



Scientific deceit

Stephen John¹ 

Received: 19 December 2017 / Accepted: 9 November 2018
© The Author(s) 2018

Abstract

This paper argues for a novel account of deceitful scientific communication, as “wishful speaking”. This concept is of relevance both to philosophy of science and to discussions of the ethics of lying and misleading. Section 1 outlines a case-study of “ghost-managed” research. Section 2 introduces the concept of “wishful speaking” and shows how it relates to other forms of misleading communication. Sections 3–5 consider some complications raised by the example of pharmaceutical research; concerning the ethics of silence; how research strategies—as well as the communication of results—may be misleading; and questions of multiple authorship. The conclusion suggests some more general conclusions.

Keywords Deceit · Value free ideal · Research ethics · Group authorship · Social epistemology

Communication is the lifeblood of science. Scientists communicate their research findings to one another, to policy-makers, to publics, through journal articles, presentations, blogs, and informal chats. A rich philosophical tradition studies which communicative norms promote the best communal epistemic outcomes (Kitcher 1993; Zollman 2007). However, scientific speech acts are more than conduits for the flow of information. They are governed by normatively rich rules and expectations; fraudulent reports and incompetent reports have similar epistemic effects, but the first are ethically wrong whereas the second are not. Two kinds of obligations constrain and guide proper scientific speech. First, general negative obligations incumbent upon all speakers; for example, not to lie, or not to cause unnecessary harm. Second, role-specific obligations, which fall on scientists, but not others: scientists are obliged to declare “conflicts of interest” whereas advertisers are not. Although any account of scientists’ communicative obligations will be controversial, there are plausible candidates for inclusion on any list. For example, the “Universal Ethical Code for Scientists”

✉ Stephen John
sdj22@cam.ac.uk

¹ Department of History and Philosophy of Science, Free School Lane, University of Cambridge, Cambridge CB2 3RH, UK

(Department for Innovation, Universities and Skills 2007), promulgated by the UK Government, enjoins scientists, “do not knowingly mislead, or allow others to be misled, about scientific matters”, thereby suggesting a negative obligation—do not mislead—and a positive, role-specific obligation—to prevent or combat misleading scientific speech. In turn, these obligations follow from an obligation incumbent on all speakers not to deceive. What, though, is scientific deceit, and how does it relate to everyday deceit? How do concerns about deceit relate to decisions not to speak at all? And can we even make sense of deceit at all in contexts involving collaborative research and authorship?

In this paper I analyse scientific deceit in terms of a phenomenon I call “wishful speaking”. This analysis has implications both for the ethics of science communication in general, and for debates over the proper role of non-epistemic values in scientific practice. Section 1 further motivates my project through the case-study of ghost-managed pharmaceutical research; although such practices seem deceitful, there are significant difficulties in understanding them using our everyday concept. The rest of the paper responds to these problems. In Sect. 2, I sketch a general account of scientific deceit in terms of wishful speaking, analysing its (dis)similarities to everyday forms of deceit. In Sect. 3, I turn to problems arising from failures to publish research, arguing that phenomena such as publication bias can be ethically problematic regardless of their net epistemic consequences. In Sect. 4, I introduce the concept of “proleptic wishful speaking” to understand not only problems in the communication of research findings but in practices and strategies of research itself. Finally, Sect. 5 shows how the theoretical tools introduced in the previous sections can dissolve worries about identifying the locus of responsibility in an era of radically collaborative research.

Before going on, two caveats may help the reader. First, my arguments suggest that we can use the ethics of communication to shed a new light on what might seem like distinctively epistemic problems. At the same time, however, I remain committed to the increasingly unfashionable view that we can, and should, distinguish between the epistemological propriety and ethical propriety of scientific research projects. Unfortunately, I lack the space for a full-blown defence of this foundational commitment. I do, however, hope that my arguments suggest an unexpected route for its defence: as important for maintaining *ethical*, rather than *epistemic*, integrity. Second, my arguments draw on a broadly Kantian account of the ethics of deception, centred on the notion of disrespect for autonomy. However, Kant’s own first-order views on the ethics of lying are notoriously restrictive, and, at a more foundational level, one might worry that notions of dignity and respect fail to provide a fully convincing account of the wrongfulness of deceit. I agree! However, I suggest that, despite its flaws, a broadly Kantian account of the ethics of communication captures important aspects of how we think about deceit in general—not least because of the historical influence of Kant’s work—and therefore, is an obvious starting-point for thinking about scientific deceit specifically. Of course, it is not the only—or maybe even the best—starting-point, but threshing out these foundational issues is beyond the scope of this paper. I stress, then, that my conclusions are both modest and provisional: I am not asserting how we *must* think about scientific deceit, but exploring how we *might*.

1 Motivating the problems

Some authors think that we can understand scientific communication on a model of intellectual (and social) equals in on-going conversation (Fricker 2002). However, a model which might perhaps fit the Seventeenth Century Royal Society may be ill-suited to the social-epistemic landscape of modern science. Not only are there often high non-epistemic stakes at play, but shifting groups of actors come together to write papers aimed at multiple audiences with minimal opportunities for dialogue. It is unclear whether familiar ethical norms, such as injunctions against deceit, can be transposed from everyday contexts to such contexts. This section makes these concerns vivid by introducing a case study, ghost management of pharmaceutical research.

A 1999 paper in the *Journal of the American Medical Association* (Flanagin et al. 1998) suggested that 19% of the articles published in six leading medical journals in 1996 had “honorary authors”—they listed as “authors” people who did not meet guidelines for being listed—and 11% had “ghost authors”—there were researchers whose contributions qualified them as “authors”, but who were not listed. Other studies confirm that honorary and ghost authorship are endemic in medical research. There are many causes of this phenomenon (Horton 2002), but it is particularly prevalent in pharmaceutical industry sponsored research. For example, 75% of reports of pharmaceutical studies in Scandinavia had ghost authors (Gøtzsche et al. 2007).

Not only do the authorship patterns of pharma- and non-pharma-funded research differ, but so do the results reported. Healy and Cattell (2003) compared 96 studies on the efficacy and safety of Sertraline, 55 of which were, ultimately, funded and controlled by the manufacturer, Pfizer. All 55 of these articles reported positively. The other 41 were far more mixed. Again, these findings are not unusual, nor confined to low prestige journals: for example, studying the *British Medical Journal*, Kjaergard and Als-Nielsen (2002) found that articles sponsored by the pharmaceutical industry tended to be far more positive about the funder’s products. Underlying the patterns in pharma-funded research is a mini-industry of “Contract Research Organisations” (CROs), sub-contracted by larger corporations. As Sergio Sismondo (2009, 2011) has documented, these organisations “ghost manage” research processes: they help identify which questions should be addressed; they identify target journals and manage the timing of submissions; they help industry-based researchers to write up papers, before smoothing the process whereby certain authors’ names, typically those from industry, are replaced by those of compliant academics.

Ghost managed research has spawned a small, but rich philosophical literature. Broadly, that literature can be arranged along two dimensions: first, whether it treats ghost managed research primarily as ethically problematic or as epistemically problematic; second, whether it treats ghost managed research as distinctively problematic, or as symptomatic of (or a contributor to) broader problems in biomedicine. On the ethical side, Ben Alamassi (2014) argues that ghost authorship undermines patients’ informed consent and Daniel Hicks (2014) argues that pharma-industry researchers have ethically improper aims. By contrast, Justin Biddle (2007) argues that some ghost managed research leads to epistemically flawed results. At a more general level, Kukla (2012) discusses how ghost authorship points to epistemic issues arising from

changing authorship norms, whereas Stegenga (2011) discusses how the asymmetric publication pattern relates to the reliability of meta-analysis in biomedical research.

