



The aim of belief and the aim of science

(El objetivo de la creencia y el objetivo de la ciencia)

Alexander BIRD*

King's College London

ABSTRACT: I argue that the aim of belief and the aim of science are both knowledge. The 'aim of belief' is to be identified with the product of a properly functioning cognitive system. Science is an institution that is the social, functional analogue of a cognitive system, and its aim is the same as that of belief. In both cases it is knowledge rather than true belief that is the product of proper functioning.

KEYWORDS: aim of belief; aim of science; functionalism; scientific progress.

RESUMEN: *Argumento en este artículo que el objetivo de la creencia y el objetivo de la ciencia son ambos conocimiento. El "objetivo de la creencia" debe identificarse con el producto de un sistema cognitivo que funcione correctamente. La ciencia es una institución que es el análogo social y funcional de un sistema cognitivo, y su objetivo es el mismo que el de la creencia. En ambos casos, es el conocimiento, más que la creencia verdadera, el producto del correcto funcionamiento.*

PALABRAS CLAVE: *Objetivo de la creencia, objetivo de la ciencia, funcionalismo, progreso científico.*

* **Correspondence to:** Alexander Bird. Department of Philosophy, Philosophy Building, King's College London, Strand, London WC2R 2LS, United Kingdom. –alexander.bird@kcl.ac.uk– <http://orcid.org/0000-0002-6244-1579>

How to cite: Bird, Alexander. (2019). «The aim of belief and the aim of science»; *Theoria. An International Journal for Theory, History and Foundations of Science*, 34(2), 171-193. (<https://doi.org/10.1387/theoria.19351>).

Received: 26 February, 2018; Final version: 18 January, 2019.

ISSN 0495-4548 - eISSN 2171-679X / © 2019 UPV/EHU



This article is distributed under the terms of the
Creative Commons Attribution 4.0 International License

1. Introduction

Aristotle begins his *Metaphysics* by telling us:

All people by nature desire to know.

Science is the institutional embodiment of this desire.^{1,2} In this paper I argue that the constitutive aim of science is the production of *knowledge*; in particular I argue that the aim of science is knowledge rather than truth.

The core argument for the knowledge aim of science draws a parallel with the aim of belief. Later in the paper, I argue that for belief to have an aim to be governed by norms of correctness stemming from the fact that our belief-producing (i.e. cognitive) systems have a function. A correct belief is one produced by a properly functioning cognitive system. And a belief produced by a properly functioning system is knowledge. So beliefs aim at knowledge.

In order to draw the same conclusion for science, I first propose, following Émile Durkheim and Talcott Parsons, that we should understand institutions such as science as having functions. Science, in particular, has a cognitive function.

In passing I note that linguistic evidence and our intuitions about particular cases also support the conclusions of the functionalist arguments I develop. For example, our intuitions about which developments constitute successful, progressive science, thereby indicating what amounts to meeting the aim of science, suggest that knowledge but not anything less than knowledge meets the aim of science. Likewise, our intuitions about correctness of belief also do not acknowledge beliefs less than knowledge as correct.

2. The aim of science

I claim that the aim of science is the production of knowledge. This claim has two elements. First that science has a single aim at all. And secondly that the epistemic status of what science aims at is knowledge. Here I am principally concerned with establishing that

¹ As it stands, this slogan is an exaggeration. We can desire to know all sorts of things, and we might institutionalise this desire. The fact that Wisden in an institution aiming at delivering knowledge of cricket statistics to fans does not make Wisden a form of science. So the kind of knowledge that is in question needs to be circumscribed. Hoyningen-Huene (2013) offers an account of scientific knowledge in terms of systematicity that is relevant here. Furthermore, ‘science’ must be understood in a broad sense to include the humanities and other disciplines that aim at knowledge, rather as ‘Wissenschaft’ is understood in German. If we want to restrict ‘science’ to the natural sciences, then further restrictions will need to be made. In this paper I am more interested in what science has in common with other knowledge-producing institutions, and less in what distinguishes science from them.

² Is science essentially social, as this claim implies? Can’t it be conducted in a solitary way? I do not think so. Yes, a scientist might carry out solitary fieldwork but this will be pursued against a background of socially produced and acquired knowledge. Perhaps someone without any scientific background could engage in an activity that looks much like science. We might call that activity ‘scientific’ to mark that fact. But they would be engaging in proto-science at most, not science itself.

science aims at knowledge rather than truth. Other possible epistemic aims, such as understanding or 'problem-solving' will need to be discussed elsewhere. This discussion is for the most part orthogonal to a discussion of the *scope* of science's aim. Van Fraassen argues that science aims at the empirical adequacy of its theories. This in effect restricts the scope of science to the observational. Nonetheless, my question still arises. Does science aim at theories with observational consequences that are true? Or does it aim at theories whose observational consequences are known to be true? My discussion applies also to deciding between those alternatives.

I will first argue that it is plausible to think of science as having a single aim and that this is an epistemic aim. I then note the suggestive linguistic evidence that science is concerned with producing knowledge rather than true belief. I next ask, what counts as success in science? Science is successful when it produces knowledge, not merely true belief. Hence, science has the aim of producing knowledge.

This will motivate and support the conclusion of section 3, in which I elaborate on the idea of science as an institution and argue that it has a cognitive function.

2.1. THE MULTIPLE AIMS OF SCIENCE AND OF SCIENTISTS

The aim or aims of science are not to be read off directly from the aims of individual scientists. The aims of the latter are varied. Some scientists are motivated by the nature of the work itself. Others may seek the fame of a Nobel Prize. Yet others might be looking for the practical applications of their research and for lucrative patents. Some but not all may have lofty ideals of expanding our knowledge of the world. None of this conflicts with the claim that there is a single constitutive aim of science. Institutions have aims that may well differ from the aims of those working in them. The armed forces have the aim of protecting the nation from its external enemies; the health service has the aim of maintaining and restoring the health of citizens. Yet soldiers and sailors, nurses and doctors may have a multitude of intentions in taking on these roles, whether it is seeking camaraderie, an exciting profession, personal fulfilment, or just needing a reliable income. A commercial enterprise might have the aim of making profits for its shareholders, without any of its employees having that aim—so long as it is structured in such a way as to incentive employee behaviours that in combination are disposed to be profit-making.

The complexities of the institution of science and its interactions with other institutions are nevertheless consistent with its having a single constitutive aim. The army puts on parades for visiting heads of state, the health service builds hospitals, and both employ large numbers of people. But none of these is a constitutive function of the institution in question. An army that didn't parade would still be an army. But if it did nothing towards defending the nation, because it lacked troops for example, it would not be an army any longer. The health service builds hospitals because these are necessary for treating the sick. If a health service built hospitals as an end in itself and didn't treat any patients in them, then it would be a health service in name only. So organisations and institutions can have constitutive aims or functions, even though there are many different kinds of thing that they actually do. The things that they typically do are means to the ends that are their constitutive aims (building hospitals), though they need not be (parading for dignitaries). Likewise science and scientific organisations can engage in activities that are not constitutive of science and its aim. Much science aims at producing technology for various social

or commercial purposes. So not only have scientists studied the mechanisms of photosynthesis, but they have also sought to replicate photosynthesis artificially in order to generate energy from sunlight in an environmentally sustainable way. Although science in fact does much of the latter kind of thing, it could do fail to do anything of the sort and still be science. Most ancient science and a large proportion of current science has no practical or technological application. Although it was part of Bacon's depiction of Salomon's House that through its acquisition of knowledge it contributed to the economy of Bensalem (the mythical nation he describes in *The New Atlantis*), it was a complaint levelled at the Royal Society that it failed to live up to this expectation.⁴ But that failure was not thereby a failure of the Royal Society to do science. On the other hand, had the Royal Society made no effort to extend knowledge, then it would not have been a scientific society. The Royal Society could have instructed Robert Hooke to buy equipment and to devise experiments purely for the ephemeral entertainment of its fellows, with no intention of understanding why such things happened or of generalizing from them to propositions about nature. Had it done so then whatever it would have been doing would not have been science.

