to deliver a strengthened Code of Practice on Disinformation (European Commission 2022). The American Library Association (2005) has issued a 'Resolution on Disinformation, Media Manipulation, and the Destruction of Public Information'. The UK Government has recently published a call for evidence into how to address the spread of disinformation via employing trusted voices. These are, of course, only a few examples of disinformation-targeting initiatives. If all of these and others are to stand any chance at succeeding, we need a correct analysis of disinformation. The practical rationale is false.

12.2.2 The Theoretical Rationale

Natural language gives us clear hints as to the non-factivity of information: we often hear people utter things like 'the media is spreading a lot of fake information'. We also utter things like 'the library contains a lot of information' – however, clearly, there will be a fair share of false content featured in any library (Fallis 2015). If this is correct, the argument goes, natural language suggests that information is not factive – there can be true and false varieties thereof. Therefore, disinformation is a species of information.

The first problem with the natural language rationale is that the cases in point are underdeveloped. Take the library case: I agree that we will often say that libraries contain information in spite of the likelihood of false

content. This, however, is compatible with information being factive: after all, the claim about false content, as far as I can see, is merely an existential claim. There being some false content in a library is perfectly compatible with it containing a good amount of information alongside it. Would we still say the same were we to find out that this particular library contains only falsehoods? I doubt it. If anything, at best, we might utter something like: ‘this library contains a lot of fake information.’

Which brings me to my more substantial point: natural language at best cannot decide the factivity issue either way and at worst suggests that information is factive. Here is why: first, it is common knowledge in formal semantics that, when a complex expression consists of an intensional modifier and a modified expression, we cannot infer a type–species relation – or, indeed, to the contrary, in some cases, we might be able to infer that a type–species relation is absent. This latter class includes the so-called privative modifiers such as ‘fake’, ‘former’, and ‘spurious’, which get their name from the fact that they license the inference to ‘not *x*’ (McNally 2016). If so, the fact that ‘information’ takes ‘fake’ as modifier suggests, if anything, that information is factive, in that fake acts as privative: it suggests that it is not information to begin with. As Dretske (1981) well puts it, mis/disinformation is as much a type of information as a decoy duck is a type of duck (see also Mingers (1995) and Floridi (2004, 2005a, 2005b) for defences of factivity). If information is factive and disinformation is not, however, the one is not the species of the other. The theoretical rationale is false: meaning and disinformation come apart on factivity grounds. As Dretske well puts it:

signals may *have* a meaning, but they *carry* information. What information a signal carries is what it is capable of telling us, telling us truly, about another state of affairs. [...] When I say I have a toothache, what I say means that I have a toothache whether it’s true or false. But when false, it fails to carry the information that I have a toothache. (Dretske 1981, 44, emphases in original)

Natural language semantics also gives us further, direct reason to be sceptical about disinformation being a species of information: several instances of dis-prefixed properties that fail to signal type–species relations – disbarring is not a way of becoming a member of the bar, displeasing is not a form of pleasing, and displacing is not a form of placing. More on this below.

- (2) **Assumption §2:** Disinformation is a species of misinformation (e.g. Floridi 2007, 2008, 2011, Fallis 2009, 2015).

Misinformation is essentially false content, and the mis- prefix modifies as 'badly', 'wrongly', 'unfavourably', 'in a suspicious manner', 'opposite or lack of', or 'not'. In this, misinformation is essentially non-information, in the same way in which fake gold is essentially non-gold.

As opposed to this, for the most part, dis- modifies as 'deprive of' (a specified quality, rank, or object), 'exclude', or 'expel from'. In this, paradigmatically,² dis- does not negate the prefixed content, but rather it signals un-doing: if misplacing is placing in the wrong place, displacing is taking out of the right place. Disinformation is not a species of misinformation any more than displacing is a species of misplacing. To think otherwise is to engage in a category mistake.

Note also that disinformation, as opposed to misinformation, is not essentially false: I can, for instance, disinform you via asserting true content and generating false implicatures. I can also disinform you via stripping you of justification via misleading defeaters.

Finally, note also that information/misinformation exists out there, whereas disinformation is us-dependent: there is information/misinformation in the world without anyone being informed/misinformed (Dretske 1981), whereas there is no disinformation without audience. Disinformation is essentially audience-involving.³

- (3) **Assumption §3:** Disinformation is essentially intentional/functional (e.g. Fetzer 2004b, Floridi 2007, 2008, 2011, Mahon 2008, Fallis 2009, 2015).