The concerns these authors raise are important, and I discuss several of them below. However, ghost-managed research practices seem problematic for a reason these authors do not discuss: they involve deliberately convincing audiences of ill-established claims. They seem deceitful. Although the wrongness of deceit is linked to epistemic concerns, it seems that ghost managed research would be wrong even if it had no effect on the overall stock of true belief; because it abuses audiences' trust. This concern is also separable from other ethical concerns about researchers' motivations or negative effects: a claim is deceitful, and, hence *prima facie* ethically wrong, even if it is made for good reasons and has positive consequences.

Unfortunately, developing the claim that ghost-managed research is deceitful is difficult, for three inter-related reasons. First, standard cases of deceit involve making claims which are epistemically deficient in some way, but, typically, industry-sponsored papers are "scientifically sound", in the sense that they meet the standards of biomedical science (Sismondo 2011). It is failure to report results, rather than the results which are reported, which are problematic. But scientists routinely fail to publish results, and such choices are often ethically fine; no-one is deceived by a failure to publish another model for calculating cancer risk. What differentiates permissible from deceitful silence?

Second, consider what Kukla (2012) calls "micro inductive risk" problems. In many research contexts, researchers may have scope to make methodological choices, i.e. choices about how to research which are not constrained by prevalent epistemic standards. For example, guidelines and common practice may dictate that studies on exposure to dioxins must involve experiments on rats, but not specify which species. Researchers' non-epistemic goals may, then, play an important role not only in how researchers communicate their research, but how they research at all; for example, researchers may decide which rat species to use on the basis of the results they hope to obtain. There is good evidence that biomedical researchers employed by the pharmaceutical industry sometimes make methodological choices this way (Wilholt 2009). These features of industry sponsored research certainly seem problematic. Concerns about deceit, however, seem narrowly limited to the question of how we communicate results, not how they are established in the first place. As such, they seem to be silent on these features of ghost-managed research. Obviously, one concept need not provide an over-arching account of every aspect of some phenomenon. However, there is something troubling about the fact that the concept of deceit is silent on micro-inductive risk choices, because such choices seem to involve a breakdown between justification and communication, whereby a desire to *assert* some claim structures *research* into that claim.

Third, standard models of deceit assume a single speaker, with an intention of causing false belief in her audience. Ghost managed papers, however, are typically the work of multiple "authors" and "contributors", making it hard to identify an individual "speaker". As Kukla points out, something similar is true of nearly all modern scientific research papers. But absent speakers, there is no deceit. Even worse, without speakers talk of the ethics of communication is empty. This is not to say that the ethics of science communication should focus on individual vice and virtue, rather than institutional

structures. However, claims that the institutional structures governing bio-medical publishing are problematic are concerned with how far actual speakers are incentivised, encouraged or enabled to perform the kinds of speech acts which ideal agents should perform. For structures governing speech to be problematic there must be speakers whose speech is structured.

I have only sketched some problems here. I return to them later. Still, the general tension should be clear: on the one hand, we think of scientific communication in ethical terms; on the other, our everyday concepts do not map well onto the communicative structures of science. Either we need a complex account of concepts such as scientific deceit or we need to rethink our ethical outrage. I now develop the first response. First, I outline an account of (relatively) straightforward cases of scientific deceit in terms of “standard wishful speaking”. I then argue that, even if ghost-managed research is not a straightforward case of deceit, it involves phenomena—“extended wishful speaking”, “socially-distributed wishful speaking” and “proleptic wishful speaking”—which share ethically salient characteristics with the paradigm concept, showing how we might meaningfully use ethical language to think about complex scientific speech.

2 Standard wishful speaking

Let us start with simpler cases of scientific deceit. Consider a hypothetical example: an epidemiologist knows that her evidence is not sufficient, relative to norms of epidemiological inference, to justify the claim “this vaccine is safe”. However, she also knows that important non-epistemic benefits would follow were parents to believe that claim (even an unsafe vaccination would do more good than harm). Therefore, she asserts the non-established claim, in the hope of maintaining herd immunity. Such speech certainly seems *prima facie* wrong, because deceitful, even if does have good consequences. What, though, makes it deceitful?

First, the scientist’s claim is not well-established relative to the epistemic standards of the scientific community. By epistemic standards, I mean principles regulating the type and degree of evidential support claims must enjoy before they can be “accepted”, i.e. treated as assumptions in further theoretical or practical reasoning (John 2015). Second, the scientist’s speech is motivated by perlocutionary knowledge, knowledge of the likely consequences of others believing what she says, and a desire to secure those consequences.¹ Neither of these conditions is independently sufficient for scientific speech to count as deceitful. Scientists may make ill-established claims for all sorts of reasons—inference is difficult—but such honest mistakes, while bad scientific speech, are not deceitful. Scientists make scientific claims for all sorts of reasons linked to the (foreseeable) consequences of doing so: to gain fame or fortune, to impress funders or to sell drugs. However, such acts are not necessarily deceitful, as long as what the scientist says is well-established. (Indeed, plausibly, one trick of a well-ordered set of scientific institutions is to ensure that, whatever a scientist’s personal motivations,

¹ The term “perlocutionary knowledge” derives from Austin’s speech act theory (Austin 1975). Franco (2017) has developed a more general account of scientific assertion in terms of speech act theory. I note below how my account differ from his.

she is encouraged or incentivised to assert claims only when they are well-established (Kitcher 1993). What makes the epidemiologist's speech act deceitful is that she makes an *ill*-established claim *because* she wishes to secure certain perlocutionary outcomes. I will call such behaviour "wishful speaking":

A scientist engages in wishful speaking when she makes a claim which is not well-established relative to relevant epistemic standards for accepting claims, and where her motivation for making that claim is the predicted non-epistemic benefits that follow from others believing (or believing that she believes) it, regardless of its truth.

I will now discuss two ways in which even standard wishful speaking in simple cases differs from everyday forms of deceit. These differences will be important when we turn to the far more complex cases exemplified by ghost-managed research.

All deceit involves some kind of misrepresentation (Eggers 2009). In everyday cases of deceit, a deceiver misrepresents her own mental states. For example, in lying, one misrepresents one's own beliefs; in packing a suitcase when one does not plan to go on holiday, one misrepresents one's intentions. On my proposed account, scientific deceit does not involve misrepresenting a speaker's mental states, but, rather, misrepresenting the relationship between some proposition, some evidence and some set of epistemic standards. Why hold that scientific deceit differs in this regard from everyday deceit? Consider two cases. Some professional geologists have strong Creationist convictions (Newton 2011). However, such a geologist speaks properly *as a scientist* when she reports claims about the deep past which are, in fact, well-established, even if she doesn't strictly *believe* them. Conversely, imagine a climate scientist who sincerely believes that climate change is worse than feared, but, as yet, the evidence for that claim has not been peer-reviewed. If the epistemic standards for assertion require that evidence has been publicly validated, then she acts wrongly in stating *in her capacity as a scientist* that climate change is worse than feared, even if she believes this (indeed, even if this claim is likely to be true).² What matters for assessing whether a claim is appropriate is whether it meets public standards, not whether it is sincerely believed (or sincerely believed with a high degree-of-certainty or known).³

Scientific speech is, then, an instance of what Sorensen (2012) calls "role-based assertion", where the propriety of some assertion is relative to the social role of the speaker. What the individual scientist believes is irrelevant to her speech, much as it is irrelevant to an actor's success in performing the role of Romeo that truly he hates Juliet.⁴ Framing scientific deceit as a failure to meet one's role responsibilities provides an explanation of why scientific deceit may seem more problematic than other forms of deceit, say, that of the used car salesman. All forms of deceit are *prima facie* wrongful because they involve treating others' beliefs as a mere means to an end. Scientific

² This is a simplified version of debates over whether the Intergovernmental Panel on Climate Change should use non-peer-reviewed studies (see O'Reilly et al. 2012).

³ These arguments relate to Lackey's work on testimony (Lackey 2007), and to norm-based accounts of trust in experts (see Nickel 2013). I explore these links in John (2018).

