

1 Objectivity: Processes and Products

Consider some quotations from recent news articles:

[The court's judgment] gives judicial licence for women and men who speak up for objective truth and clear debate to be subject to aggression, bullying, no-platforming and economic punishment (Bowcott, 2019).

When presented with lies versus truth, every journalist is facing a test of conscience. Are you being truly objective and just or are you being selective and partial? What does your decision reveal? TRUTH AIN'T LIE! (Kuo, 2019).

Teachers, particularly of the sciences, have smuggled political positions into what should be factual and objective subjects (Seaman, 2019).

All three of these quotes vividly convey a sense that it is important to be 'objective' – to speak up for objective truth, to engage in objective journalism or to maintain the objectivity of teaching – but, also, that objectivity is fragile, and under attack. In turn, they relate to broader social worries that we are moving into a post-truth world, dominated by 'fake news' and a disregard for science and rationality. But they are puzzling: they suggest that certain things – selectivity, partiality, political commitments – *threaten* objectivity, but is there anything *more* to objectivity than the absence of these factors? To complicate matters further, all three quotes arise in the context of heated ongoing debates. The first quote is a response to a ruling over transgender rights; the second is from China's Ministry of Foreign Affairs, berating coverage of China's domestic policies; and the third is from a climate sceptic, decrying the teaching of climate science in schools. Can talk of objectivity really resolve politically charged debates, or is it just a mask for social and political agendas?

There is an obvious response to these concerns: to turn to the philosophy of science. After all, many philosophers of science have held that a distinctive feature of scientific inquiry is that it aims at objectivity, and that we should assess and criticise putatively 'scientific' practices in terms of whether or not they are objective. Unfortunately, the concept of objectivity is not straightforward. For example, in an influential article questioning the objectivity of the medical sciences, Stegenga defines an objective process as being 'not sensitive to researchers' personal idiosyncrasies and biases' (Stegenga, 2011, 499). This definition sounds plausible; it seems fairly clear that processes which are affected by biases are not objective. But, on reflection, it raises a series of questions: couldn't biases sometimes lead us to true claims? Don't we value the insights of idiosyncratic geniuses? Would our processes be any better at knowledge production if they were affected by shared, rather than idiosyncratic, biases? Can any process be entirely insensitive to researchers' idiosyncrasies?

In turn, these epistemological problems intersect with broader debates. On the one hand, they relate to metaphysical debates over realism and representation: in what sense is a quality like the efficacy of drugs ‘out there’, waiting to be discovered, and to what extent does it depend on us? On the other hand, they relate to debates over the proper relationship between science, society, politics and economics: can medical science ever be objective when so much is funded by the pharmaceutical industry? Would we be better-off were medical science nationalised?

We have no shortage of reasons, both practical and theoretical, to understand debates over the nature of scientific objectivity. The aim of this Element is to provide an overview of some of them. Of course, objectivity is a vast topic for a short study; as such, my coverage is partial. Specifically, I focus on four topics: the turn to trust in recent work on objectivity (Section 1); the relationship between biases, values and objectivity (Section 2); the relationship between objectivity and social structures (Section 3); and the relationship between objectivity and the notion of epistemic perspectives or situated knowledges (Section 4). Inevitably, I do not discuss these topics in as much depth as I would like, and there are many topics I would like to discuss but which I do not touch on at all; I warn the reader of any major gaps as I go along. The topic of objectivity is not only vast, but also controversial. I have tried to provide a balanced account of various debates, but I see no point in hiding my own opinion on some key debates. I do, however, try to warn the reader where I insert myself into the text. As we will see in Section 2, though, it is an interesting question whether being open about my own biases is enough for my claims to be objective!

1.1 Proliferating Senses of Objectivity

We can identify three constraints on any plausible account of scientific objectivity; it should be concerned with *representation*, pick out something *valuable* and it should be *viable*. Before going on, I shall explain these desiderata.

First, I take it that any account of objectivity should have some *fidelity* with our everyday uses of the term. In turn, I suggest that this implies that objectivity talk is primarily concerned with *practices of representation*. Although we use the term objective in a variety of ways – to refer to certain sorts of factual claims, or to certain processes for generating factual claims or to virtues of individual scientists – these uses are united by a sense that objectivity is concerned with representing the world; we care that evidence amalgamation methods in medicine are objective because, if they are not, we doubt that they will be accurate guides to figuring out whether drugs will work or not; we care