The most widely spread assumption across disciplines is that disinformation is intentionally spread misleading content, where the relevant way to think about the intention at stake can be quite minimal, as having to do with content that has the function to mislead (Fallis 2009, 2015). I think this is a mistake generated by paradigmatic instances of disinformation. I also think it is a dangerous mistake, in a world of the automated spread of disinformation that has little to do with any intention on the part of the programmer, to operate with such a restricted concept of disinformation. To see this, consider a black-box artificial intelligence (AI) that, in the absence of any intention to this effect on the part of the designer, learns how to and proceeds to widely spreading false claims about COVID-19

² Not essentially, however. Disagreeable and dishonest are cases in point, where the dis- prefix modifies as 'not-'. The underlying rationale for the paradigmatic usage, however, is solidly grounded in the Latin and later French source of the English version of the prefix (the Latin prefix meaning 'apart', 'asunder', 'away', 'utterly', or having a privative, negative, or reversing force).

³ See Grundmann (2020) for an audience-orientated account of fake news.

vaccines in the population in a systematic manner. Intention is missing in this case, as is function: the AI has not been designed to proceed in this way (no design function), and it does not do so in virtue of some benefit or another generated for either itself or any human user (no etiological function). Furthermore, and most importantly, AI is not the only place where the paradigmatic and the analytic part ways: I can disinform you unintentionally (where, furthermore, the case is one of genuine disinformation rather than mere misinformation). Consider the following case: I am a trusted journalist in village V, and, unfortunately, I am the kind of person who is unjustifiably very impressed by there being any scientific disagreement whatsoever on a given topic. Should even the most isolated voices express doubt about a scientific claim, I withhold belief. Against this background, I report on V TV (the local TV station in V) that there is disagreement in science about climate change and the safety of vaccines. As a result, whenever V inhabitants encounter expert claims that climate change is happening and vaccines are safe, they hesitate to update accordingly.

A couple of things about this case: first, this is not a case of false content/misinformation spreading – after all, it is true that there is disagreement on these issues (albeit very isolated). Second, there is no intention to mislead present in the context, nor any corresponding function. Third, and crucially, however, it is a classic case of disinformation spreading.

Finally, consider the paradigmatic spread of conspiracy theories. Their advocates are, paradigmatically, believers, and their intention is to inform rather than mislead. Since spreading conspiracy theories is a central case of disinformation spread – indeed, I submit, if our account of disinformation cannot accommodate this case, we should go back to the drawing board – we need a new account of the nature of disinformation that does not require any intention or function to mislead.

12.3 A Knowledge-First Account of Disinformation

In what follows, I will offer a knowledge-first account of disinformation that aims to vindicate the findings of the previous section.

Traditionally, in epistemology (e.g. Dretske 1981) and philosophy of information alike, the relation between knowledge and information has been conceived on a right-to-left direction of explanation (i.e. several theorists have attempted to analyse knowledge in terms of information). Notably, Fred Dretske thought knowledge was information-caused true belief. More recently, Luciano Floridi's (2004) network theory involves an

argument for the claim that, should information be embedded within a network of questions and answers, then it is necessary and sufficient for it to count as knowledge. Accounts like these, unsurprisingly, encounter the usual difficulties in analysing knowledge.

The fact that information-based analyses of knowledge remain unsuccessful, however, is not good reason to abandon the theoretical richness of the intuitive tight relation between the two. In extant work (Simion and Kelp forthcoming), I have developed a knowledge-based account of information that explores the prospects of the opposite, left-to-right direction of explanation: according to this view, very roughly, a signal s carries the information that p iff it has the capacity to generate knowledge that p .⁴ On this account, then, information carries its functional nature up its sleeve, as it were: just like a digestive system is a system with the function to digest and the capacity to do so under normal conditions, information has the function to generate knowledge and the capacity to do so under normal conditions (i.e. given a suitably situated agent).