⁴ This is one point at which this paper differs from Franco, who, following Searle, treats sincerity as a necessary condition for proper assertion (Franco 2017, p. 172). My discussion in Sect. 4 suggests why this distinction is important for understanding radically collaborative research contexts. See, also, John (2018).

deceit, however, also involves a violation of social expectations: in general, audiences are willing to defer to scientists because they expect that, in virtue of their role-specific obligations, scientists will report claims which they have epistemic reasons to accept (whereas they do not necessarily expect the same of used car salesmen). In deceiving, the scientist not only disrespects others, but violates this trust. In what follows, I will assume that scientists have a strong *prima facie* obligation not to deceive. Of course, I do not mean that all cases of scientific deceit are all-things-considered impermissible. That is a question for the ethics of deceit in general (see Saul 2012, for rich discussion), but I will assume that scientific deceit is *normally* impermissible.

These comments on epistemic trust suggest a second distinction between scientific deceit and everyday deceit. When we think about cases of everyday deceit, we typically focus on cases where agents make “bald” assertions, say “it is raining”. By contrast, scientific findings are often framed in “hedged” terms, reporting the epistemic likelihood of some claim relative to some body-of-evidence. For example, our hypothetical epidemiologist is unlikely to say “the vaccine is safe”, but “according to our evidence, it is very likely the vaccine is safe”. I suggest that this feature of scientific communication has implications for the relationship between lying and misleading in the scientific context. (Before doing so, it is important to note that scientists do sometimes make bald claims. For empirical research, see Chiu et al. (2017) and Vinkers et al. (2015). However, when scientists are less careful, the issues I explore below are irrelevant; the following epicycle of argument applies only when scientists do make hedged claims).⁵

Sometimes, a consequence of a scientist making a hedged claim is that some audience member(s) act on that claim as if it were a bald claim; for example, on the basis of the hedged claim “according to our evidence, it is very likely the vaccine is safe”, parents might act on the bald claim, “the vaccine is safe”. Sometimes the relevant bald claim will be false; the vaccine will turn out to be unsafe. Cases where the initial scientific claim was, itself, beyond reproach, are tragic, but the unavoidable consequence of the fallibility of scientific inquiry; as such, they do not raise questions about the ethics of deceit. Consider, however, cases where a scientist makes a hedged claim which is false—for example, the evidence does not, in fact, strongly support the claim that the vaccine is probably safe—with the intention that audiences come to act on some corresponding bald claim—for example, that the vaccine is safe—and where, in fact, the corresponding bald claim is false. In this case, the belief which the scientist intends the audience to form is not the claim which, strictly, she asserts; as such, it might seem that, strictly, she is not guilty of wishful speaking. However, this result is ethically preposterous. It would imply that scientists could disown responsibility for *any* consequences of their claims—regardless of how ill-established they are—through the fancy casuistical claim that her audience inferred a claim beyond what she strictly *said*. For ethical purposes of thinking about scientific deceit, then, we should not draw a strong distinction between what scientists strictly-and-literally say and the claims which they imply (which is why my definition of wishful speaking is framed in terms of the claims which scientists “make”, hence allowing that these claims may extend beyond the claims they strictly “assert”).

⁵ I am grateful to an anonymous referee for pressing me on this issue, and much of the rest of this section.

These comments relate to a familiar tradition in the ethics of deceit, which distinguishes between lying, where a speaker asserts a claim which is (she believes) false, and misleading, where a speaker asserts a claim which is (she believes) true but where she intends that her hearer will draw a false inference (Saul 2012; Webber 2013; Rees 2014). Some writers argue that misleading speech is less ethically egregious than downright lying, because a misled hearer is (partially) responsible for her false beliefs in a way in which a lied-to speaker is not (Webber 2013). I have argued that, in cases of scientific deceit, there is no important ethical difference between what is said and what is implied. Why? First, there are good reasons to doubt the ethical salience of hearers' responsibility anyway (see Rees 2014). Second, even if the distinction is salient in some cases, it only makes sense when there is rough epistemic symmetry between speakers and hearers, such that hearers *can* (reasonably be expected to) identify misleading speech. Such epistemic symmetry is missing in most cases where scientists speak to non-scientists; even if non-experts have the technical training to question scientists' claims, which is unlikely, they may lack the institutional or financial power to do so. Therefore, regardless of the salience of the distinction in everyday cases, the ethics of science communication should not draw a strong ethical distinction between what scientists strictly *say* and what they *imply*.⁶

In this section, I have sketched a general framework for thinking about scientific deceit, and how it differs from everyday cases. This has important consequences for first-order debates. For example, accusations of insincerity, commonly encountered in debates over climate change, may divert attention from the more important questions of whether standards have been met (Lane 2014). Second, scientists' attempts to avoid responsibility for the outcomes of their speech acts by claiming public misinterpretation of what they strictly *said* may be legally sound, but morally problematic. For example, we cannot exonerate the seismologists who made claims about the risks of the 2009 earthquake in l'Aquila from charges of irresponsible behaviour simply by investigating whether their claims were, strictly-and-literally, true.⁷ Third, they provide us with a way of understanding more specific rules. For example, conflicts of interest policies concern cases where "professional judgment concerning a primary interest (such as research validity) tends to be unduly influenced by a secondary interest (e.g. motive)" (Resnik and Elliott 2013). Understanding "judgment" in terms of the scientist's public speech act, rather than her inner judgment, and "secondary interest" in terms of some perlocutionary end, we can interpret such policies in terms of a concern that conflicts might lead to "wishful speaking". However, my arguments so far have assumed a single speaker, shorn of social context. Do they help us in more complex cases such as contemporary biomedical research?

⁶ This is the second key point on which my analysis differs from Franco's account of scientific assertion. Franco treats cases where audiences fail to grasp the strict meaning of scientists' claims as a form of "misfire" (Franco 2017, p. 177). This rests on an assumption that the locutionary content of an illocutionary act is what was (strictly) "said". I suggest that, for ethical reasons, we should instead expand the scope of what was said to include what Franco would view as perlocutionary consequences of speech. Again, this may seem a trivial difference, but it is important for tracing the ways in which speech may be deceitful, as opposed to being wrong because of its other effects.

⁷ My thoughts here are directly influenced by Anthony Woodman's discussion of this case in a (sadly unpublished) M.Sc. essay (Woodman submitted).

3 Socially extended wishful speaking

Clearly, in ghost-managed research, scientific claims are made because they further non-epistemic interests (of the drugs companies). However, as I noted above, much—perhaps all—scientific assertion is guided by concern for perlocutionary consequences—be it promoting social welfare, esteem or financial reward—but this doesn't necessarily render the relevant assertions unethical.⁸ To show ghost managed research is deceitful, we need to show that claims made for non-epistemic reasons fail to meet conventional epistemic standards. Although some cases of industry-sponsored research clearly violate epistemic norms (Biddle 2007), in general, the vast majority of industry-sponsored papers abide by the methodological norms prevalent in biomedical research. Indeed, industry-sponsored papers are, often, “better” than non-industry funded papers: they more closely adhere to standard procedures, are more clearly written, use more sophisticated statistical tools, and so on.⁹ CRO executives boast about their impressive acceptance rates at top journals (Sismondo 2009). Of course, tools for assessing “scientific soundness” are crude. However, there is no reason to believe that, at the level of individual papers, industry-sponsored research is systematically of lower epistemic quality than non-industry-funded research. (Note that I return to some complexities around the notion of “micro inductive risk” in the next section). As such, it is unclear that industry sponsored research counts as deceitful in the sense of being a case of wishful speaking. What makes ghost managed research problematic, it seems, is not so much the papers which are published, but choices not to publish other papers. It is silence, rather than speech, which is deceitful. When, though, are silences deceitful? In this section, I first address this question, before developing an account of how co-ordinated patterns of silence and speech can be described using the terminology of wishful speaking. Finally, I show how using the terminology of wishful speaking helps us better understand concerns about publication bias, and different forms of objectivity.