that journalists are objective because they are supposed to tell us what is actually happening. Of course, in some cases, objectivity may be tied-up with virtues other than representational accuracy; for example, you might think it is important that a judge is ‘objective’ in the sense that she is impartial; or we might think it is important that scientists building a climate model are ‘objective’, even if the sense in which a model ‘represents’ the world is complex. Still, I suggest that, even in these trickier cases, talk of objectivity is often bound up with concerns about representation: the judge should be impartial because partiality gets in the way of reaching what, according to the law, would be the ‘right’ answer; a good climate model need not be accurate in every regard, but its goodness is, at least in part, a function of the accuracy of its predictions. By contrast, talk of objectivity seems a little strange in cases where there is no link to practices of representation. For example, it would seem odd to say that an engineer redesigning a toaster should be ‘objective’ in her work (although, tellingly, it does make sense to say that she should be objective in *reporting* the results of her work – say, whether the toaster really uses less electricity). Below, I will return to the obvious worry that notions of accurate representation are metaphysically tricky; still, that worry is separate from whether we think of scientific objectivity in terms of accurate representation.

In this sense, *scientific* objectivity needs to be distinguished from a different ‘purely procedural’ sense of objectivity, as following any old rule blindly. For example, imagine an examiner who blindly follows the exam marking guidelines, placing her own sense of candidates’ ability to one side. We might say that she acts objectively, even if we share her suspicion that those guidelines are fatally flawed as a guide to candidates’ ability. When we talk about scientific processes as objective, though, we typically imply something more: that these processes help us represent the world well, rather than that they can be followed in a robotic manner. (Much more on this distinction to follow.)

Second, objectivity is epistemically *valuable*; we have good reason to ensure epistemic practices *are* objective and to trust objective practices. Of course, that is not to say that objectivity is always valuable; perhaps there are cases where we should bend the truth for the sake of some greater good. Nor is it to deny that there are important criticisms of objectivity, for example from feminist critics (see Section 4). Still, in general, it seems that we value objectivity, and any decent account of the concept should capture this thought.

Third, the concept of objectivity is supposed to be action-guiding, in the sense that it provides a yardstick for assessing and changing epistemic practices. As such, objectivity must be *viable*; even if it is impossible for a person, process or claim to be fully objective, it must be possible to be more or less objective. A nice question, to which we will return shortly, is how to relate this third

desideratum to the first: the stronger our account of what is required for an ‘objective’ representation of the world, the harder it is to see how the notion of objectivity could guide action.

Beyond these three desiderata, matters get more complex. Megill (1994) thinks that there are four key senses of objectivity; Gaukroger lists five senses (Gaukroger, 2012, chapter 1); Douglas (2004) may seem more modest, suggesting three ‘modes’ of objectivity (in terms of processes of interaction with the world, features of individuals’ thought processes and social procedures), but then suggests that each mode is further sub-divided into different forms. Most terrifyingly of all, in an influential article, Marianne Janack lists 13 senses of objectivity she has found in the literature:

- (1) objectivity as value neutrality;
- (2) objectivity as lack of bias, with bias understood as including: (a) personal attachment; (b) political aims; (c) ideological commitments; (d) preferences; (e) desires; (f) interests; (g) emotion;
- (3) objectivity as scientific method;
- (4) objectivity as rationality;
- (5) objectivity as an attitude of ‘psychological distance’;
- (6) objectivity as ‘world-directedness’;
- (7) objectivity as impersonality;
- (8) objectivity as impartiality;
- (9) objectivity as having to do with facts;
- (10) objectivity as having to do with things as they are in themselves; objectivity as universality;
- (11) objectivity as disinterestedness;
- (12) objectivity as commensurability;
- (13) objectivity as intersubjective agreement (Janack, 2002, 275).

In short, everyone agrees that the term objectivity is used in a lot of ways (even if they disagree over how many).

Of course, it is possible that a term is used in various ways, but with a shared core meaning. For example, looking at Janack’s list, we might think that the concept of objectivity as lack of bias (sense 2) stems from a concern that bias gets in the way of creating ‘intersubjective agreement’ (sense 13); that ‘impersonality’ (sense 7) is important for much the same reason; and so on. However, there are reasons to think that objectivity talk might not just be complex, but, in Douglas’ phrase, ‘irreducibly complex’: that there may be no guarantee that all of the different senses of objectivity will coincide (Douglas, 2004).