Against this background, I find it very attractive to think of disinformation as the counterpart of information: roughly, as stuff that has the capacity to generate or increase ignorance (i.e. to fully/partially strip someone of their status as knower, or to block their access to knowledge, or to decrease their closeness to knowledge). Here is the account I want to propose:

Disinformation as ignorance-generating content (DIGC): X is disinformation in a context C iff X is a content unit communicated at C that has a disposition to generate or increase ignorance at C in normal conditions.

Normal conditions are understood in broadly etiological functionalist terms (e.g. Graham 2012, Simion 2019b, 2021a) as the conditions under which our knowledge-generating cognitive processes have acquired their

⁴ My co-author and I owe inspiration for this account to Fred Dretske's excellent book *Knowledge and the Flow of Information* (1981). While Dretske himself favours the opposite direction of analysis (knowledge in terms of information), at several points he says things that sound very congenial to our preferred account and that likely played an important role in shaping our thinking on this topic. On page 44 of this book, for instance, Dretske claims that '[r]oughly speaking, information is that commodity capable of yielding knowledge, and what information a signal carries is what we can learn from it'. Sandy Goldberg pointed out to me that Gareth Evans as well may well have had something in the vicinity in mind in his *Varieties of Reference* (1982), in the chapter on communication, when he said that we can exploit epistemic principles about knowledge transmission in testimony cases to derive the semantics of the words used in those knowledge-transmitting cases (roughly, the words mean what they must if such knowledge is to be transmitted).

function of generating knowledge. The view is contextualist in that the same communicated content will act differently depending on contextual factors such as the evidential backgrounds of the audience members, the shared presuppositions, extant social relations, and social norms. Importantly, as with dispositions more generally, said content need not actually generate ignorance in the context – after all, dispositions are sometimes masked.

Now, generating ignorance can be done in a variety of ways – which means that disinformation will come in diverse incarnations. In what follows, I will make an attempt at offering a comprehensive taxonomy of disinformation. (The ambition to exhaustiveness is probably beyond the scope of this chapter, or even of an isolated philosophical project such as mine; however, it will be useful to have a solid taxonomy as a basis for a fully-fledged account of disinformation: at a minimum, any account should be able to incorporate all varieties of disinformation we will have identified.⁵) Here it goes:

- (1) Disinforming via spreading content that has the capacity of *generating false belief*. The paradigmatic case of this is the traditionally recognised species of disinformation: intentionally spread false assertions with the capacity to generate false beliefs in hearers.
- (2) Disinforming via misleading defeat. This category of disinformation has the capacity of stripping the audience of held knowledge/being in a position to know via defeating justification.
- (3) Disinforming via content that has the capacity of inducing epistemic anxiety (Nagel 2010). This category of disinformation has the capacity of stripping the audience of knowledge via belief defeat. The paradigmatic way to do this is via artificially raising the stakes of the context/introducing irrelevant alternatives as being relevant: ‘Are you really sure that you’re sitting at your desk? After all, you might well be a brain in a vat’; or ‘Are you really sure he loves you? After all, he might just be an excellent actor, in which case you will have wasted years of your life.’ The way in which this variety of disinforming works is via falsely implicating that these error possibilities are relevant in the context when in fact they are not. In this, the audience’s body of evidence is changed to include misleading justification defeaters.

⁵ See Simion (2019a, 2021a, 2021b), Simion and Kelp (2022), and Kelp and Simion (2023a) for knowledge-centric accounts of trustworthiness and testimonial entitlement. See Kelp and Simion (2017, 2021) for functionalist accounts of the distinctive value of knowledge.

- (4) Confidence-defeating disinformation. This has the capacity to reduce justified confidence via justification/doxastic defeat: you are sure that your name is Anna, but I introduce misleading (justification/doxastic) defeaters, which gets you to lower your confidence. You may remain knowledgeable about p : 'my name is Anna' in cases in which the confidence lowering does not bring you below the knowledge threshold. Compatibly, however, your knowledge – or evidential support – concerning the correct likelihood of p is lost: you now take/are justified to take the probability of your name being Anna to be much lower than it actually is.
- (5) Disinforming via exploiting pragmatic phenomena. Pragmatic phenomena can be easily exploited to the end of disinforming in all of the ways above: true assertions carrying false implicatures will display this capacity to generate false beliefs in the audience. I ask: 'Is there a gas station anywhere near here? I'm almost out of gas.' And you reply: 'Yeah, sure, just one mile in that direction!'; knowing perfectly well that it's been shut down for years. Another way in which disinformation can be spread via making use of pragmatic phenomena is by introducing false presuppositions. Finally, both justification and doxastic defeat will be achievable via speech acts with true content but problematic pragmatics, even in the absence of generating false implicatures.