The relationship between silence and deceit is complex, particularly in scientific contexts. It does not follow from the claim that we should assert *only* well-established claims that we should assert *all* well-established claims; there are ethical reasons not to publish findings about how to build a bomb even when—perhaps, particularly when—they are well-established (Kitcher 2003). Less dramatically, scientific research programmes may often throw up uninteresting or inconclusive results or go down blind alleys. Scientists may have slightly unsavoury reasons for not publicising these results—for example, to avoid appearing incompetent—but failing to communicate is not always deceitful. Therefore, the mere fact that ghost-management suppresses research does not show that they deceive. However, some scientific silences so seem deceitful. Consider, for example, a different version of the vaccination example. A senior epidemiologist is aware that the claim that a vaccine is safe has not yet been well-established. She studiously fails to intervene in on-going debate, knowing that her silence will be taken to imply that the vaccine is safe.

⁸ See Williams, (2002, pp. 141–142) for the importance of distinguishing whether scientists have non-epistemic motivations for making assertions and whether their non-epistemic motivations undermine the trustworthiness of the content of those assertions.

⁹ This is hinted at in the statistical analysis by Healy and Cattell (2003), and developed by Sismondo (2009).

How, then, can we distinguish, in general, between permissible and impermissible scientific silences? I suggest that scientific silences are problematic when audiences' expectations about scientists' communicative obligations are such that the audience will infer from silence that some claim is well-established when, in fact, it is not.¹⁰ Given the audience's expectations, the epidemiologist deceives, but an individual, unknown scientist who chooses not to publish a null result on a topic of little significance does not. In the previous section, I argued that the lying/misleading distinction is irrelevant to the ethics of science communication; so, too, there seems no significant moral difference between deceiving through scientific speech and through silence.

Expanding wishful speaking to include "wishful silences" does not automatically imply that ghost managed research is a problematic violation of communicative norms. No-one forms a belief that a particular drug is safe simply because any particular trial was not published. However, we can think of deceit not only in terms of one-off decisions, but in terms of communicative strategies. Consider an everyday case: an old man asks me how his son is doing; I know the son died this morning, but I reply, "I saw him yesterday and he seemed fine".¹¹ This true report might count as deliberately misleading, given that I know the old man will infer that his son is alive. Imagine, however, that the old man is naturally suspicious. So, I add a second (true) statement: "his wife called me earlier and she sounded happy". Call this "extended misleading"; neither statement on its own would lead the old man to a false belief, but jointly they suffice for doing so. Although extended misleading is more semantically complex than familiar cases of misleading it seems morally equivalent.

Patterns of communication may deceive in a similar, albeit even more complicated, way in scientific contexts. Imagine, for example, yet another epidemiologist. Rather than simply reporting that the vaccine has no side effects, instead on day 1, she says "study 1 showed no side effects", on day 2, she says "study 2 showed no side effects", and, on day 3, her last morning on the Today programme, "there is no plausible known mechanism by which this vaccine could lead to these side effects". Assume each is true. None of these claims independently implies that the vaccine is safe, but a competent hearer might infer from all three of them that the vaccine's safety has been established. However, the epidemiologist fails to report another true claim—"animal studies have shown the vaccine has side effects"—such that no respectable epidemiologist would judge the *overall* body of evidence suffices to accept that the vaccine is safe. (I simply assume that this is how biomedical researchers do—or should—combine evidence). In this case, the speaker has abused the hearer's assumption that the speaker's testimony is governed by a norm of "quantity", to assert as much relevant information as possible to implicate a false claim about the overall state of knowledge (Grice 1975).¹² As such,

¹⁰ Whether silences can be misleading has spawned a long debate in medical ethics; see Jackson (2001, esp. Chaps. 4 and 7). For epistemological issues, see Goldberg (2011).

¹¹ The example here is a variation on a standard case used in the literature; I do not know where it originates.

¹² Note an overlap here between the ethics of communication and the norms of belief. According to Carnap's "principle of total evidence", estimates of probability should be made relative to all of the available evidence (Carnap 1947). On my account, when a speaker fails to abide by ethical norms governing communication, she fails to make available to a hearer evidence which the latter ought (in an epistemic sense) use in her reasoning.

although each claim she makes is well-established, she engages in “extended wishful speaking”:

In extended wishful speaking a scientist makes a series of assertions, each of which is individually well-established relative to local epistemic standards, but where the pattern of communication is such as to lead a reasonable hearer to infer a claim which is not well-established, and this communicative strategy is adopted in order to secure the perlocutionary consequences of listeners believing the relevant implied claim or believing that the speaker believes that claim (regardless of its truth).¹³

Consider, now, a slightly more complex variant of the first case. The epidemiologist colludes with her research team, and, rather than appear on the radio herself each morning, instead sends colleague 1 to report on day 1, colleague 2 to report on day 2, and colleague 3 to report on day 3. By this point, our everyday vocabulary of “deceit” is surely stretched. However, such a case is structurally analogous to “extended wishful speaking” in ethically relevant ways: in both cases, some agent co-ordinates a pattern of communication with the end of leading audiences to believe ill-established claims, and the means to this end involves abusing audiences’ expectations. As such, we can describe it as “socially extended wishful speaking”:

In socially extended wishful speaking, one or more scientists makes a series of assertions, each of which is individually well-established relative to local epistemic standards, but where the pattern of communication is such as to lead a reasonable hearer to infer a claim which is not well-established, and this communicative strategy occurs as a result of some agent or agents’ intention to secure the perlocutionary consequences of listeners believing the relevant implied claim or believing that the speaker believes that claim (regardless of its truth).

Of course, when multiple individuals, each with multiple intentions, interact, it may be difficult to decide whether some pattern of communication was chosen because some agent (or agents) wished to secure certain ends. Furthermore, in such cases of collusion, questions of culpability and blame become extremely complex (Kutz 2007). However, worries about the very possibility of intentional action aside, there is no reason to doubt that some agents can intentionally engineer a pattern of communication (by themselves or other agents) with the aim of promulgating certain beliefs, and we can have evidence of such behaviour. Consider debates over the link between cigarette smoking in the late 1950s and early 1960s (Proctor 2012). We have excellent evidence that patterns of epidemiological publication were causally related to the intentions of some powerful agents to ensure that publics believed ill-established claims (or, perhaps more accurately, failed to believe well-established claims). In short, socially-extended wishful speaking is a particularly egregious form of agnotology.

In cases of ghost managed research is that various pieces of evidence—ranging from what CROs advertise at conferences to the asymmetrical publication pattern—imply that socially extended wishful speaking has occurred; patterns of communication have

¹³ The notion of “communicative strategy” borrows from Neil Manson’s analysis of political spin (Manson 2012).

been engineered such that, even if each paper is epistemically impeccable, a reader who bases her overall assessment of drug efficacy or safety on published evidence infers a conclusion which is not well-established relative to the actually-collected evidence. Although socially extended wishful speaking is removed from the paradigm case of wishful speaking, which is removed from everyday deceit, all three phenomena share ethically salient features. Hence, I suggest that ghost managed research breaches ethical norms governing communication.

One reason why it can be hard to prove that socially extended wishful speaking has occurred is that certain sorts of communicative patterns which might lead audiences to form ill-established beliefs might occur for reasons other than intentional manipulation of communicative norms; most obviously, when institutional arrangements fail to incentivise scientists to perform speech acts they ought to perform. Plausibly, this is true of biomedical research more generally; an unintended result of institutional factors is publication bias, non-publication of negative or null results, which makes it difficult to assess the likely efficacy and safety of various drugs (Stegenga 2011). Therefore, quite apart from concerns about ghost managed research, we have good epistemological reasons to change biomedical publishing more generally; for example, through requiring pre-registration of trials or publication of all trial data. It may seem, then, that concerns about ghost managed research are simply part-and-parcel of a larger set of concerns about the epistemic trustworthiness of contemporary biomedicine, and that responding to the latter concerns will automatically solve the former. However, I have tried to show that such cases are ethically problematic, regardless of their socio-epistemic consequences. Why?