Broadly, these reasons stem from the fact that contemporary notions of objectivity have a complex history. As Loraine Daston and Peter Galison

(2007) suggested in their monumental historical study, our understanding of the epistemic virtues that should guide practices of representation has shifted over time. Botanical illustrators or map-makers of the eighteenth century were guided by an ideal of ‘truth to nature’, according to which representations should capture the ideal type of a specimen, through removing blemishes and quirks. In the nineteenth century, however, the advent of technologies such as photography motivated a shift towards ‘mechanical objectivity’, in which scientists were expected to show restraint and let nature ‘speak for itself’. In turn, in the twentieth century, this ideal gave way to a notion of the expert as exercising a kind of ‘trained judgement’, which allowed her to see the underlying patterns in data. These differing concepts do not neatly replace one another as we progress towards the one true view. Rather, they co-exist and jostle in shaping our sense of good representation, and, hence, what counts as objective. Of course, Daston and Galison’s story is contestable, but the moral is simple enough: the concept of objectivity has a complex history, and, as Nietzsche claimed, ‘only that which has no history can be defined’.

1.2 Processes and Products

Despite these complexities, it seems possible to distinguish two main approaches to understanding objectivity: first, as a feature of certain kinds of epistemic products; second, as a feature of certain kinds of epistemic processes (Reiss and Sprenger, 2017). I will explore these two approaches and their relationship, and sketch why, over the last few decades, philosophers of science have increasingly focused on the second.

On the first approach, some epistemic product – say, a factual belief or claim – is ‘objective’ or ‘objectively true’ if it represents some feature of the world as it really is, rather than as it appears to us; in Koskinen’s nice formulation, it is knowledge of ‘the object untainted by the distortions caused by our subjectivity’ (Koskinen, 2018, 3). To motivate this general worry, consider the concern that claims such as ‘grass is green’ are not fully objective, because ‘greenness’ is a ‘secondary quality’, that is one which depends, in some sense, on human observers (Menziés and Price, 1993). The view of objectivity as somehow related to representing reality apart from humans is appealing, but horribly hard to pin down. As such, it is often expressed in metaphorical language: for example, in Thomas Nagel’s metaphor of objectivity in terms of a ‘view from nowhere’, achieved through detachment from ‘the contingencies of the self’ (Nagel, 1986) or in Bernard Williams’ concept of the ‘absolute conception’ of the world (Williams, 1985).

There are multiple versions of the second ‘process’ account of objectivity. However, they typically all have at their core a sense of objective processes as more or less rule-governed activities, relatively immune to features of the individuals who undertake them; for example, consider Stegenga’s characterisation of objectivity in medical science as insensitivity to personal or idiosyncratic biases. On such views, objectivity is primarily a feature of the ways in which we investigate the world, rather than the products of our investigations. Note that, on the process view, the product of some investigation need not count as part of the furniture of the Universe to count as objective. For example, claims about drug effectiveness might not appear in the basic ontological inventory, but be ‘objective’ as long as they result from certain sorts of processes which minimise bias.¹

Distinguishing these two senses of objectivity is central to many philosophical disputes. Consider a non-scientific example: some philosophers worry that moral claims – for example, that murder is wrong – are not objective. Often, such claims involve an (implicit or explicit) juxtaposition with scientific claims. Worries about moral objectivity are often motivated by the (alleged) fact of widespread moral disagreement (Mackie, 1977). Broadly, however, we can distinguish two kinds of moral anti-realist worry. One is that moral claims are not objective in the sense that there are no moral facts to which they can correspond (consider the argument for moral scepticism along the lines that ‘were there moral facts, disagreement would not be so widespread’). The second is that moral claims are not objective in the sense that there is no widely shared criterion for resolving moral disagreement (consider the argument for moral scepticism that ‘were there a widely shared criterion, disagreement would be more easily resolved’). We might respond to the second of these worries – show that there is some criterion which we do, or should, all use to solve moral disputes – without responding to the first – that is without showing that moral properties are out there in nature. For some, establishing such a criterion would suffice for moral objectivity, whereas for others, a worry would remain.

We see similar worries in philosophy of science. Consider a striking example from twentieth century philosophy of science. Thomas Kuhn’s *Structure of Scientific Revolutions* argued for an understanding of science as involving two stages (Kuhn, 1962). During normal science, scientists are engaged in ‘problem solving’ within some paradigm, a shared set of rules, principles and exemplars; during periods of revolution, a paradigm gives way to a second paradigm,

¹ There are multiple ways of distinguishing these two forms of objectivity; for example, see Axtell’s pragmatist distinction between ‘ontological’ *versus* ‘cognitive’ senses of objectivity (Axtell, 2015, 2–4)

radically changing the problems that scientists work on, how they work on them and what counts as success. Kuhn's work is often claimed to have challenged beliefs about the objectivity of science.² However, in these discussions, we often find two distinct issues run together. One set of concerns is that Kuhn's claims about the role of paradigms in shaping scientific practice implies that we always make claims about the world from *within* a particular position. Specifically, Kuhn seems to suggest that, because all scientific research is paradigm-bound, it is extremely difficult to compare paradigms as better-or-worse; what counts as a key problem and as a way of resolving a problem in *my* paradigm may differ fundamentally from what counts as a key problem or successful solution in *your* paradigm. If so, it seems that there is no way of comparing paradigms in terms of which one is closer to the world as it 'really' is. Kuhn's work seems, then, to threaten the objectivity of scientific products (or strictly, our ability to know whether our products are objectively true or not).