2.2. SCIENCE AS KNOWLEDGE—THE LINGUISTIC EVIDENCE

We should first note the linguistic evidence in favour of the view that the nature and aim of science concern knowledge. The Oxford English Dictionary gives a definition of science as 'The kind of organized knowledge or intellectual activity of which the various branches of learning are examples'. This definition is a disjunction of two ways of conceiving science: as a set of propositions and as an activity. Regarded as a special set of propositions, it is notable that the OED describes science as organized knowledge, not as organized belief. With this propositional conception in mind, it is much more common to use 'knowledge' than 'belief' to describe science. 'Science is organised knowledge', also asserted by Herbert Spencer (1911, 61-2), has a naturalness not shared by 'science is organised belief' or even 'science is organised true belief'.⁵ If science, under the propositional conception, is a body of knowledge, then how should we think of science when conceived

³ Douglas (2014), however, argues that dissolving the pure/applied distinction for science allows for a better understanding of scientific progress.

⁴ Bacon (1620) writes, 'We have three [fellows of Salomon's House] that bend themselves, looking into the experiments of their fellows, and cast about how to draw out of them things of use and practise for man's life.'

⁵ Such instances may be multiplied, as exemplified by these definitions: from the UK's Science Council 'Science is the pursuit and application of knowledge and understanding of the natural and social world following a systematic methodology based on evidence' (<https://sciencecouncil.org/about-science/our-definition-of-science/>); from the American Physical Society 'Science is the systematic enterprise of gathering knowledge about the universe and organizing and condensing that knowledge into testable laws and theories' (https://www.aps.org/policy/statements/99_6.cfm); Encyclopaedia Britannica 'Science, any system of knowledge that is concerned with the physical world and its phenomena and that entails unbiased observations and systematic experimentation. In general, a science involves a pursuit of knowledge covering general truths or the operations of fundamental laws' (<https://www.britannica.com/science/science>) (the above all accessed on 10/08/2018). While we may quibble about the details of any of these definitions, it is notably that how natural it is to use the term 'knowledge' to characterise science, and, in the last definition, the goal of science.

of as an activity? The intuitive answer is that science-as-an-activity aims at achieving or augmenting science-as-a-set-of-propositions. That is, the activity of science aims at producing or adding to the body of scientific knowledge.

Etymology also provides linguistic evidence for the knowledge view of the aim of science. In English and the Romance languages, the relevant terms are derived from the Latin *scire* meaning *to know*: *science* (English), *la science* (French), *la scienza* (Italian), *la ciencia* (Spanish). In Germanic (including Nordic) languages, the terms for science all relate to those languages' words for *know*: *Wissenschaft / wissen* (German), *vetenskap / veta* (Swedish), *wetenschap / weten* (Dutch), *vísindi / vita* (Icelandic); likewise in the Finno-Ugric languages: *tiede / tietää* (Finnish), *tudomány / tudás* (Hungarian); and in Irish science is *eolaíocht* and I know is *Tá eolas agam*. We find the same pattern beyond Europe: *bilim / bilmek* (Turkish), *uloom / ilm* (Arabic). The Farsi (Persian) word for science, *elm*, derives from the Arabic for knowledge, *ilm*. The words *madda* and *episteme* denote science in modern Hebrew and modern Greek respectively, but refer to knowledge in the ancient versions of those languages. Not all languages make the link between science and knowledge directly—in most Slavic languages the word for science is linked to the word for teaching or for learning, though Czech (*věda / vědět*) and Croatian (*znanost / znati*) are like the Germanic languages in this respect. Like Polish, Russian, and Serbian, Chinese links science to learning, rather than to knowledge *per se*. But in no language of which I am aware is there a link between the words for science and for truth.

Linguistic evidence provides no knock-down argument. But it is highly suggestive. Supporters of the truth view of the aim of science need not deny that science delivers knowledge (although some do). They can maintain that those who aim at truth will want to use reliable methods of achieving that truth and when they do, their beliefs will be knowledge. That makes knowledge a kind of by-product of science. If that is the case it is odd that our words for science are so closely related to our words for knowledge, but never to our words for truth. Why should our words for science derive from its by-product? It is much more plausible that those words link to knowledge because knowledge is the essence of science.

2.3. PROGRESS AND KNOWLEDGE

When is science successful? What counts as good science? Because the aim of an activity sets conditions for the success of that activity, answering those questions will provide an insight into the aim of science.

Imagine a scientific community that tests its theories by an unreliable method (such as astrological divination). By using that method it comes to believe theory T; coincidentally T is true. Does that count as successful science? René Blondlot and a good number of other French scientists believed that he had discovered a new kind of ray, hitherto unknown to physics—N-rays. In due course it became clear that these scientists were subject to a form of self-deception, in large part fostered by a patriotic desire to further the reputation of French science. There are of course no N-rays. But now imagine a variation on the actual story, whereby it turns out that there are in fact rays corresponding to Blondlot's theoretical description of N-rays. Call this theory 'T'. But this is a coincidence—Blondlot's methods and reasoning gave him no good reason to think that there are such rays. So Blondlot correctly believed significant novel truths of the form: there are rays with the property φ .

But he had no justification for those beliefs and so they do not count as knowledge. Would Blondlot's science have been successful under those circumstances? Would it have met the minimum standards for good science?

It seems clear that Blondlot's science would not have been good or successful science in such a scenario. So the truth of T is not enough. Clearly what is lacking is adequate justification. Good science requires not just that our theories be true but that we have adequate grounds for believing them, that our methods are reliable, and that they do not rest on self-deception or other irrational causes. The best explanation for those additional conditions is that they are required for science to be knowledge.

Elsewhere I have argued that science progresses when it accumulates knowledge (Bird 2007a; Bird 2008; see Rowbottom 2008, 2010; Cevolani and Tambolo 2013; Niiniluoto 2014; Dellsén 2016 for critical responses). If scientific progress were a matter of increasing truth or truthlikeness, then, under the imaginary scenario where T is (coincidentally) true, the French scientific community at first made progress by believing in T. Furthermore, when it gave up belief in T, it then regressed. But that is not correct. The community made no progress by coming to believe T on unreliable grounds. And giving up an unreliably tested theory is not regressive—it is, if anything, progressive. Again, the best explanation of this verdict is that progress requires not true belief but knowledge. Because T was not known to be true, adding a belief in T was not progressive and giving up a belief in T was not regressive. So scientific progress is the accumulation of knowledge rather than truth. Progress in some activity is related to the aim of that activity thus:

If an activity A aims at goal X, then A makes progress insofar as it gets closer to achieving X or does more of X or does X better.

So if science aimed at truth, then science would be making progress by adding to truth or truthlikeness. The latter is false, as we have just seen. By *modus tollens* therefore, it cannot be that the science aims at truth. On the other hand, if science aims at knowledge, then scientific progress would be the accumulation of knowledge. So we should conclude that science aims at knowledge rather than truth.

3. *Science as a cognitive institution*

Although the institution of science is complex and multi-faceted, I have argued that it is plausible that it has a constitutive aim, the production of knowledge. In this section I examine in more detail what it is for science to have an aim. I first sketch reasons for thinking that science should be regarded as an institution that is itself the possessor of knowledge and furthermore that in virtue of science knowing things, societies, whose members are mostly not scientists, can know the same things. These conclusions support (but are not essential to) the argument that follows, that science has a cognitive *function*. Drawing on sociological theory, I propose that science stands to society much as a cognitive faculty stands to an organism; the accepted theories of science thus correspond to belief. This sets us up for a discussion of the aim of belief in the following sections. The aim of belief, I will argue, is knowledge, thus supporting the conclusion that the aim of science is knowledge,

3.1. SCIENCE AS A KNOWING INSTITUTION

Groups can be the possessors of knowledge and belief in a way that does not reduce to a summary of the corresponding propositional attitudes of the members of those groups (Gilbert 1987; Tuomela 1992). Indeed the knowledge of a group does not even supervene on the mental states (including knowledge) of the group's members (Bird 2010b). Such groups include scientific groups, such as a research team. They can also be larger groups, such as a scientific community defined by field. Groups or collectives that are epistemic agents can be larger still. The institution of science as a whole can be regarded as such a collective entity. As such it is often the subject of knowledge attributions, particularly of the form 'Science knows that p '.⁶ Furthermore, in virtue of the discoveries of science, we make statements of the form 'We know that p '.⁷ One might think of the relationship between 'science knows' and 'we know' as a case of testimony. Statements such as 'Science tells us that the universe is over 13 billion years old' would suggest this. But, I think, given certain conditions regarding the close relationship between science and the rest of society, this relationship is constitutive. That is, in virtue of science itself knowing that p and science having a certain role in society, then society at large knows that p .