First, there is a moral difference between responding to incentives created by the institutions of biomedical publishing and manipulating those institutional structures independent of the epistemic effects of those incentive structures more generally. Consider an analogy with conflicts of interest policies. There is some evidence that requiring scientists to state conflicts of interest may have the perverse consequence of making them less alert to possible sources of bias (Intemann and de Melo-Martin 2009). This is a good reason to rethink those policies. Still, if those policies are in place, there is a difference between deliberately flouting them and accidentally falling prey to bias, and these differences are relevant to how we think about moral blame and culpability. Similarly, “socially extended wishful speaking” and sets of interlocking institutional norms may have broadly similar socio-epistemic effects—publication bias—but the former is ethically problematic in a way in which the latter is not.

Second, there is no a priori reason why instituting systems and norms which cannot be “gamed” will necessarily have the best epistemic and practical consequences overall. Even if, in my case, we have both ethical and epistemic reasons to be concerned about current publication systems, proposals motivated by ethical concerns about unfair abuse of the system do not necessarily lead to the system which has the best social-epistemic consequences or vice versa. Drawing this distinction is particularly important in scientific contexts, such as biomedicine, which are closely related to regulatory contexts. Ted Porter has pointed out that in a “public measurement system”, such as those used by state bureaucracies, “standardisation” (that like cases be—and/or be seen to be treated alike) and “proper surveillance” are important values. These values give rise to a “strong incentive to prefer readily standardizable measures to highly

accurate ones” (Porter 1994, p. 391). Indeed, Porter even suggests that were one manufacturer to use “state-of-the-art” analysis to assess toxicity, this would be “viewed as a vexing source of interlaboratory bias, and very likely an effort to get more favorable measures by evading the usual protocol, not as a welcome improvement in accuracy” (Porter 1994, p. 391).

To generalise, there may be a systematic tension between having epistemic rules and institutions which allow for “procedural objectivity”—systems which cannot be gamed by various actors, through means such as extended wishful speaking—and systems which allow for greater objectivity in the sense of mirroring nature more accurately (see McGill 1994). We may have to choose between systems which minimise the chance of deceit and systems which promote the generation and spread of knowledge. In line with the exploratory, provisional aims of this paper, I do not have an account of how to make such choices; however, it is important to recognise that ethical and epistemic concerns might not always coincide.

4 Proleptic wishful speaking

In the previous section, I noted that it is a mistake to assume that just because published research is sponsored by the pharmaceutical industry, it must, therefore, be epistemologically suspect. Rather, much pharma-sponsored research is consistent with prevailing epistemic standards. Nonetheless, epistemic standards do not determine *every* methodological choice. As Sect. 1 noted, pharma-sponsored researchers often resolve these “micro-inductive risk” problems in ways which favour establishing certain hypotheses; for example, they might choose model organisms which are known to respond well to certain chemicals. These practices appear problematic. Furthermore, they are clearly guided by communicative interests; methodological choices are made based on an assessment of the value of being able to say, to regulators or publics, that certain results have been shown. Hence, it seems that a decent account of the ethics of communication should help us think through these cases. However, as it stands, the concept of wishful speaking is simply silent on the question of how research *should* be conducted, but, rather, seems to rely on an independent account of when claims are well-established. In this section, I address these challenges, developing a concept of “proleptic wishful speaking” which helps us think through the proper role of non-epistemic interests in research contexts.

Consider another simplified case: an epidemiologist values the consequences of parents believing a vaccine is safe (regardless of whether or not it is safe); however, she also does not want to make any ill-established claims (after all, as we noted above, straightforward wishful speaking is rightly vilified.) She starts a new research project, during which she makes a series of methodological choices which increase the chance that she will be able to conclude that, relative to her evidence, the vaccine is safe (she is more likely to disregard evidence that the vaccine is unsafe than evidence that the vaccine is safe, uses model organisms which tend not to show side-effects, and so on). Imagine, further, that each of these individual decisions is epistemically permissible, relative to the standard norms of biomedical testing. At the end of this project, then, the claim “the vaccine is safe” is well-established, in the sense outlined

in Sect. 2. The researcher then makes the claim to the public. This researcher has not engaged in wishful speaking in any of the senses outlined so far. However, there is a clear resemblance between her research strategies and standard wishful speaking: in standard wishful speaking a speaker makes an ill-established claim because she values the consequences of others believing that claim, regardless of its truth; in our case, a researcher decides on research strategies because she values the consequences of others believing some claim, regardless of its truth. Call this phenomenon “proleptic wishful speaking”:

A scientist engages in proleptic wishful speaking when she makes methodological choices in the course of research where her motivation for making those choices is a desire to show a specific claim is well-established, and where she wants to show this claim is well-established because of the predicted non-epistemic benefits that follow from others believing it, regardless of its truth.

The concept of proleptic wishful speaking describes one sort of strategy for making methodological choices. Other strategies are possible: for example, the scientist could just choose the cheapest option, or do whatever her colleagues do, or even just flip a coin! Is there anything wrongful with proleptic wishful speaking? Above, I suggested that, even if scientific deceit is particularly wrongful because it violates audiences’ expectations, all deceit is *prima facie* wrongful, and gave a (broadly) Kantian account of why: deceitful speech disrespects hearers by treating their beliefs as mere means to our ends. When we deceive others, we seek to influence them into behaving certain ways. By contrast, when we speak truly, we improve others’ ability to choose in ways which reflect their own values. Deceitful speech fails to respect others in their capacity as autonomous agents. One can contest this account of the *prima facie* wrongness of deceit. However, if plausible, it suggests an ethical objection to proleptic wishful speaking; in adopting such research strategies, a researcher treats the communication of results as a tool for changing others’ behaviour, rather than as a tool for promoting their ability to make informed choices. This attitude towards communication is the same attitude that renders straightforward deceit problematic.¹⁴ Hence, even if proleptic wishful speaking differs from standard wishful speaking, they are ethically problematic for similar reasons. And, insofar as we think that everyday deceit is particularly problematic in scientific cases, then it seems particularly problematic for scientists to display the attitudes involved in proleptic wishful speaking.

Of course, real-life cases, such as those encountered in pharmaceutical research, are far more complex than my hypothetical epidemiologist. In these cases, research choices may typically be affected through the kinds of complex co-ordination discussed in the previous section (Sismondo 2009). As such, strictly, we are dealing with forms of

¹⁴ Better to understand the concept of an attitude, consider Scanlon’s distinction between the “permissibility” and “meaning” of an action (Scanlon 2008, Chap. 1). In Scanlon’s example, the right thing to do when you encounter a drowning child is to jump in the river to save him. That you do so out of a desire to appear in the newspaper may reflect badly on you, affecting the meaning of your action, but is irrelevant to its permissibility. My suggestion here is that there is a similar structure at play in cases of proleptic wishful speaking: as long as each choice was genuinely epistemically underdetermined, then, even if we engage in proleptic wishful speaking, the end claim is well-established, i.e. epistemically permissible. However, in getting to that claim via proleptic wishful speaking, we act in an ethically problematic manner because we express problematic attitudes; the meaning is problematic.

socially extended proleptic wishful speaking. Nonetheless, these suggestions imply an important *ethical* objection to how pharma-sponsored researchers solve micro-inductive risk problems: that their strategies imply a problematic attitude towards hearers' autonomy, in a way in which, for example, choosing the cheapest option, or even flipping a coin does not.

To clarify these claims, consider an alternative approach; writing about Wilholt's example of lab rat choice (and similar cases in pharmaceutical research) Winsberg et al. (2014) object that such strategies reduce "producer risk" at the expense of "consumer risk"; i.e. the researchers decrease the chance that the pharmaceutical company will obtain a "bad" result from the perspective of maximising their profits, but only by increasing the chance that some dangerous drug will be permitted on the market, a "bad" result from the public's perspective (Winsberg et al. 2014). I agree that the wrong involved in these cases is not straightforwardly epistemic, but also has an important ethical dimension. However, I disagree on how to characterise the ethical wrong. On their approach, it seems that the problem with industry-sponsored research is not that researchers engage in proleptic wishful speaking, but proleptic wishful speaking based on an incorrect assessment of the value of audiences believing certain claims. On the approach I have sketched, by contrast, proleptic wishful speaking is *prima facie* wrongful, regardless of whether it is aimed at reducing consumer or producer risk.