A second set of concerns, however, focuses on Kuhn's account of paradigm change. Famously, Kuhn (apparently) suggested that there was no general pattern or logic to paradigm changes, but, rather, that shifts occurred as the result of contingent sociological factors, suggesting a version of what is sometimes called Planck's principle, that 'a new scientific truth does not triumph by convincing its opponents and making them see the light, but rather because its opponents eventually die and a new generation grows up that is familiar with it' (Planck, 1950, 33). As Kuhn himself later acknowledged, in his seminal work he seemed to suggest that the *process* of theory change is not objective, but a matter of mob psychology (Kuhn, 1977).

Of course, Kuhn's claims about how paradigms structure science and about how paradigms change are linked. However, they can be distinguished, with important implications for how we think about objectivity. For example, in later work, Kuhn (1977) presented an account of change based around the idea that scientists choose between paradigms on the basis of their possession of certain sorts of epistemic virtues. Unfortunately, he suggested, given that there are a plurality of epistemic virtues, and no obvious way in which to rank them, different scientists might come to different epistemic judgements about which paradigm to prefer. However, despite these problems, he suggested that the process of paradigm change could be 'objective', as the community could provide judgements about which virtues to favour, and, hence, which paradigms to adopt, that were not swayed by individual investigators' preferences. We might concede that Kuhn's later theory does show that scientific change can be

² For a short but stimulating account of how Kuhn's work relates to broadly Kantian notions of objectivity, see Gaukroger, 2012, chapter 5

‘objective’, in the sense that it is not affected by idiosyncrasies and biases, related, perhaps, to Daston and Galison’s notion of trained judgement. However, even granting this, we might still be worried that Kuhn’s work challenges the objectivity of science, insofar as it implies that our paradigms always structure our experience, such that we can never step back and ask which scientific theory represents the world as it really is.

As the example of Kuhn’s shifting views suggests, it is *possible* to provide an account of objective scientific processes that does not necessarily respond to concerns about the objectivity of scientific products. Still, you might think that when we are talking about scientific objectivity, the product sense is more important than the process sense; that the ultimate aim of science is to help us somehow pull back the filters we impose to show us the world as it really is. From this perspective, Kuhn missed something in his response to his critics: the important worry about his work is that theory change might not get us closer to a true picture, rather than that it is affected by idiosyncrasies. However, although the conception of objectivity in terms of revealing the fundamental structure of the world may still be popular in some circles, it does not seem central to recent philosophy of science.³

An enjoyable article by Elisabeth Lloyd (1995) provides a neat summary of many of the concerns about the notion that science is – or ought to be – in the game of achieving the ‘view from nowhere.’ For Lloyd, what she calls talk of the ‘really Real’ or ‘big O’ objectivity is a ‘philosophical folk story’ (Lloyd, 1994, 353), which fails on multiple grounds. It is premised on an implausible reductionist metaphysics, of the world as arranged in hierarchical layers wrapped up with a discredited religious worldview; even if this metaphysical view is correct, it may be impossible for us, limited human creatures, to see the world as it really is; even if we could achieve the view from nowhere, it would be impossible for us to know that we had. Using the terminology above, Lloyd suggests that the notion of objectivity as the ‘really Real’ is not *viable*. To add to Lloyd’s attack, we might note that even if we can achieve ‘big O’ objectivity, it is not clear that this conception has fidelity with our everyday uses of the term. For example, no-one thinks that the current UK inflation rate is part of the metaphysical furniture of the Universe. If anything depends on human perceptions, interactions and evaluative practices, then ‘inflation’ does. Nonetheless, we can and do talk of more or less objective ways of constructing inflation indices or measuring inflation or reporting changes to inflation. In short, even if

³ It is worth being careful here: outside philosophy of science, there is still a tendency to think that *true* objectivity involves something like talking about what is really there, rather than merely having certain sorts of non-biased processes. Quite why this should be the case is interesting, but not my concern here.

the absolute conception were *viable*, it is not clear that it captures our everyday sense that a wide range of things – not just the basic things – can be represented more or less objectively.