What all of these ways of disinforming have in common is that they generate ignorance – by generating either false beliefs, knowledge loss, or a decrease in warranted confidence. One important thing to notice, which was also briefly discussed in the previous section, is that this account, and the associated taxonomy, is strongly audience-involving, in that disinformation has to do with the capacity to have a particular effect – generating ignorance – in the audience. Importantly, though, this capacity will heavily depend on the audience's background evidence/knowledge: after all, in order to figure out whether a particular piece of communicated content has the disposition to undermine an audience in their capacity as knowers, it is important to know their initial status as knowers. Here is, then, on my view, in more precise terms, what it takes for a signal to carry a particular piece of disinformation for an audience A :

Agent disinformation: A signal r carries disinformation for an audience A wrt p iff A 's evidential probability that p conditional on r is less than A 's unconditional evidential probability that p and p is true.

What is relevant for agent disinformation with regard to p is the probability that p on the agent's evidence. And A 's evidence – and, correspondingly, what underlies A 's evidential probability – lies outwith A 's skull: it consists in probability raisers that A is in a position to know. Recall the account defended in Chapter 7:

Evidence as knowledge indicators: A fact e is evidence for p for S iff S is in a position to know e and $P(p/e) > P(p)$.

In turn, we have seen that, on this account, a fact e being such that I am in a position to know it has to do with the capacity of my properly functioning knowledge-generating capacity to take up e :

Being in a position to know: S is in a position to know a fact e if S has a cognitive capacity with the function of generating knowledge that can (qualitatively, quantitatively, and environmentally) easily uptake e in cognisers of S 's type.

This completes my account of disinformation. On this account, disinformation is the stuff that undermines one's status as a knower. It does so via lowering their evidential probability for p – the probability on the p -relevant facts that they are in a position to know – for a true proposition. It can, again, do so by merely communicating to A (semantically, pragmatically, etc.) that not- p when in fact p is the case. Alternatively, it can do so by (partially or fully) defeating A 's justification for p , A 's belief that p is the case, or A 's confidence in p .

One worry that the reader may have at this point goes along the following lines: isn't the account in danger of over-generating disinformation? After all, I might be wrong about something I tell you through no fault of my own; isn't it harsh to describe me as thereby spreading disinformation? Furthermore, every true assertion that I make in your presence about p being the case may, for all I know, serve as (to some extent) defeating evidence for a different proposition q , which may well be true. I truthfully tell you it's raining outside, which, unrelatedly and unbeknownst to me, together with your knowledge about Mary not liking the rain, may function as partial rebutting defeat for 'Mary is taking a walk' – which may well, nevertheless, be true. Is it now appropriate to accuse me of having thereby disinformed you? Intuitively, that seems wrong.⁶

Here are also a couple of parallel cases from Sandy Goldberg (in conversation): say that S is widely (but falsely) thought to be an inveterate

⁶ Many thanks to Sandy Goldberg, Julia Staffel, and Martin Smith for pressing me on this.

liar, so that whenever S says that p , everyone immediately concludes that $\neg p$. Prior to encountering S, A has a credence of 0.2 in p (this is what A's evidence supports prior to encountering S's testimony). S testifies (truly) that p , and A, who, like everyone else in the entire community, takes S to be a liar, drops her credence in p to 0.1. If S's reputation as a liar is assumed to be common knowledge, such that everyone knows of it and would update accordingly, it seems that S's true testimony would meet my analysis of agent disinformation. Conversely, one can imagine cases in which S says something false, explicitly aiming to disinform, but others (who know of S's lying ways) come to true conclusions on the basis of the fact that S said so. On my account, this will not count as a case of disinformation.

Three things about these cases: first, note, once more, that intentions don't matter for disinforming. As such, restricting disinforming via defeat to intentional/functional cases will not work for the same reasons that created problems for the intention/function condition on disinformation more broadly – we want an account of disinformation to be able to predict that asserters generating doubt about, for example, climate change via spreading defeaters to scientific evidence, even if they do it without any malicious intention, are disinforming the audience.