Here is a related case. We say 'North Korea knows [*or* the North Koreans know] how to build an atomic bomb'. Yet, only a restricted group of scientists and engineers has any direct part in this knowledge.⁸ Normal North Koreans and even leading members of the North Korean government are not members of this restricted group. Yet they are included in the collective (the North Korean nation or people) to whom the attribution of knowledge is made. So we have a case where a broader collective can be said to know something in virtue of the fact that some sub-collective knows something, without there being any transfer of knowledge by testimony.⁹

The usual relationship between science and society is not like that between North Korean nuclear weapons engineers and the North Korean people. In the usual case, scientific knowledge is typically made available so that individuals and organisations in wider society can access and make use of it. Of course, most individual members of a society will not

⁶ Examples of science as the subject of knowing (or not knowing): James Franklin's (2009) book *What Science Knows and How It Knows It*; 'Why science will never know everything about our universe', article in *Forbes* <www.forbes.com/sites/startwithabang/2017/02/10/why-science-will-never-know-everything-about-our-universe/#76932ff55513> accessed 12.02.2018; 'What science knows about why people are gay', article in *Newsweek* <www.newsweek.com/what-science-knows-about-why-people-gay-751927> accessed 12.02.2018; '9 things science knows about baseball' article in *Discover* <discovermagazine.com/galleries/zen-photo/b/baseball-science> accessed 12.02.2018.

⁷ For example: 'Science, evolution, and what we know: 2018 edition', article in *The Huffington Post* <www.huffingtonpost.com/entry/science-evolution-and-what-we-know-2018-edition_us_5a5bc6e4b0a233482e0cb7> accessed 13.02.2018; 'What we know', American Association for the Advancement of Science website on climate change <whatwknow.aaas.org> accessed 13.02.2018; 'How do we know what stars like our Sun are made of?' BBC website <www.bbc.co.uk/earth/story/20160128-how-do-we-know-what-stars-like-our-sun-are-made-of> accessed 13.02.2018.

⁸ And even then they only have this knowledge collectively rather than individually.

⁹ If there are concerns that this case is about knowing *how* then we can replace 'knows how to build an atomic bomb' with 'knows that p ' for some proposition p that is central to atomic bomb making.

be able to understand let alone themselves make use of the esoteric discoveries of science. Nonetheless, they are connected to those discoveries in numerous ways. They may be consumers of products designed using science-based technology or as citizens subject to government policies devised on the basis of social science research. The North Korean example can be seen a special case where the North Korean people (in the view of its government) benefit from the social need for military power that is met by possessing knowledge of nuclear technology. The connections are not just through technology and applied science. We draw on science as consumers of popular science in books, newspaper articles, and television programmes. And in principle those with sufficient resources and aptitude can study science so that they can directly access the knowledge that science produces. Given the multiple ways in which the knowledge that science produces affects or is available to wider society, for many purposes we can regard the 'we' in 'we know that *p*' as encompassing all or most of humanity.

3.2. SCIENCE AS AN INSTITUTION WITH A COGNITIVE FUNCTION

According to a conception of institutions derived from Émile Durkheim (1893; 1894) and Talcott Parsons (1961), we should see society as analogous to an organism. Institutions are like the organs and physiological systems of an organism. And like the organs of an animal or plant, institutions have functions. And such functions can be constitutive of the organs in question. A wing has the function of enabling flight. And it has that function essentially. That function determines which part of an animal is its wing. We can identify the wings of an albatross and the wings of a bee as wings, because they are analogues even though they are not homologues. That is, they are both wings because they perform the function of enabling flight.¹⁰ Likewise, what determines the identity of an institution, such as the law or the army or the higher education sector, is the function it performs in society.

Herbert Spencer (1874) suggested functional analogies between particular functional systems in animals and in societies, for example between the central nervous system and the government, since both played an overall regulatory function, coordinating the interactions of other organs/institutions. While no-one would suggest that every animal organ has a social analogue or vice-versa, it is nonetheless possible to see some parallels as sufficiently strong that we may hypothesize that they belong to the same functional type. Below I shall discuss the cognitive faculties of a person, how they are integrated, and how their function is to produce knowledge (or true belief). In parallel, we can say that various social institutions and practices perform cognitive functions. That is, they serve to collect, generate, and distribute information of a certain kind. Governments, for example, need information in order to carry out their functions. Therefore

¹⁰ It is *not* that they occupy the same locus in the structure of the two organisms: they do not. A bat's wings and a bird's wings and a human arm are homologues, because they are evolved from the same structure in a common ancestor. Whereas a bee's wing has no evolutionary ancestry in common with those wings; it is a wing nonetheless. Penguins, therefore, do not have wings. Zoologists describe their forelimbs as flippers, as in a dolphin, since the function they perform is that of propelling and guiding the animal through water. They are vestigial wings. But vestigial wings are no more wings than a retired postman is a postman or someone's ex-spouse is still (thereby) a spouse.

governments have created various organisations whose function is to collect that information, from spy agencies to census bureaus. Non-governmental organisations, such as businesses and charities, likewise have divisions whose function is to create or collect relevant kinds of information.

So there are formally created institutions or sub-components of institutions with cognitive functions. Science as a whole, however, is not a formally created institution. Science has grown up more or less informally, across national boundaries, inspired in large part by the curiosity of individuals. That organic development of science does not preclude it from having a social function. Above I characterized science as the institutional embodiment of the human desire to know. The desire to gain knowledge from others, as well as the prestige associated with displaying one's own knowledge, led to the exchange of knowledge and thence, as Kuhn (1962) describes, also to the sharing of methods and values. This leads to the development of an institution, that is to say a group whose members show a degree of social cohesion; they are bound together by mechanisms of what Durkheim called 'solidarity'. On the one hand there is organic solidarity arising from a division of intellectual labour—scientists depend on the results produced by other scientists.¹¹ And on the other there is mechanical solidarity—scientists share values and methods, and a conception of themselves as scientists (or, earlier, 'natural philosophers' or 'men of science'). Although it is itself an informal institution, science from the very beginning has been pursued in formal institutional settings, from Aristotle's Lyceum and the Mouseion at Alexandria through the madrassas and universities of medieval Europe. And states took an interest, supporting and creating scientific institutions of their own, both for reasons of prestige as well as in the hope that the knowledge generated might be of practical benefit—as increasingly it came to be.

As we shall see, the outputs of a properly functioning cognitive faculty of an organism are varied in the use to which they are put: knowledge as an input into practical reasoning, leading to action; as an input to theoretical reasoning, leading to further knowledge; knowledge to be stored in memory (short term or long term). Likewise the outputs of science are varied in use: knowledge for practical purposes, as evidence for theorizing, and to be added to the store of shared knowledge. Nonetheless, a common nature underlies the variety of uses of these outputs is: the outputs all ought to be knowledge, and will be if the systems in question are functioning correctly.

In conclusion, then, institutions in a society have functions just as faculties and organs within an organism have functions. While there need not be any specific parallels at a greater level of detail, it turns out that there are faculties and institutions which have a cognitive function. The commonality between faculties with cognitive functions and institutions with cognitive functions is not an accident: the fundamentality of information collection, production, and processing for any complex biological or social entity makes this inevitable (Burgin 2009). In the next section I therefore look more closely at cognitive faculties in an individual organism.

¹¹ One does not want to overstate parallels between science and commerce. Nonetheless, it is intriguing that commerce is another largely informal, supra-national institution characterised by division of labour, linking individuals many thousands of miles apart, that developed early on through a process of exchange.

4. *Belief has an aim*

In this section I articulate an approach to answering the question ‘what is the aim of belief?’¹² In the next section I continue the argument to show that the aim of belief is knowledge.