Do I really think it is preferable to make methodological choices by flipping a coin, rather than by the aim of reducing consumer risk? All-things-considered, probably not; rather, just as deceit may be all-things-considered permissible, the *prima facie* wrongness of proleptic wishful speaking may be outweighed by other concerns. I do not pretend to have a full account of how to balance all of the relevant ethical considerations. Furthermore, it may be that some forms of proleptic wishful speaking can be ethically better than others; for example, if researchers draw on non-epistemic values which they know their audience shares.

Still, the concept of wishful speaking has a role to play in helping us understand one, distinctively ethical, way in which justification can be distorted by non-epistemic values: not only by paying too little account of consequences or too much account of the wrong consequences, but by paying too great an attention to consequences at all.

Even if the practical implications of this view are unclear, it may have interesting theoretical implications. It has become common to say that non-epistemic values play some legitimate role in justification (Hicks 2014). However, it is far less clear that anyone agrees on what that role is. The case of industry-sponsored research neatly illustrates some problems for thinking about these issues: the ways in which the non-epistemic concern for profit structures pharma-sponsored research seem problematic. However, even defenders of the "epistemic priority thesis" agree that non-epistemic values can play a proper justificatory role when epistemic considerations underdetermine research strategy (Steel 2013) [and, of course, opponents of that thesis hold that they can play an even more expansive role (Brown 2017)]. Hence, it seems all sides of debate must agree that if there is something wrong in how industry-sponsored researchers resolve methodological underdetermination, then the problem is broadly ethical. One way of filling out these concerns is to say that the profit motive is necessarily inimical to the goals of biomedical research (Hicks 2014). However, such value judgments are by no means uncontroversial (Reiss 2017).

One way in which we might resolve these disputes would be to establish whether using the profit motive to resolve methodological choices tends to lead to better or worse outcomes in terms of population health and well-being. The arguments above suggest, however, that such information does not provide a full account of the ethical terrain. Rather, on my proposed account, it is problematic for researchers to appeal to non-epistemic values in ways which constitute wishful speaking, regardless of the further consequences (positive or negative) of so doing. Therefore, recognising that concerns about deceit might affect the propriety of research strategy captures the intuitive concern that there is something problematic about how industry-sponsored researchers resolve micro-inductive risk problems, without requiring us to establish claims about the likely effects of these choices. Given the sheer difficulty of establishing such causal claims, this seems to me an advantage of my approach. However, this approach also has a cost: it implies that there may be ethical reasons to structure research choices in ways which we know are sub-optimal in terms of their ethical consequences; for example, by seeking to reduce cost, rather than reduce consumer risk. This may seem puzzling. However, it is, I suggest, simply a variant on the familiar tension between deontic and consequentialist ethical considerations. Recognising the possibility of such tensions is important for thinking about the broader debate over the proper role of non-epistemic values in scientific research. Often, we place ethical limits on permissible research, even when doing so hinders the pursuit of knowledge (Brown 2017). My claim here, then, is that, just as ethical considerations—say, concerning individuals' rights to bodily integrity—can sometimes over-rule epistemic considerations, so, too, they can sometimes over-rule other ethical considerations. Recognising this complexity can help us see how a concern to insulate methodological choices from certain sorts of non-epistemic value judgments need not involve denying the relevance of non-epistemic values to proper justification; rather, they can be an expression of that relevance. The previous section warned against thinking that our epistemic and ethical concerns must coincide when thinking about scientific communication; the arguments above warn against thinking that disparate ethical concerns must always coincide when thinking about the complex interplay between justification and communication.

5 Many speakers

A common concern about ghost-managed papers is that they misascribe authorship, in the sense that they do not follow the rules set out by such bodies as the International Committee of Medical Journal Editors for listing “authors”. This is a form of deceit, but one which I will not dwell on, because, more interestingly, these institutional norms are themselves contested (Bhopal et al. 1997; Horton 1998; Biagioli 1998). In turn, these debates point to a problem for thinking about the concept of deceit specifically, and for an ethics of science communication more generally: given the complexities of modern scientific research, it is often difficult to identify who (or what) counts as “the” speaker of a paper. This problem poses a fundamental challenge to the arguments of this paper. I have tried to show how we can develop an account of scientific speech which is both sensitive to the institutional realities of modern science and which captures distinctively ethical concerns. However, even a complex ethics of communication—which

focuses on patterns of speech and how speech can be manipulated by second parties, and on how values influence justificatory practices—still requires that we can identify speakers. There must be someone speaking for there to be socially extended wishful speaking or proleptic wishful speaking to occur. The production and promulgation of sentences may cause harms, but unless we can identify someone who is making those claims, it is a mistake to think about those harms in terms of communicative obligations. This section first sketches these problems in more detail, then responds to them.

Imagine that, sitting in splendid isolation, I write a paper and send it to a journal which publishes it. To make matters simple, imagine that the editors do not even require a single correction (this is a distant possible world!) Unfortunately, due to a printing error, my name is replaced with yours. In this case, the printer's error may lead others to the erroneous conclusion that you are the speaker of the claims in my paper. Consider, now, a potted summary of the modern bio-medical research process [as expertly described in Kukla (2012)]. Trials may be run by many different teams of researchers in different countries. Results may then be pooled and shared, and analysed and re-analysed by people who did none of the original research, who possess different kinds of expertise and knowledge. Preliminary results and early drafts may be shared with some of the original researchers, before the whole is tidied-up by a professional “medical writer” for submission. Unlike in the case of my perfect paper, it seems that in this case we face a question of who or what counts as the “author” in the first place?¹⁵

Broadly, we face three options: a single human individual; a set of human individuals; a corporate group (“the research team”). Given how research papers are assembled, it seems odd to pick out a single individual as “the” speaker of all of the paper's claims, ruling out the first option. The second and third options also both seem problematic. It is entirely plausible that none of the individual contributors believe—still less know—each and every claim in a collaboratively-created paper [see Horton (1998) for some examples, initially suggested as philosophical possibilities by Hardwig (1991)]. Such concerns may push us towards thinking of the paper as the articulation of the beliefs or knowledge of the group as a corporate whole. Unfortunately, even if groups can have minds of their own—which is contestable—the group of “authors” behind papers are rarely sufficiently well-integrated to count as corporate agents (Kukla 2012).

Faced with such puzzles, it may seem tempting to think that treating debates over authorship as resting on a metaphysical question over who counts as “the” speaker of the claims in a paper is a philosophical error. Really, one might say, authorship debates are *solely* concerned with social questions rather than some fact of the matter.¹⁶ For example, if authorship norms do (or should) signal “credit”, we might favour including all contributors as “authors”; if they assign “responsibility”, then, given legal considerations, more restrictive norms may seem appropriate. Disputes over authorship norms often concern conflicts might, then, simply be recast as debates over

¹⁵ As Huebner et al. (2017) stress, these problems do not necessarily arise in all cases of co-authorship, but are a distinctive feature of research processes where there is neither an epistemic to-and-fro between authors nor a single authorial “controller”.

¹⁶ This line seems implicitly assumed in Kukla et al. (2014).

which functions matter most.¹⁷ However, this approach has costs. Scientific papers certainly look like they contain assertions. A sentence such as “this drug is safe”, read in a journal, seems like a claim about the drug’s safety, rather than a complicated kind of evidence for its safety. If so, either scientific papers must have speakers or they are speech acts without real speakers. The latter option has odd consequences, both theoretical and practical: without speakers we lose our grip on the locus of communicative responsibility.