It is now far more common for philosophers of science to conceive of objectivity primarily in terms of a feature of processes. One important driver of this shift is the kinds of metaphysical and epistemological concerns expressed by Lloyd.⁴ A second important driver is the link that many philosophers draw between objectivity and trust. Following a suggestion made by Arthur Fine (1998) in an attack on the Nagel/Williams picture, many philosophers stress the pragmatic functions of objectivity talk as a way of signalling epistemic trustworthy sources of information.⁵ Specifically, Fine's key intervention has two moves: first, he dismisses the notion that we can have any way of getting to the 'really Real' on a wide variety of grounds; second, he suggests that we can best understand objectivity talk in terms of its function, as helping us identify which kinds of scientific process are useful to us. Ultimately, then, he suggests that objectivity is simply 'that in the process of inquiry which makes for trust in the outcome of inquiry'.

It is easy to see how a focus on trust can lead us to a focus on process, by thinking about the problems non-experts face in assessing expert testimony (Goldman, 2001). Imagine that you must decide whether or not to trust some informant. Typically, we turn to informants precisely when we ourselves cannot distinguish true from false claims in some domain; if we knew what was true, we wouldn't need informants! Simply saying that some informant's epistemic product is 'objectively' true does not give us any further reason to believe that informant; it is a bit like responding to 'why should I believe what you just said?' by answering 'because it is really, really true'. By contrast, saying that the product was arrived at by an 'objective' process does give us some reason to accept the informant's claim: we can check whether she really followed that process, we can check whether we think that process really is objective; and so on. Thinking about objectivity as a feature of processes seems to help us in our everyday task of placing and withholding trust, in a way in which thinking about it as a feature of products does not. For practical purposes, then, there are excellent reasons to focus on objectivity of process, rather than product.

⁴ In the name of fairness, as Lloyd herself notes, it is unclear whether even apparent defenders of 'big O' objectivity were ever really committed to very strong metaphysical and epistemological claims

⁵ Although Fine's suggestion that we focus on the link between objectivity and trustworthiness can be detached from any particular metaphysical views about the 'view from nowhere', Fine also had a horse in that race, suggesting that many metaphysical disputes about realism rest on a kind of mistake; see, for example Fine, 1984

The claim that objectivity is bound up with notions of trust provides us, then, with reasons to focus on the objectivity of processes, rather than products. Before moving on, it is worth quickly noting an important ambiguity in the literature on objectivity and trust. Very broadly, we can distinguish two senses of ‘trust’: first, a rich, thick sense where ‘trust’ is bound up with broadly ethical concerns and face-to-face relationships; second, a thin sense, where ‘trust’ means something more like ‘a willingness to rely’. This distinction has been central in broader philosophical debates about the nature of trust (Hawley, 2014). Which sense is relevant to thinking about objectivity? Consider some examples: a scientist writing a paper who must decide whether to use reports from co-authors scattered across the globe (Kukla, 2012); a non-scientist watching a documentary who must decide whether to accept the experts’ claims that climate change is serious (John, 2018). It seems natural to describe both of these cases as involving decisions to place epistemic trust, and as cases where a key question involves the objectivity of informants, but, in neither case does the relevant relationship seem to involve deep, thick ethical assumptions and commitments. Therefore, I suggest that, in thinking about trust and objectivity in science, it is more sensible to interpret trust in the second, thinner sense. Indeed, similar concerns have led Inkeri Koskinen (2018) to suggest that, in analysing scientific objectivity, we should replace ‘trust’ talk with ‘reliance’. Koskinen has an important point. However, as Daston and Galison point out, historically, notions of objectivity have been caught up with notions of virtue. Regardless of whether we *ought* to think of scientific objectivity as having an ethical component, we often *do* think that way. Using the language of trust helps remind us of these important complexities and ambiguities.

With this caveat in place, which processes count as objective, though? Heather Douglas (2004) has distinguished three kinds of processes – processes of engagement with the world, individual thought processes and social processes that produce results – each of which can be characterised in terms of being more or less objective along various dimensions. For example, she suggests that one use of ‘objective’ is to describe processes of engagement with the world that involve ‘manipulation’ of objects (‘manipulable’ objectivity), where manipulability may come in degrees; another sense is to describe social processes that allow for strong interaction and contestation between scientists (‘interactive objectivity’), which, again, can come in degrees; and so on. Perhaps *the* key question in the recent literature is whether there is some common epistemic thread running through the different senses of objectivity beyond a link to trust? Douglas suspects not: that objectivity is an irredeemably complex concept. Others disagree; for example, Koskinen (2018) argues that all senses of objectivity are united in that they are concerned with strategies for the