The key claims of this section are these:

- The normativity implicit in ‘the aim of belief’ is teleological: the teleology of belief sets correctness conditions for belief.
- That teleology derives from the fact that cognitive systems have a function.
- The function of a cognitive system is to produce true belief

Note that I do not claim that the normativity deriving from the functional teleology of belief has any implications for a deontology of belief. That is, they do not tell us (not directly anyway) what we ought or ought not believe.¹³

4.1. TELEOLOGY AND THE AIM OF BELIEF

Since Bernard Williams (1973) first used the phrase ‘the aim of belief’ philosophers have sought to articulate precisely what it might mean. On the one hand, one may take the term ‘aim’ seriously, as implying a goal or intention. Such is the teleological (or ‘teleologist’) interpretation. On the other hand one may prefer to understand ‘aim’ more metaphorically, as referring to a *norm* of belief (Wedgwood 2002; Gibbard 2005). This is the normativist interpretation.

While these are the standard opposing views about the aim of belief, that should not be taken to imply that the teleologist must deny that there are any norms of belief. The ordinary term ‘aim’ implies both normativity and teleology. If a darts player aims at the triple twenty, then he intends the dart to hit that part of the board; the purpose of his throw is to get that score. That is the teleology of ‘aim’. If he hits the triple twenty, his throw was successful; if he hits the triple five, then his throw failed. This is the normativity of ‘aim’. The two are related, since what counts as success or failure (normativity) of the throw depends on the intention of the player (teleology). Had the player intended to hit the triple five, then hitting the triple five would have been a successful shot.

While the teleology of aim generates norms of success and failure, not all norms come from teleology. A person’s behaviour is governed by ethical norms, even if they have no intention of acting ethically. The same goes for constitutive norms—norms that govern the nature and existence of some process. The rules of chess, for example, tell us what chess is and determine whether some activity constitutes playing a game of chess (Wedgwood 2002). Such norms, like ethical norms, allow us to evaluate an outcome independently of

¹² I have previously sketched this functionalist approach in Bird (2007, 2010a). Papineau (2013) holds that the ‘aim of belief’ can be understood in terms of biological function, while Sullivan-Bisset (2017) takes a biological functionalist approach to epistemic normativity. For a related view see Millikan (1993).

¹³ McHugh (2014) and Kiesewetter (2017) have argued that a functionalist account of the kind I offer cannot account for the normative demands of rationality.

a subject's intentions. If a player is bribed to throw a match, they still count as winning if they do what the rules say constitutes winning, despite lacking an intention so to do.

The norms governing belief *qua* belief set the conditions for the correctness of a belief.¹⁴ So when it is claimed that a belief is correct precisely when it is true, that claim is not as trivial as it sounds. One might think that belief *per se* does not generate any norms.¹⁵ Or one might think, as I do, that it generates norms beyond truth. If there are such norms these may not only be constitutive of belief—they may also individuate belief (Fassio 2018). That is, they may not only be necessary conditions of belief they may also be sufficient to distinguish belief from other attitudes. I take constitutive and individuating norms to describe essences. It is the essence of belief that it is governed by these norms. I shall argue that a belief is not correct (does not satisfy all the norms of belief *qua* belief) unless it is also knowledge. I note at this point that the truth norm does not seem to be individuating, since guessing is also subject to the truth norm (Owens 2003). This criticism does not affect the knowledge norm for belief, since guessing is not subject to the knowledge norm. A guess is successful if the proposition guessed to be true is in fact true; it does not need to be known—indeed a guessed proposition cannot be knowledge.

So the debate between the teleologists and the normativists may be framed as one regarding the source of norms—whether they derive from something teleological, such as an intention, or from some other non-teleological source (or are basic).

Intention, however, does not seem to be the source of the norms of belief. I have mentioned that the most widely favoured norm for belief is truth whereas I think the correct norm is knowledge. Often neither truth nor knowledge is the intended outcome of believing. For one thing, our beliefs are not typically subject to our will, as Williams (1973) pointed out; they usually just occur, most obviously in ordinary perception. Since one cannot ascribe an intention to a subject's believing in such a case, it cannot be that any norm is generated by such an intention. Furthermore, insofar as we can influence the contents of our beliefs, we might not always intend them to be true. If Jane fears that she might be captured and subjected to an interrogation with a reliable lie-detector or truth drug, she may intend that certain beliefs she has about sensitive matters should be false, and she might try to bring that about. Wedgwood (2002) holds that even in such cases, where falsity is intended and may even be morally good, it remains the case that the truth norm still applies and the false belief fails to satisfy it. Thus those involved in this debate do not think that the norm of belief can be derived from the intentions of the believer.

For this and other reasons, most discussions of the 'aim of belief' have assumed that the normativity of belief's aim does not arise from teleology, and so have preferred normativism. Nonetheless, teleology may still be the source of the norms of belief, for not all teleology is a matter of intention. Velleman (2000) asserts the thesis that 'belief is an accept-

¹⁴ One might draw a distinction between *standards* of correctness and *norms*, as does Sullivan-Bisset (2017), on the ground that norms generate oughts whereas not all standards do (see also Bykvist and Hattiangadi 2013). If that is a valid distinction, then I am concerned with standards of correctness when I talk of norms.

¹⁵ Papineau (2013) denies that there are any distinctive norms of belief. But he takes a norm of belief to be something that generates an obligation on a potential believer, whereas I take a norm of belief to be a standard of correctness, generated by the (biological) function of belief-producing systems. The latter view is neutral on whether such norms also lead to obligations on believers.

ance regulated in ways designed to ensure that its content is true'.¹⁶ We can understand this as expressing a teleology that is independent of intention. Consider someone who has accepted Pascal's conclusion in his famous wager. Like Jane, he intends to believe something that he in fact strongly doubts (that God exists). Pascal's advice was to ensure that he is exposed only to the evidence (including the testimony of theists) that favours belief in God. So although his intention aims at falsity, he exploits the fact that his belief-forming mechanisms are, as Velleman says, oriented towards ensuring that their contents are true (e.g. they are determined by the available evidence). And, clearly, the mechanisms of perception may be regarded as regulating the resulting belief in a way designed to ensure a true perceptual belief. Here 'designed' is not a matter of intention, but a matter of a function brought about by natural selection.

4.2. A FUNCTIONALIST ACCOUNT OF BELIEF

Teleology can be a source of normativity without deriving norms from intentions. The teleology of functions also generates norms. If the pineal gland does not produce melatonin (leading to sleep disorders) then the pineal gland is not functioning correctly; it is not doing what it ought to do. While Aristotle also saw teleology in nature, a naturalistic understanding of biological function is available in the light of our knowledge of evolution, as I discuss below. This fact is, perhaps, under-appreciated by those who reject a teleological approach on the ground that we can fail to intend our beliefs to be true. With the teleology of functions in mind, we may instead interpret 'the aim of belief' as meaning (at a first approximation) 'the function of belief-producing (i.e. cognitive) faculties'. On this view the use of the term 'aim' is not entirely metaphorical. For the teleological aspect of 'aim' is genuinely present in the function of a cognitive faculty.

Organisms have systems that cause them to respond to environmental stimuli. Examples include the tropisms of plants and simple animals, such as the chemotropism of plants that directly causes their roots to grow towards beneficial soil minerals. While such organisms can be said to carry information about their environments, just as a simple thermostat carries information about temperature, they cannot be said to process that information, since the information is not stored in a manner that makes it available for any purpose other than the stimulus-response system of which it is a part. More complex animals, by contrast, are able to integrate information from multiple stimuli. In the most sophisticated instances, most notably in humans, but also in higher primates, some other mammals, and some birds, this integration constitutes practical reasoning—using information (consciously or unconsciously) to make decisions about how to achieve goals. Such integration via reasoning requires that the information in question is represented in a common fashion. Practical reasoning typically involves putting together novel combinations of inputs in the form of perceptual information from several senses and information stored in memory with desires of different sorts to produce an action or intention to act. This would not be possible if the contents of all these states were represented in different ways—if the differ-

¹⁶ Velleman's earlier view has developed in the light of criticism by Shah (2003), leading to their joint view in Shah and Velleman (2005). The latter leans more towards normativism. For an alternative teleological view which resists Shah's arguments see Steglich-Petersen (2006).

ent systems produced outputs of entirely distinct types. We can therefore regard these information-generating systems as having a common function—the function of producing a certain kind of output that is the raw material for the processes of reasoning. Having the same function, these are systems of a common type: *cognitive* systems or faculties.

