



An externalist teleology

Gunnar Babcock¹ · Daniel W. McShea¹

Received: 7 January 2021 / Accepted: 24 April 2021

© The Author(s), under exclusive licence to Springer Nature B.V. 2021

Abstract

Teleology has a complicated history in the biological sciences. Some have argued that Darwin's theory has allowed biology to purge itself of teleological explanations. Others have been content to retain teleology and to treat it as metaphorical, or have sought to replace it with less problematic notions like teleonomy. And still others have tried to naturalize it in a way that distances it from the vitalism of the nineteenth century, focusing on the role that function plays in teleological explanation. No consensus has seemed possible in this debate. This paper takes a different approach. It argues that teleology is a perfectly acceptable scientific notion, but that the debate took an unfortunate misstep some 2300 years ago, one that has confused things ever since. The misstep comes in the beginning of Aristotle's *Physics* when a distinction is made between two types of teleological explanation. One type pertains to artifacts while the other pertains to entities in nature. For Aristotle, artifacts are guided by something external to themselves, human intentions, while natural entities are guided by an internal nature. We aim to show that there is, in fact, only one type of legitimate teleological explanation, what Aristotle would have considered a variant of an artifact model, where entities are guided by external fields. We begin with an analysis of the differences between the two types of explanation. We then examine some evidence in Aristotle's biological works suggesting that in his account of the natural-artifactual distinction, he encountered difficulties in trying to provide teleological accounts of spontaneous generation. And we show that it is possible to resolve these difficulties with a more robust version of an artifact model of teleology, in other words, with an externalist teleology. This is McShea's model, in which goal-directed entities are guided by a nested series of upper-level fields. To explain teleological behavior, this account invokes only external physical forces rather than mysterious internal natures. We then consider how field theory differs from other efforts to naturalize teleology in biology. And finally, we show how the account enables us to grapple with certain difficult cases—genes and intentions—where, even in biology today, the temptation to posit internal natures remains strong.

Keywords Teleology · Goal-directedness · Aristotle · Genes · Function · Intentions

Extended author information available on the last page of the article

1 Introduction

Modern biological sciences have made efforts to purge themselves of teleological explanations. This is because teleology is subject to the persistent criticism that such explanations depend on positing something like a “vital life force” in biological entities. In essence, if biological entities have goals then it seems there must be something special about life that isn’t present in the ontologies of the other sciences, something that imbues living entities with goals. However, positing something like a life force is a metaphysical commitment few are willing to make. If teleological explanations are only intelligible in conjunction with the existence of something akin to a life force, then biologists are right to reject teleological explanations.

This paper argues that criticisms leveled against teleology, like the one above, rely on an implicit background assumption about teleological explanations generally. This assumption is that there are really two kinds of teleological explanation: one that applies to the biological sciences, which is problematic, and another that applies to artifacts, which is unproblematic because goals are given externally via human intentions. This division between explanations was made by Aristotle in the *Physics*. It was made because it seemed obvious that external human goals direct the structure and development of artifacts, whereas natural goals, internal to biological entities, direct their growth and development. Accordingly, Aristotle defined natural entities as those entities that have their own internal sources of motion and rest. But, of course, these internal sources of motion and rest are metaphysically mysterious. And, it is not hard to see the trajectory in intellectual history that begins with Aristotle’s notion of internal sources of motion and ends with notions of vital life forces or benevolent creators. Since this division was introduced by Aristotle it has seldom been questioned by those investigating teleology.

We argue that this natural/artifactual division (NAD) is a misstep. The NAD is not necessary in order to deploy teleological explanations in biology. There is, in fact, only one legitimate kind of teleological explanation—the artifact model. This model does not require positing metaphysically objectionable life forces to explain the goal directedness of biological entities. It does, however, require a more robust understanding of the external forces that direct biological entities than we find in Aristotle’s scheme, where the only external forces are human intentions. After showing the problem with natural explanations, we also show that there is a variant of the artifact model—what we call “field theory”—that supports teleology in the sciences. Field theory offers a new kind of teleology, an *externalist teleology*.

To see how the natural model of teleological explanations goes awry, we look to Aristotle’s biological works, specifically his accounts of spontaneous generation. Aristotle appears to have encountered the limitations of the natural model when attempting to explain the biological processes underlying what he took to be spontaneous generation. In these problem cases, Aristotle saw that certain biological phenomena lacked obvious internal sources to explain their procreation. This forced him into the dilemma of having to either acknowledge that spontaneous generation was not a “natural” phenomenon, or to abandon natural teleological explanations. Having demonstrated that Aristotle may have come to question whether a hard and

fast division between nature and artifacts was necessary for teleology, we argue that field theory makes sense of the problem cases, and others more pressing in the biological sciences.

Under field theory, structures that are external to goal-directed entities, i.e., fields, produce what under an Aristotelian framework might be considered external sources of motion. This discovery reveals that many of the criticisms leveled against teleology only get traction by assuming the NAD, and it shows how eliminating the NAD rescues teleological explanations from many of the standard criticisms.