We must either rethink conditions for being a speaker, or deny that scientific papers are assertions, or hold that there are speech acts without speakers. The cleanest and simplest response to this trilemma is to adopt a minimalist understanding of who counts as a speaker (of scientific papers, at least): the speaker of a scientific claim is whichever individual or group of individuals fulfils the functional role of producing that claim. What renders a claim “scientific” is a set of markers—both internal to the paper itself (for example, that certain conventions are obeyed, certain references are cited, and so on) and institutional (for example, that the claim appears in a certain sort of journal which practises certain forms of peer review). It may be hard to identify necessary and sufficient conditions for a claim to count as scientific, but, we may gain hints by looking at the social strategies used by those who wish to *appear* scientific; for example, homeopaths or climate change deniers (Oreskes 2017).

How, though, can a group of individuals be said to be the speaker of a claim if that group lacks the coherence to count as a corporate agent? This apparent puzzle arises, I suggest, only because we think of assertion as an articulation of a belief, and, therefore, proper assertion as requiring some agent to whom we may ascribe beliefs. However, as Sect. 2 argued, norms of sincerity are irrelevant to the propriety of scientific speech. From my proposed perspective, speakers simply are those things which play the role of making claims about the world. In turn, authorship ascriptions can be understood as doing two things. First, they identify the individuals who constitute the body that play this role. Second, by publicly identifying these individuals, they serve a social purpose, of providing these individuals with incentives to ensure that this role is played well.¹⁸ Strictly, scientific speakers may be mindless; as such, it makes no sense to hold them accountable. However, authorship norms may still ensure that human individuals—who can be held accountable—have good reasons to ensure that these speakers’ words are proper.¹⁹

One aim of this paper has been to examine a particular topic in the ethics of scientific communication: the nature of deceit. A more general aim is to explore questions for a deep ethics of science communication; i.e. an account of the general ethical norms incumbent upon scientific speakers, which we might use to criticise and assess institutional norms. I have argued that, even in radically collaborative research contexts, we can identify speakers, thus saving that project. Note, of course, that these

¹⁷ Of course, these functions vary over time: see Gallison and Biagioli (2003).

¹⁸ I suggest that in his influential 2001 paper, Hardwig is led to questions about the possibility of collective scientific knowledge via the tacit assumption that we learn from scientific papers much as we learn from normal testimony and the assumption that testimony requires sincerity. Denying the latter claim allows us to avoid thinking that a long list of authors raises a puzzle about collective belief. It also has some practical benefits—see John (2018).

¹⁹ Therefore, I am more optimistic in principle—if not in practice—than Huebner et al. (2017).

speakers may not be psychologically integrated in such a way that we can describe them as having intentions. Often, then, it will be misguided to talk of such speakers as engaged in standard wishful speaking or straightforward proleptic wishful speaking. Nonetheless, the claims these speakers make may still be deliberately influenced by the intentional choices of other agents; for example, publications “spoken” by a group may be determined by the choices of their funders. Hence, even if these speakers are not themselves ethically culpable, these cases exemplify the ethical wrong of socially co-ordinated wishful speaking. We can use a distinctively ethical vocabulary—drawn from our understanding of simpler, everyday cases of deceit—to think about complex scientific communication.

6 Conclusion

In this paper, I have developed the concept of wishful speaking through a case-study of the ethical wrongs which occur when the non-epistemic interests of pharmaceutical companies structure both the pursuit and communication of results. Such influences can lead to socially extended wishful speaking and socially extended proleptic wishful speaking, phenomena which, I have argued, count as complex forms of deceit. I have also shown how this general account of scientific speech relates to the complex problem of authorship in an age of radically collaborative research, suggesting that we need to move away from the thought that a scientific speaker requires a single scientific mind. Although, as I have stressed, my first-order claims are provisional, this task is important for three reasons. First, as I noted at the start, there are pressing practical reasons to understand the nature of scientific deceit; although we routinely demand of scientists that they avoid deceit, our claims are empty if we lack a clear sense of what deceit means in the scientific context. Second, there are serious difficulties in the blithe assumption that we can simply apply concepts from everyday ethical contexts to the communicative structures of modern science. Showing that this task is possible, and clarifying how it relates to other ways in which we might think about those structures (for example, in terms of their epistemic and non-epistemic consequences) is important for shoring up the possibility of an ethics of modern science communication at all.

Third, and most generally, as Sect. 4 discussed, there is a long-running dispute over the proper role or roles of various sorts of non-epistemic values in different scientific contexts. Clearly, the concept of deceit, even tidied up into a concept of wishful speaking, cannot resolve all of these disputes. However, it does provide some ways into thinking about them. The vast majority of scientific research is undertaken to be communicated. Hence, it seems that there must be some relationship between these debates over scientific justification and the ethics of communication. One strand of research focuses on the implications of recognising the role of non-epistemic values in research for how we communicate scientific results; for example, that research should be presented so as to allow for “back-tracking” (McKaughan and Elliott 2013). A second strand works the opposite direction: by trying to unpick issues around justification by thinking about communication. For example, in different ways, Wilholt (2013) and John (2015) have argued that the needs of communication rule out some appeals to non-epistemic values on grounds of efficiency. Franco (2017) has developed a line-of-

thought implicit in Douglas (2009), that, because, as a speech act, scientific assertion does more than merely express beliefs, scientists have a responsibility to consider the non-epistemic consequences of different kinds of error in resolving inductive risk problems.

In this paper, I have outlined one way in which we might develop the second strand of research, by showing how the concept of deceit may help us understand a range of issues at the interface between science and industry. However, whereas much research on the relevance of norms of communication for debates over non-epistemic values in science focuses on the relevance of the *consequences* of communicating claims for how we should think about *justification*, I have sought to explore a different concern: the *deontological* constraints on the proper role of perlocutionary knowledge in communication. Given that the border between communication and justification is porous, this approach has implications for how we think about justification. Developing the broader implications of this approach beyond the case of pharma-sponsored research is not the task of this paper. Still, I do hope to have shown that the everyday concept of deceit might have important applications not only for what scientists should say, but for how they should reason.²⁰

Acknowledgements This paper has benefited greatly from feedback from readers including Neil Manson, Gabriele Badano, Charlotte Goodburn and Jacob Stegenga. It has also benefited from feedback from audiences in Cambridge (CRASSH work-in-progress seminar), London (Institute of Cancer Research), and Beijing (Centre for Science Communication, Peking University). I am also grateful to the Independent Social Research Foundation for supporting this research via their funding of the “Limits of the Numerical” project. However, I am most grateful to students who attended my Part II “Primary Source” seminar, “Ghost-authorship in medical research”, in the Department of History and Philosophy of Science, University of Cambridge, 2015–2016. Their questions and answers were far more interesting than anything that would occur to me alone.

Open Access This article is distributed under the terms of the Creative Commons Attribution 4.0 International License (<http://creativecommons.org/licenses/by/4.0/>), which permits unrestricted use, distribution, and reproduction in any medium, provided you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons license, and indicate if changes were made.

References

- Alamassi, B. (2014). Medical ghostwriting and informed consent. *Bioethics*, 28(9), 491–499.
- Austin, J. L. (1975). *How to do things with words*. Oxford: Oxford University Press.
- Bhopal, R., et al. (1997). The vexed question of authorship. *British Medical Journal*, 314, 1009.
- Biagioli, M. (1998). The instability of authorship: Credit and responsibility in contemporary biomedicine. *The FASEB Journal*, 12(1), 3–16.
- Biddle, J. (2007). Lessons from the Vioxx debacle: What the privatization of science can teach us about social epistemology. *Social Epistemology*, 21(1), 21–39.
- Brown, M. (2017). Values in science: Against epistemic priority. In D. Steel & K. C. Elliott (Eds.), *Current controversies in values and science*. London: Routledge.
- Carnap, R. (1947). On the application of inductive logic. *Philosophy and Phenomenological Research*, 8, 133–148.
- Chiu, K., Grundy, Q., & Bero, L. (2017). ‘Spin’ in published biomedical literature: A methodological systematic review. *PLoS Biology*, 15(9), e2002173.