Reasoning in this sense may be conscious or unconscious, reflective or unreflective. Even near instantaneous responses to new stimuli can be the results of reasoning, if they are suitably responsive to background states of the subject (which distinguishes them from reflex actions). In other cases the reasoning may be complex and need not lead directly to action. In these cases a process of reasoning may lead to an output that is of the same kind as its input. That is to say, it is an output that can itself be used as an informational input for practical reasoning. Such processes are instances of theoretical reasoning. The capacity for theoretical reasoning is thus another cognitive faculty. The outputs of theoretical reasoning need not be used only as inputs for practical reasoning, since they can be used as inputs for a further process of theoretical reasoning. And the outputs of theoretical reasoning can be stored in memory, either because the process of reasoning is extended in time (short-term memory) or because those outputs might be useful as inputs at some later time (long-term memory). Memory is thus another cognitive faculty—although it does not generate information, its function is also to output information for the processes of reasoning.

The concept of ‘function’ in use here is the concept of biological function. Just as it is the function of the pineal gland to produce melatonin, it is the function of our cognitive systems to produce a certain kind of beneficial output. Accounts of biological function (Wright 1973; Millikan 1989; Neander 1991; Walsh 1996) link a system’s function to biological fitness. The function of a system (organ, faculty, . . .) is φ if by φ -ing that system confers an advantage in a regime of natural selection.¹⁷ By producing *true* beliefs in the manner described above, human and animal cognitive systems give their possessors selective advantages over organisms that lack such systems or which have similar systems but which do not produce truth.¹⁸

4.3. AN OBJECTION: HEURISTICS AND BIASES

An objection to this argument for true belief as the function of cognitive systems argues that fitness might not be best served by truth. No cognitive system is perfect, and so the fitness conferred by a cognitive system will be determined not just by the value of accu-

¹⁷ Debates among philosophers of biology focus on whether the selective regime in question is located in the distant past, the recent past, or the present. Which view is correct need not trouble us here.

¹⁸ Sullivan-Bisset (2017) argues that cognitive faculties also have functions other than truth production. Some belief-forming mechanisms, such as those produced by self-enhancement bias and self-deception, can be adaptive because they assist the effective functioning of the person in a different way, aiding ‘self-organization’. This does not detract from the claim that the primary function of a cognitive system is truth. Biological systems with secondary functions are common. For example, a secondary endocrine function of the heart is the production of atrial natriuretic peptide (ANP), which instructs the kidneys to reduce sodium reabsorption when there is a sudden increase in blood pressure or volume. Pumping blood is still the constitutive function of the heart—producing ANP is tacked on to this function. Similarly, even if some belief-forming mechanisms have a non-epistemic function, it remains the case that the function of cognitive faculties is truth.

rate beliefs but also but the costs of inaccurate beliefs. Furthermore, there are costs associated with having or implementing highly accurate systems, and these too will have an effect on what an optimally fit system does. For example, an important use of cognition for our ancestors (in many ancestral species) will have been to tell us whether a predator is nearby. The costs of a type-II error (a false negative) will be considerable (death, serious injury), whereas the costs of a type-I error (false positive) will be rather less (effort spent in moving away, some opportunity costs in feeding etc.). Consider circumstances in which the sensory stimuli are marginally better correlated with the absence of a predator than with its presence. Given the imbalance of costs, fitness will be served by the creature's cognitive system recording the presence of a predator. On the other hand, the system would maximise truth by recording the absence of a predator (since that verdict will be right more often when the stimuli are marginally better correlated with absence). Additionally, fitness will not be served by employing additional time and effort to get the most accurate verdict, since the predator might seize its chance before this process is complete—better to respond quickly to a few salient indicators and beat a hasty retreat if they are present. The heuristics and biases research programme suggests that the demands of fitness lead to belief-forming dispositions that do not maximise the probability of our beliefs being true.¹⁹

It is important here to distinguish a heuristic from a biased mechanism.²⁰ A heuristic assists the believer in getting to the truth, but is imperfectly reliable because the costs of greater accuracy may outweigh the benefits. But that fact does not refute the proposition that the heuristic aims at truth. Any heuristic will have less fitness value than another that is more accurate but has the same costs of implementation. Nonetheless, would it not be correct to say that the function of the heuristic is not 'to produce true belief' but is something like 'to produce true beliefs subject to cost constraints'? That is not correct, since functions should be defined in terms of their output (or output-input relationship), and efficiency constraints are not elements of the output. Note in any case that any function is subject to such constraints. Any organ could be better at its function. In many cases it will be that the better version has not evolved because although its better functioning would per se be an addition to fitness, its resource needs would reduce fitness elsewhere. For example, it might be that a human heart twice as big would be much better at pumping blood, which might make one better at running. But the large heart might therefore require more calories to operate than (in our ancestors at least) could be gained by the improved ability to hunt that better running provides. Even so, one would not deny that it is strictly correct to say that the heart has the function of pumping blood *simpliciter* on that ground.

On the face of it, the objection from biased cognitive mechanisms cannot be dismissed in this way. For a biased mechanism may be more fit than an unbiased one, even when the

¹⁹ The nature of these heuristics and biases is contested. See, for example, Tversky and Kahneman (1974); Gigerenzer (1991); Baron (2014).

²⁰ Although the concepts are distinct, a mechanism can both be a heuristic and a bias. Indeed one might expect heuristics to be sometimes biased. A heuristic will accept a moderate rate of error. If errors of one sort are more costly than errors of another sort, there may be benefits from those errors being biased.

costs of implementation are the same. The biased mechanism has evolved the way it has, not despite its inaccuracy, but in order to be inaccurate in a certain way.

The answer to this objection starts by reminding us that it is always the *truth* of the animal's belief that contributes to its success. It never benefits from falsely believing that a predator is present when it isn't. And so on biological conceptions of function, which link function to what contributes to fitness, the function of a cognitive system must be to produce true belief. For example, Denis Walsh (1996, 564) proposes that, 'The/a function of a token of type X with respect to selective regime R is to m iff A 's doing m positively (and significantly) contributes to the average fitness of individuals possessing X with respect to R .²¹ Since believing what is true will contribute to fitness and believing falsely does not, a cognitive system has the function of producing true belief.

If the function of a cognitive system is to produce true belief, then how has it evolved to have a bias? As we saw, the value of bias arises because a cognitive system is imperfect, delivering false belief on some occasions. So bias is valuable only because cognitive systems may fail in their functions. Some physiological systems are evolved as fail-safe systems, which come into play once a failure of the primary system is detected. For example, the heart has a natural pacemaker in the sinoatrial node; if the signal from the sinoatrial node fails, then cells in the atrioventricular node will act as a backup pacemaker. The bias of a cognitive system is not a backup, but is a pre-emptive 'modification'—though strictly it is not a post-hoc modification but a pre-existing feature that has been retained.

Cognitive systems have evolved from stimulus-response systems. The latter do not produce beliefs, just behaviour. In a sense they are biased—such a system might produce flight behaviour out of proportion to the number of occasions on which predators are present. But this is not biased belief, just 'biased' behaviour. The cognitive systems that have evolved from particular stimulus-response systems have in some cases retained the 'biases' of the latter, in particular if they are biases that might be beneficial to fitness for the reasons under discussion. Insofar as these cognitive mechanisms are belief-forming mechanisms they aim at truth; insofar as they are behaviour-influencing mechanisms, they may be biased. The point of the evolution of cognition, the production of belief, over a suite of stimulus-response systems is to decouple the information contained in a stimulus from behaviour. It is decoupled so that this information may be integrated with information from multiple sources, such as the various senses and memory and processes of theoretical reasoning—which is to say, it allows for practical reasoning. If the only influence on the production of beliefs were the benefits generated by this process, then a mechanism would not produce biased beliefs because its influence on behaviour depends on being integrated with other beliefs. The benefits of such a system arise precisely because one cognitive system does not 'know' what outputs another is producing. Specifically *belief*-forming mechanisms evolve to enable us to optimise behaviour in epistemically novel situations which are not well suited to stimulus-response mechanisms. These are precisely the conditions under which there is no benefit to be gained from bias.