Second, note that it is independently plausible that, just as any bad deed can be performed blamelessly, one can also disinform blamelessly; if so, given garden variety epistemic and control conditions on blame, any plausible account of disinformation will have to accommodate non-knowledgeable and non-intentional instances of disinformation. Conversely, like with all intentions, intentions to disinform can also fail: one may aim to disinform and fail to do so. Indeed, it would be theoretically strange if an intention to disinform will be analytically successful.

Finally, note that we don't need to restrict the account in order to accommodate the datum that disinformation attribution, and the accompanying criticism, would sound inappropriate in the cases above. We can use simple Gricean pragmatics to predict as much via the maxim of relevance: since the issue of whether Mary was going for a walk was not under discussion, and nor was it remotely relevant in our conversational context, flat out accusing you of disinforming me when you assert truthfully that it's raining is pragmatically impermissible (although strictly speaking true with regard to Mary's actions).

Going back to the account, note that, interestingly, on this view, one and the same piece of communication can, at the same time, be a piece of information and a piece of disinformation: information, as opposed to

disinformation, is not context-relative. Content with knowledge-generating potential (i.e. that can generate knowledge in a possible agent) is information. Compatibly, the same piece of content, in a particular context, can be a piece of disinformation insofar as it has a disposition to generate ignorance under normal conditions. I think this is the right result: me telling you that p : 99 per cent of Black people at Club X are staff members is me informing you that p . Me telling you that p in the context of you inquiring as to whether you can give your coat to a particular Black man is a piece of disinformation since it carries a strong disposition (due to the corresponding relevance implicature) to generate the unjustified (and maybe false) belief in you that this particular Black man is a member of staff (Gendler 2011).

Finally, and crucially, my account allows that disinformation for an audience A can exist in the absence of A's hosting any relevant belief/credence: (partial) defeat of epistemic support that one is in a position to know is enough for disinformation. Even if I (irrationally) don't believe that vaccines are safe or that climate change is happening to begin with, I am still vulnerable to disinformation in this regard in that I am vulnerable to content that has, under normal conditions, a disposition to defeat epistemic support available to me that vaccines are safe and climate change is happening. In this, disinformation, on my view, can generate ignorance even in the absence of any doxastic attitude – by decreasing closeness to knowledge via defeating available evidence. This, I submit, is a very nice result: in this, the account explains the most dangerous variety of disinformation available out there – disinformation targeting the already epistemically vulnerable.

12.4 Conclusion

Disinformation is not a type of information and disinforming is not a way of informing: while information is content with knowledge-generating potential, disinformation is content with a disposition to generate ignorance under normal conditions in the context at stake. This way of thinking about disinformation, crucially, tells us that it is much more ubiquitous and hard to track than it is currently taken to be in policy and practice: mere fact-checkers just won't do. Some of the best disinformation detection tools at our disposal will fail to capture most types of disinformation. To give but a few examples (but more research on this is clearly needed): the PHEME project aims to algorithmically detect and categorise rumours in social network structures (such as X (formerly Twitter) and Facebook)

and to do so, impressively, in near real time. The rumours are mapped according to four categories, including 'disinformation, where something untrue is spread with malicious intent' (Søe 2016). Similarly, Kumar and Geethakumari's project (2014) had developed an algorithm that ventures to detect and flag whether a tweet is misinformation or disinformation. In their framework, 'Misinformation is false or inaccurate information, especially that which is deliberately intended to deceive [and d]-isinformation is false information that is intended to mislead, especially propaganda issued by a government organization to a rival power or the media' (Kumar and Geethakumari 2014, 3). In Karlova and Fisher's (2013) diffusion model, disinformation is taken to be deceptive information. Hoaxy is 'a platform for the collection, detection, and analysis of online misinformation, defined as "false or inaccurate information"' (Shao et al. 2016, 745). Examples targeted, however, include clear cases of disinformation such as rumours, false news, hoaxes, and elaborate conspiracy theories (Shao et al. 2016).

It becomes clear that these excellent tools are just the beginning of a much wider effort that is needed in order to capture disinformation in all of its facets rather than mere paradigmatic instances thereof. At a minimum, pragmatic deception mechanisms, as well as evidential probability-lowering potentials, will need to be tracked against an assumed (common) evidential background of the audience.