²⁰ For some sense of how this project might go, see John (forthcoming).

- Department for Innovation, Universities and Skills. (2007). *Rigour, respect, responsibility: A universal ethical code for scientists*. London: Her Majesty's Stationery Office.
- Douglas, H. (2009). *Science, policy and the value free ideal*. Pittsburgh: University of Pittsburgh Press.
- Eggers, P. (2009). *Deceit: The lie of the law*. London: Informa Law Publishing.
- Flanagin, A., et al. (1998). Prevalence of articles with honorary authors and ghost authors in peer-reviewed medical journals. *Journal of the American Medical Association*, 280(3), 222–224.
- Franco, P. (2017). Assertion, non-epistemic values, and scientific practice. *Philosophy of Science*, 84(1), 160–180.
- Fricker, E. (2002). Trusting others in the sciences: A priori or empirical warrant? *Studies in History and Philosophy of Science Part A*, 33(2), 373–383.
- Gallison, P., & Biagioli, M. (Eds.). (2003). *Scientific authorship*. Cambridge: Harvard University Press.
- Goldberg, S. (2011). The epistemic division of labour. *Episteme*, 8, 112–125.
- Gøtzsche, P., et al. (2007). Ghost authorship in industry-initiated randomised trials. *PLoS Medicine*, 4(1), e19.
- Grice, P. (1975). Logic and conversation. In P. Cole & J. Morgan (Eds.), *Syntax and semantics 3: Speech acts* (pp. 41–58). New York: Academic.
- Hardwig, J. (1991). The role of trust in knowledge. *Journal of Philosophy*, 88(12), 693–708.
- Healy, D., & Cattell, D. (2003). Interface between industry, authorship and science in the domain of therapeutics. *British Journal of Psychiatry*, 183, 22–27.
- Hicks, D. (2014). A new direction for science and values. *Synthese*, 191(14), 3271–3295.
- Horton, R. (1998). The unmasked carnival of science. *The Lancet*, 351(9104), 688–689.
- Horton, R. (2002). The hidden research paper. *Journal of the American Medical Association*, 287(21), 2775–2778.
- Huebner, B., Kukla, R., & Winsberg, E. (2017). Making an author in radically collaborative research. In T. Boyer-Kassem & C. Mayo-Wilson (Eds.), *Scientific collaboration and collective knowledge*. Oxford: Oxford University Press.
- Intemann, K., & de Melo-Martin, I. (2009). How do disclosure policies fail? *The FASEB Journal*, 23, 1638–1642.
- Jackson, J. (2001). *Truth, trust and medicine*. London: Routledge.
- John, S. (forthcoming). Science, truth and dictatorship: Wishful thinking or wishful speaking? *Studies in History and Philosophy of Science*.
- John, S. (2015). Inductive risk and the contexts of communication. *Synthese*, 192(1), 79–96.
- John, S. (2018). Epistemic trust and the ethics of science communication: Against transparency, openness, sincerity and honesty. *Social Epistemology*, 32(2), 75–87.
- Kitcher, P. (1993). *The advancement of science*. Oxford: Oxford University Press.
- Kitcher, P. (2003). *Science, truth and democracy*. Oxford: Oxford University Press.
- Kjaergard, L., & Als-Nielsen, B. (2002). Association between competing interests and authors' conclusions: Epidemiological study of randomised clinical trials published in the BMJ. *British Medical Journal*, 325, 7358.
- Kukla, R. (2012). Author TBD. *Philosophy of Science*, 79(5), 845–858.
- Kutz, C. (2007). *Complicity: Ethics and law for a collective age*. Cambridge: Cambridge University Press.
- Lackey, J. (2007). Norms of assertion. *Noûs*, 41(4), 594–626.
- Lane, M. (2014). When the experts are uncertain: Scientific knowledge and the ethics of democratic judgment. *Episteme*, 11(01), 97–118.
- Manson, N. C. (2012). Making sense of spin. *Journal of Applied Philosophy*, 29(3), 200–213.
- McKaughan, D. J., & Elliott, K. C. (2013). Backtracking and the ethics of framing: Lessons from voles and vasopressin. *Accountability in research*, 20(3), 206–226.
- Megill, A. (1994). Introduction: Four senses of objectivity. In A. Megill (Ed.), *Rethinking objectivity* (pp. 1–20). Durham: Duke University Press.
- Newton, S. (2011). Creationism creeps into mainstream geology. *Earth*. <https://www.earthmagazine.org/article/creationism-creeps-mainstream-geology>.
- Nickel, P. (2013). Norms of assertion, testimony and privacy. *Episteme*, 10(02), 207–217.
- O'Reilly, J., Oreskes, N., & Oppenheimer, M. (2012). The rapid disintegration of consensus: The west antarctic ice sheets and the intergovernmental panel on climate change. *Social Studies of Science*, 2012(42), 709–731.
- Oreskes, N. (2017). Systematicity is necessary not sufficient: On the problem of facsimile science. *Synthese*. <https://doi.org/10.1007/s11229-017-1481-1>.

- Porter, T. (1994). Making things quantitative. *Science in Context*, 7, 389–407.
- Proctor, R. N. (2012). The history of the discovery of the cigarette–lung cancer link: Evidentiary traditions, corporate denial, global toll. *Tobacco Control*, 21(2), 87–91.
- Rees, C. F. (2014). Better lie! *Analysis*, 74(1), 59–64.
- Reiss, J. (2017). Meanwhile, why not biomedical capitalism? In D. Steel & K. C. Elliott (Eds.), *Current controversies in values and science*. London: Routledge.
- Resnik, D., & Elliott, K. (2013). Taking financial relationships into account when assessing research. *Accountability in Research: Policies and Quality Assurance*, 20(3), 184–205.
- Saul, J. M. (2012). *Lying, misleading, and what is said*. Oxford: Oxford University Press.
- Scanlon, T. (2008). *Moral dimensions*. Cambridge: Harvard University Press.
- Sismondo, S. (2009). Ghosts in the machine. *Social Studies of Science*, 39(2), 171–198.
- Sismondo, S. (2011). Corporate disguises in medical science. *Bulletin of Science, Technology and Society*, 31(6), 482–492.
- Sorensen, R. (2012). Lying with conditionals. *The Philosophical Quarterly*, 62(249), 820–832.
- Steel, D. (2013). Acceptance, values, and inductive risk. *Philosophy of Science*, 80(5), 818–828.
- Stegenga, J. (2011). Is meta-analysis the platinum standard of evidence? *Studies in History and Philosophy of Science Part C: Studies in History and Philosophy of Biological and Biomedical Sciences*, 42(4), 497–507.
- Vinkers, C. H., Tjldink, J. K., & Otte, W. M. (2015). Use of positive and negative words in scientific PubMed abstracts between 1974 and 2014: Retrospective analysis. *BMJ*, 351(December), h6467.
- Webber, J. (2013). Liar! *Analysis*, 73(4), 651–659.
- Wilholt, T. (2009). Bias and values in scientific research. *Studies in History and Philosophy of Science Part A*, 40(1), 92–101.
- Wilholt, T. (2013). Epistemic trust in science. *The British Journal for the Philosophy of Science*, 64(2), 233–253.
- Williams, B. (2002). *Truth and truthfulness*. Princeton, NJ: Princeton University Press.
- Winsberg, J., Huebner, B., & Kukla, R. (2014). Accountability and values in radically collaborative research. *Studies in History and Philosophy of Science Part A*, 46, 16–23.
- Woodman, A. (submitted). *Facilitating informed consent*. M.Sc. Degree, Department of History and Philosophy of Science, University of Cambridge.
- Zollman, K. J. (2007). The communication structure of epistemic communities. *Philosophy of Science*, 74(5), 574–587.