²¹ Walsh's view is a dispositional theory, in that it defines function in terms of current contribution to fitness. Other accounts will emphasize past contribution to fitness. That difference, important though it is, does not bear on the current point.

4.4. COGNITION AND THE AIM OF BELIEF

Belief is subject to norms of correctness. These conditions are constitutive and individuating of belief. So a belief *qua* belief is what it ought to be if and only if it satisfies these norms. The normativist takes this to say all there is to say about the aim of belief. For the normativist the claim ‘A belief is correct if and only if it achieves the aim of belief’ is trivial. For the teleologist this is a more substantial claim because ‘the aim of belief’ is construed realistically—the norms arise from the teleology of belief. The teleology of belief, according to my argument, is found in the function of a cognitive system (rather than the intention of an individual believer). Functions are teleological and the source of the norms of belief.

A cognitive system has a function just as other systems of an organism have functions. The concept of function in use here is biological; it is nonetheless also normative. The heart’s function is to pump blood. If it fails to do so, it is malfunctioning, it is not doing what it ought to do. Metabolic pathways are biological systems for the synthesis of biologically important molecules. The porphyrin synthesis pathway has the function of producing haem, a key component of blood. If that pathway fails (for example because of an enzyme deficiency), then it will not produce haem, and a serious disease may result. The product of the porphyrin synthesis pathway *ought* to be haem, but sometimes it is not.²²

So to identify the correctness conditions of belief, we need first to identify the function of cognitive systems. As we have seen, that function is the production of true beliefs as the input into reasoning, since it is the truth of beliefs that makes them adaptive.

Since the function of a cognitive system is the production of true belief might, it might appear then that the appropriate norm or standard of correctness for belief is truth, and so we can conclude that belief aims at truth. This is not, however, quite correct, as I explain in section 5.2 below. For even if Φ -ing is the function of a system S, it does not follow that S is functioning *properly* when it does Φ .²³

5. *Belief aims at knowledge rather than truth*

In the preceding section, I argued that ‘the aim of belief’ is to be understood in terms of correctness conditions for belief, and that these are set by belief-producing (i.e. cog-

²² For example, a failure in the pathway may lead to sideroblastic anaemia. The pathway produces erythroblasts that are abnormal—megaloblasts—rather than normoblasts. Erythroblasts *ought* to be normoblasts but sometimes, in cases of sideroblastic anaemia, they are not.

²³ This functionalist approach will be attractive especially to naturalized epistemologists. Neta (2007) promotes an explicitly functional account of knowledge, while a similar view is articulated by Kornblith (2002). Bishop and Trout (2004) also develop a naturalistic approach to epistemology, but seek to bypass questions concerning the nature of knowledge. Neta appeals, as I do, to the function of cognitive systems, saying that knowledge is a matter of achieving the ‘goal-state’ of a cognitive system. He says that the goal-state of a cognitive system can differ between inquisitive and non-inquisitive creatures. It is not clear to me that this distinction is justified by satisfactory account of function. In any case, I don’t think that the goal state of a cognitive system includes reliability (non-inquisitive creatures) or being reasonably based on adequate grounds (inquisitive creatures). The goal-state is truth; knowledge is a matter of the goal-state being achieved by the proper functioning of a cognitive system.

nitive) systems having a function. This function is to produce true beliefs. It is important to note, however, that this does not mean that the aim of belief and the norm of correctness for belief is true belief. That is because a malfunctioning system can bring about the appropriate sort of output by accident—producing the functionally appropriate output, true belief in this case, is not enough for a system to be functioning correctly. So, I shall argue in this section, the correctness conditions for a belief (the norms of belief *qua* belief) are that it fulfils the function of cognitive systems is the right way. Those correctness conditions are fulfilled by knowledge but not by states less than knowledge.

I precede the main argument concerning the proper functioning of a cognitive system with considerations from the epistemology of Moorean and lottery beliefs.

5.1. KNOWLEDGE IS THE AIM OF BELIEF—EPISTEMOLOGICAL CONSIDERATIONS

Above, in section 2.3, I argued that our conception of ‘correct’ science is that it produces knowledge, and that we do not regard accidentally true science as successful science. Likewise, I first argue that our conception of correct believing is not satisfied merely by true belief, whereas it is satisfied by knowledge. This is revealed by certain epistemological phenomena surrounding Moorean beliefs and lottery beliefs.

Consider Elliott who does not know that dogs bark. He believes (Whiting 2013, 188):

(M) Dogs bark, but I don’t know that dogs bark.

Elliott’s belief (M) is true. But there is something wrong with this belief. He should not believe (M). If the norm for belief is truth, then since (M) is true, Elliott does satisfy the norm, and so there should be nothing to criticize about the belief (M).

On the other hand, if knowledge is the norm for belief, then (M) cannot meet the norm, for the norm would require him to know the first conjunct, that dogs bark. But he does not know that dogs bark. Furthermore, meeting the norm for the first conjunct (knowing that dogs bark) is inconsistent with meeting the norm for the second conjunct (knowing that he doesn’t know that dogs bark). So the belief norm does not explain why it is wrong to believe (M) whereas the knowledge norm does explain this.

Now consider Stanley who has one ticket in a million ticket fair lottery. The draw has taken place but Stanley has not been informed of the outcome. Stanley believes (Whiting 2013, 191):

(L) My ticket didn’t win.

which is in fact true. Nonetheless Stanley shouldn’t outright believe (L). Again the truth norm does not explain this, since believing (L) meets that norm. Whereas the knowledge norm does explain why Stanley should not outright believe (L) since that belief does not amount to knowledge.

If the incorrectness of believing (L) is not apparent, then consider a time shortly before the lottery draw. Stanley correctly thinks:

(P) My ticket might win the lottery.

He also believes:

(L*) My ticket will not win the lottery.

Both are true propositions. However, it would be wrong for Stanley to believe the true conjunction:

(P&L*) My ticket might win the lottery but it will not win the lottery.

If believing (P) and (L*) are both correct in this context, then believing their conjunction should also be correct. But it is not. Since it is clearly alright to believe (P), the problem with (P&L*) must be (L*). If it is wrong to believe (L*) ahead of the draw it is wrong to believe (L) after the draw since Stanley has received no new information in the meantime. Moreover, (P&L*) itself another case of a true proposition that it is wrong to believe, which the truth norm cannot explain but the knowledge norm can.

If Stanley believes (L) or (L*) he might act on that belief, for example by throwing his ticket away or not making the slightest effort to find out the result of the draw, even if the prize is significant. But those would be irrational actions. Imagine that the lottery is organised by a philanthropist who offers a €1,000,000 prize but sells the tickets for €0.50 each. So Stanley could reasonably buy a ticket on the ground that its expected value is greater than the cost of buying it. If it is reasonable to believe (L)/(L*) and so to act on it by throwing the ticket away, then he is guaranteed to lose money in a sequence of reasonable beliefs and actions.

In conclusion, we have two cases that are explained by knowledge being the norm for correct belief that cannot be explained by truth being the norm for correctness.²⁴

5.2. KNOWLEDGE AND PROPERLY FUNCTIONING COGNITIVE SYSTEMS

I have been arguing that belief aims at knowledge. One might argue that it is *prima facie* just as plausible that belief aims at truth, because, as I have acknowledged, the function of our cognitive systems is to produce true belief. True belief, it would seem, would have the same selective advantage as knowledge.

Timothy Williamson (2000) would deny the last assertion, on the ground that knowing has different behavioural consequences from true belief. For true belief is consistent with belief being formed unreliably. In such cases true belief is less robust than knowledge. While I am sympathetic to Williamson's position on this point, I pursue here a different line of argument.

If the function of a biological (e.g. cognitive) system *S* is *X*, there is more to the correct functioning of *S* than bringing about *X*. The normativity of function includes both what the function produces and how it produces it. Consequently, the corresponding norm is not merely that the system should produce *X*. For example, what is the function of the heart? To pump blood. That simple answer is good enough for most purposes. But it does not suffice to tell us when a heart is functioning properly. Consider Elsa the electri-

²⁴ But see Whiting (2013) for a defence of the truth norm against these arguments.

cian. While on a job, her heart stops beating in the normal way. She falls over and her chest touches some live cables she is working on and the electric shock she receives stimulates her heart muscles so that it pumps her blood. In this case her heart has brought about the supposedly defining activity of the heart's function, pumping blood. Yet, this is not a case of a properly functioning heart but is a case of a malfunctioning heart. A properly functioning heart is one whose pumping is brought about in a normal way. (It is up to physiologists to tell us what this is—the electrical stimulation comes from the sinus node, not from accidentally touching live cables.)

For this reason, the following:

The aim of belief is X iff the function of cognitive systems is X

is only approximately true. For an additional correctness norm (and so an additional component of the aim of belief) is that the cognitive system should produce X in the right way.

So while it may not be wrong to say 'belief aims at truth', like 'the function of the heart is to pump blood', such statements fail to capture the full normativity of 'aim' and 'function'. When epistemologists discuss the 'aim of belief', it is this normativity they are trying to capture and account for. They are asking for the conditions under which a belief can be said to be 'correct', which is to say that it satisfies all the norms it is subject to qua belief. The functionalist construal of 'aim' shows that a correctly functioning heart and a correctly functioning cognitive system are also subject to norms concerning the way in which pumping and true belief are brought about. As a result a correct pumping of the blood is one brought about by the heart's being stimulated by the sinus node, and a correct belief is a true belief brought about by a similarly proper, normal process.

What then is a true belief brought about by a similarly proper, normal process? An unjustified true belief and a justified true Gettier belief are beliefs not brought about in the normal way. In such cases our cognitive systems are malfunctioning despite producing true beliefs. According to Plantinga (1993) and Boyce and Plantinga (2012), *knowledge* is true belief brought about by a properly functioning cognitive system in an appropriate environment. For our purposes that seems right.²⁵ If unjustified beliefs and justified true but Gettier beliefs are examples of the products of malfunctioning cognitive systems or a well-functioning cognitive systems in an inappropriate environment, then it highly plausible that to be the product of a properly functioning system in an appropriate environment, the belief must amount to knowledge.²⁶

²⁵ One does not have to take proper functioning to provide an analysis of knowledge to hold that to meet the norms belief is subject to as the output of a cognitive system, a belief must amount to knowledge. The proper functioning view has been criticized on the ground that proper functioning is not necessary for knowledge, as shown by the Swampman case (Sosa 1993). The right answer to this, if you think that Swampman can in due course have knowledge, is to adopt a current fitness account of biological function (Walsh 1996). Although Swampman's 'biological systems' are not themselves evolved they can participate in evolution, and they do have current fitness and this determines their function. For more on proper functioning accounts of knowledge, warrant, and justification see Boyce (2018).

²⁶ The proper functioning account of knowledge has also been criticized for being insufficient for knowledge. Supposedly it is subject to its own Gettier-style cases (Boyce 2018). This depends on the details of one's version of proper functionalism. See Bergmann (2006) for another approach.

My view can therefore be summarized thus: it is the essence of a properly functioning cognitive system that it produces knowledge as its output. I note that this coheres with other views concerning knowledge. For example, John Hyman (1999) argues that *S* knows that *p* precisely when the fact that *p* can be used by *S* as a reason for something. For example, knowing that it is raining is a matter of being able to use the fact that it is raining as a reason for my taking an umbrella when I go outside, or as a reason for thinking that this is the third wettest June since 1900. According to my view, the function of the cognitive faculties is just that, to provide a link between the subject and the relevant facts so that they may be used as the inputs (reasons) in practical and theoretical reasoning. It also implies (and explains) the truth of Williamson's (2000, 47) claim that (mere) belief is botched knowing. Knowing is what a properly functioning cognitive system produces. False belief and true belief that isn't knowledge are what is produced by a cognitive system that is not functioning properly, when it botches its job. Finally, since the output of a properly functioning cognitive system is knowledge and such outputs are themselves the inputs into reasoning, it follows that the appropriate input into reasoning is also knowledge. Elsewhere I have defended the view that we can give a functional account of evidence—evidence is that which is the appropriate input into a knowledge-producing inference (and instance of theoretical reasoning). So the account I am giving here of the role of knowledge in cognitive system supports the view that evidence is knowledge (Williamson 1997; Bird 2018).

5.3. DEMANDINGNESS

Properly functioning cognitive systems in normal circumstances produce knowledge. The caveat 'in normal circumstances' acknowledges that to produce knowledge even a properly functioning cognitive system requires that the world should cooperate rather than present a well-executed *trompe l'oeil* or elaborate fraud, for example. Now consider conditions that are propitious in this respect, but in which knowledge regarding some proposition is not possible because relevant evidence is limited or contrary. A properly functioning cognitive system will not be disposed to produce belief in these conditions. A person with normal vision presented with a red tomato in good lighting conditions just cannot believe that it is green. A competent scientist who starts with an open mind on whether smoking causes cancer cannot just believe it does not solely on the basis of knowing about a healthy nonagenarian who has smoked heavily for seventy years.

The fact that a properly functioning cognitive system produces knowledge therefore explains the phenomenon that McHugh (2011) calls *demandingness*: the fact that 'you cannot, deliberately and in full awareness, form a belief in a proposition if you regard your evidence for that proposition as less than *sufficient*, where sufficiency involves more than having better or stronger evidence for the proposition than for its negation' (see also Owens 2003). McHugh is not claiming that beliefs with insufficient evidence are never possible. The exceptions, however, are pathological and clearly involve incorrect means of belief formation, such as self-deception. In cases where the preponderance of evidence is against a proposition, cannot the belief aim account for demandingness? After all, a belief in such circumstances is more likely to violate such a norm than satisfy it. McHugh (2011, 373, fn 9) defends demandingness even for such cases. In any case, I think that a stronger statement of demandingness holds, according to which one cannot form beliefs (deliberately

and in full awareness etc.) when the evidence favours a proposition but only very slightly. I subscribe to this stronger form.

Let us say you have the aim of achieving X. You can do action A with the intention of achieving X even if A has a relatively small chance of achieving X. But you cannot do A with the intention of achieving X if you know that A will not achieve X. For example, you want to go out of your house to the shops and back and you aim to do so without getting wet. You see that the clouds are dark and that rain is quite possible. On the other hand there are brighter clouds on the horizon and these seem to be coming in your direction. If the intention of avoiding getting wet is important you can wait until the skies are clear of the darkest rainclouds, reducing the chances of being rained on. If getting wet is not quite so important you can leave sooner. On the other hand, if it is actually raining now, so that it is clear that leaving your house will inevitably lead to you getting wet, then you cannot leave with the intention of not getting wet. Likewise, if the aim of belief were truth, then one could believe that *p* with that aim even if the evidence is only marginally in favour of its being true that *p*. On the other hand, if the aim of belief were knowledge, then one could not believe that *p* with that aim if the evidence is only marginal. For in such a situation it is clear that belief could not achieve that goal. Belief on the basis of limited evidence, even if the evidence is favourable and the belief is true, does not amount to knowledge. The truth aim of belief cannot explain demandingness whereas the knowledge aim does explain it.

6. Conclusion

What is the aim of belief? This is standardly construed as a question about the correctness conditions for belief—the set of norms concerning belief *qua* belief. These norms I have argued are supplied by the fact that the systems that produce beliefs, cognitive systems, have a function. What are the conditions under which a cognitive system is functioning correctly? When it produces *true belief in the right way*. Most plausibly true belief produced in the right way is knowledge. So the aim of belief is knowledge.

A functionalist view of social institutions says that institutions are like biological systems in having functions. Science is a social cognitive system and so the conditions of its proper functioning are the same as those for a biological cognitive system, i.e. knowledge. This conclusion is further supported by the consideration that progress in science is constituted by the accumulation of knowledge, implying that the aim of science is the production of knowledge.

REFERENCES

- Bacon, Francis 1620. *Novum Organum*. In B. Vickers (Ed.), *Francis Bacon the major works*. Oxford: Oxford University Press. Oxford World's Classics edition 2008.
- Baron, Jonathan. 2014. Heuristics and biases. In E. Zamir and D. Teichman (Eds.), *The Oxford handbook of behavioral economics and the law*, pp. 3-27. Oxford: Oxford University Press.
- Bergmann, Michael. 2006. *Justification without awareness: A defense of epistemic externalism*. Oxford: Oxford University Press.
- Bird, Alexander. 2007. Justified judging. *Philosophy and Phenomenological Research* 74: 81-110.

- Bird, Alexander. 2007a. What is scientific progress? *Noûs*, 41: 64-89.
- Bird, A. 2008. Scientific progress as accumulation of knowledge: a reply to Rowbottom. *Studies in History and Philosophy of Science*, 39: 279-81.
- Bird, Alexander. 2010a. The epistemology of science—a bird's-eye view. *Synthese* 175: 5-16.
- Bird, Alexander. 2010b. Social knowing. *Philosophical Perspectives* 24: 23-56.
- Bird, Alexander. 2018. Evidence and inference. *Philosophy and Phenomenological Research* 96(2): 299-317.
- Bishop, Michael and J. D. Trout 2004. *Epistemology and the Psychology of Human Judgment*. Oxford: Oxford University Press.
- Boyce, Kenneth. 2018. Proper functionalism. *Internet Encyclopedia of Philosophy*. <http://www.iep.utm.edu/prop-fun/> accessed 24 February 2018.
- Boyce, Kenneth and Alvin Plantinga. 2012. Proper functionalism. In A. Cullison (Ed.), *The Continuum Companion to Epistemology*, pp. 124-140. London: Continuum.
- Burgin, Mark. 2009. *Theory of Information: Fundamentality, diversity and unification*. Singapore: World Scientific.
- Bykvist, Krister and Anandi Hattiangadi 2013. Belief, truth, and blindspots. In T. Chan (Ed.), *The Aim of Belief*, pp. 100-22. Oxford: Oxford University Press.
- Cevolani, Gustavo and Luca Tambolo 2013. Progress as approximation to the truth: A defence of the verisimilitudinarian approach. *Erkenntnis* 78: 921-35.
- Dellsén, Finnur. 2016. Scientific progress: Knowledge versus understanding. *Studies in History and Philosophy of Science* 56: 72-83.
- Douglas, Heather. 2014. Pure science and the problem of progress. *Studies In History and Philosophy of Science Part A* 46: 55-63.
- Durkheim, Émile. 1893. *De la division du travail social*. Paris: Alcan. Translated as *The Division of Labor in Society*, by W. D. Halls, New York, NY: The Free Press, 1984.
- Durkheim, Émile. 1894. *Les règles de la méthode sociologique*. Paris: Alcan. Translated as *The Rules of Sociological Method*, by W. D. Halls, and edited with an introduction by S. Lukes, New York: The Free Press, 1982.
- Fassio, Davide. 2018. The aim of belief. *Internet Encyclopedia of Philosophy*. <http://www.iep.utm.edu/beli-aim/> accessed 25 February 2018.
- Franklin, James. 2009. *What science knows and how it knows it*. New York, NY: Encounter Books.
- Gibbard, Allan. 2005. Truth and correct belief. *Philosophical Issues* 15: 338-50.
- Gigerenzer, Gerd. 1991. How to make cognitive illusions disappear: Beyond “heuristics and biases”. *European Review of Social Psychology* 2: 83-115.
- Gilbert, Margaret. 1987. Modelling collective belief. *Synthese* 73: 185-204.
- Hoyningen-Huene, Paul. 2013. *Systematicity: The nature of science*. Oxford: Oxford University Press.
- Hyman, John. 1999. How knowledge works. *Philosophical Quarterly* 49: 433-51.
- Kiesewetter, Benjamin. 2017. *The normativity of rationality*. Oxford: Oxford University Press.
- Kornblith, Hilary. 2002. *Knowledge and its place in nature*. New York: Oxford University Press.
- Kuhn, Thomas. S. 1962. *The structure of scientific revolutions*. Chicago, IL: University of Chicago Press.
- McHugh, Conor. 2011. What do we aim at when we believe? *Dialectica* 65: 369-92.
- McHugh, Conor. 2014. Fitting belief. *Proceedings of the Aristotelian Society* 114: 167-87.
- Millikan, Ruth G. 1989. In defense of proper functions. *Philosophy of Science* 56: 288-302.
- Millikan, Ruth G. 1993. Naturalist reflections on knowledge. In R. Millikan (Ed.), *White Queen psychology and other essays for Alice*, pp. 241-64. Cambridge, MA: MIT Press.
- Neander, Karen. 1991. Function as selected effects: A conceptual analyst's defense. *Philosophy of Science* 58: 168-84.
- Neta, Ram. 2007. How to naturalize epistemology. In V. F. Hendricks and D. Pritchard (Eds.), *New Waves in Epistemology*, pp. 324-53. Basingstoke: Palgrave Macmillan.
- Niiniluoto, Ilkka. 2014. Scientific progress as increasing verisimilitude. *Studies in History and Philosophy of Science Part A* 46: 73-7.

- Owens, David J. 2003. Does belief have an aim? *Philosophical Studies* 115: 283-305.
- Papineau, David. 2013. There are no norms of belief. In *The aim of belief*, pp. 64-79. New York, NY: Oxford University Press.
- Parsons, Talcott. 1961. *Theories of society: Foundations of modern sociological theory*. New York, NY: Free Press.
- Plantinga, Alvin. 1993. *Warrant and proper function*. New York, NY: Oxford University Press.
- Rowbottom, Darrell P. 2008. N-rays and the semantic view of scientific progress. *Studies in History and Philosophy of Science* 39: 277-8.
- Rowbottom, Darrell P. 2010. What scientific progress is not: Against Bird's epistemic view. *International Studies in the Philosophy of Science* 24: 241-55.
- Shah, Nishi. 2003. How truth governs belief. *The Philosophical Review* 112: 447-82.
- Shah, Nishi and J.D. Velleman. 2005. Doxastic deliberation. *The Philosophical Review* 114: 497-534.
- Sosa, Ernest. 1993. Proper functionalism and virtue epistemology. *Noûs* 27: 51-65.
- Spencer, Herbert. 1874-96. *The principles of sociology*. Volumes I-III. London: Williams and Norgate.
- Spencer, Herbert. 1911. *Essays on education and kindred subjects*. London: Dent.
- Steglich-Petersen, Asbjørn. 2006. No norm needed: On the aim of belief. *The Philosophical Quarterly* 56: 499-516.
- Sullivan-Bisset, Ema. 2017. Biological function and epistemic normativity. *Philosophical Explorations* 20(sup 1): 94-110.
- Tuomela, Raymond. 1992. Group beliefs. *Synthese* 91: 285-318.
- Tversky, Amos et al. 1974. Judgment under uncertainty: Heuristics and biases. *Science* 185: 1124-31.
- Velleman, J. David. 2000. On the aim of belief. In *The possibility of practical reason*, pp. 244-81. Oxford: Oxford University Press.
- Walsh, D.M. 1996. Fitness and function. *British Journal for the Philosophy of Science* 47: 553-74.
- Wedgwood, Ralph. 2002. The aim of belief. *Philosophical Perspectives* 16: 267-97.
- Whiting, Daniel. 2013. Nothing but the truth. In T. Chan (Ed.), *The aim of belief*, pp. 184-203. Oxford: Oxford University Press.
- Williams, Bernard. 1973. Deciding to believe. In *problems of the self*, pp. 136-51. Cambridge: Cambridge University Press.
- Williamson, Timothy. 1997. Knowledge as evidence. *Mind* 106: 1-25.
- Williamson, Timothy. 2000. *Knowledge and its limits*. Oxford: Oxford University Press.
- Wright, Larry. 1973. Functions. *The Philosophical Review* 82: 139-68.

ALEXANDER BIRD is Peter Sowerby Professor of Philosophy and Medicine at King's College, London. He was previously the Professor of Philosophy at the University of Bristol and Reader at the University of Edinburgh. His recent publications include: "Against Creativity" (with Alison Hills) *Philosophy and Phenomenological Research*, "Understanding the Replication Crisis as a Base Rate Fallacy" *British Journal for the Philosophy of Science*, and "Fundamental Powers, Evolved Powers, and Mental Powers" *Proceedings of the Aristotelian Society Supplementary Volume*.

ADDRESS: Department of Philosophy, Philosophy Building, King's College London, Strand, London WC2R 2LS, United Kingdom. Email: alexander.bird@kcl.ac.uk.