2 The natural/artifactual distinction

While Aristotle was not the originator of teleology, the account he provides of teleology is perhaps the most historically influential one. His main argument for teleological explanations is found in book two of the *Physics*. It has to do with what he takes to be explanatory failures of the alternatives. Theories of change, like those put forward by Democritus or Empedocles, do not have the explanatory resources to account for the phenomena we observe in the world. They either end up being too deterministic or too random to explain the regularity in much of the change we observe, even though this regularity is not entirely predictable either. Oak trees grow in regular ways, but we don't know precisely how many branches any given oak will have. Aristotle thought teleology best captured processes of change and motion. His argument for teleological explanation deserves greater attention than it is given here. But that is not the focus of this paper. Our interest in Aristotle's account of teleology has to do with a fundamental distinction that he makes at the very beginning of Book II of the *Physics*:

Some things are due to nature; for others there are other causes. Of the former sort are animals and their parts, plants, and simple bodies like earth, fire, air and water – for we say that these and things like them are due to nature. All these things plainly differ from things which are not constituted naturally: each has in itself a source of change and staying unchanged, whether in respect of place, or growth and decay, or alteration. A bed, on the other hand, or a coat, or anything else of that sort, considered as satisfying such a description, and in so far as it is the outcome of art, has no innate tendency to change... (*Physics*, trans. Charlton, 192b9-20)

Aristotle distinguishes between nature and art (or artifacts) because “these things plainly differ”. And there is evidence that Aristotle is not alone in this sentiment (see Rose et al. 2018). This difference in entities is important because it marks two different kinds of teleological explanations for Aristotle. Artifacts receive their ends, goal directedness or *telos*, from artisans. In this way, artifacts change, move or are directed by an artisan who is *external* to the artifact. A piece of metal takes the shape of a hammer because an artisan has a design in mind for something that will be able to secure nails. In contrast, Aristotle argues that natural entities have their own, internal sources of change, motion or ends. In this way, natural entities move

and change without needing to be pushed or pulled by anything from the outside. This internal nature contains a form that causes an entity to move towards a future end. Thus, the natural/artifactual division (or the “NAD”) offers two different models of teleological explanation. On the “natural” model an entity contains its own goals, or has ends built in, so to speak, to its internal nature, whereas on the “artifact” model, goals or ends are set by something external to the entity.

Most all the subsequent discussions surrounding teleology since Aristotle introduced the NAD have been influenced by it in one way or another, insofar as attention is given to the differences and similarities between organisms and artifacts. This remains true today (e.g. see Matthen, 1997; Lewens, 2004; Basl & Sandler, 2013). Discussions of the NAD often revolve around whether uses of the natural model should be understood as metaphors for the artifact model,¹ and to what extent there are literal parallels to be drawn between the two (e.g. Lewens, 2004; Wright, 1976). What continues to dog biology is its use of teleological language while simultaneously denying the existence of anything that smacks of a metaphysically dubious Aristotelian internal nature that might be construed non-metaphorically. If teleological language is taken literally, most see this as deeply problematic. For example, in *Not by design: retiring Darwin's watchmaker*, biologist John O. Reiss notes that “if we treat the intentional teleology involved in acts and the making of artifacts as the paradigm for teleology in general, the issue for teleological conceptions in biology is whether nature can have intentions” (Reiss, 2009, p. 10). Reiss’ conclusion about nature having intentions is a resounding “no”. But the tension remains present so long as teleological language is deployed in biology.

Notice that in the quote above Reiss, like Aristotle, understands the teleology of artifacts as being dependent upon the intentions of an external agent, like Wright (1976) and Boorse (1976). And Reiss does not take issue with this artifact model:

We commonly say, for example, ‘I went to the store to buy bread’ and consider this an explanation for our actions. Almost everyone agrees that from a scientific standpoint this is not problematic and involves no reverse causation, or cause that comes after its effect (a common concern of critics of teleology) (Reiss, 2009, p. 9).

What he finds problematic is the application of the artifact model (what he calls “intentional teleology”) to the biological sciences because he believes it unavoidably ascribes intentions to nature. He is not alone in this sentiment. Let’s briefly review the criticisms leveled against teleological explanations.

In the post-Darwin era, teleological explanations look problematic if they’re meant to work in the way Aristotle set out with the NAD. This is because while the artifact model appears scientifically unproblematic, the natural model requires positing internal sources of goal directedness that are metaphysically mysterious unless

¹ The use of ‘artifact model’ here is different than Lewens’ use of the term. Lewens’ notion of an “artifact model” is when biologists “talk of organisms as though they were designed objects” (Lewens, 2004, p. 2).

ascriptions of goal-directedness in nature are meant to be understood metaphorically. There is nothing puzzling about an artifact, like a hammer, because its shape, genesis and actions are all clearly derivative of one's intentions to secure nails. And, because we all routinely experience such intentional states there is no need to justify their causes or existence. The hammer has explanatory origins in a familiar, intuitive place. However, if the natural model is meant to mirror the artifact model in the way that Aristotle meant (Leunissen, 2010, pp. 16–8), then the origin of a natural entities' intentional source is obscure. There is no empirical evidence to suggest there are special vital life forces or internal natures, unless we're willing to locate these in something like genes (discussed later). Nor do we want explanations that imply backwards causation because of some future goal (see Garson, 2008, 2016). This has led to various modern efforts to naturalize teleology by describing teleological behavior without reference to such natures, but instead with reference to codes or programs that are internal to entities in the form of genes (see Boorse, 1976; Mayr, 1988) or internal mechanisms of various kinds (e.g., Nagel, 1979). In light of this, it is worth exploring what Aristotle had in mind when he says this source is "internal" to the entity, contained within its formal nature. What are these internal natures?

For Aristotle, the existence of internal natures doesn't require more than an assertion because "that there is such a thing as nature, it would be ridiculous to try to show; for it is plain that many things are of the sort just described" (*Physics*, 193a2–3). But why does he think internal natures are self-evident? When Aristotle asserts that natural things have internal goals, it seemed as though the burden of proof rested with any of his contemporaries who would have attempted to deny the existence of such goals. This is because if things, like oak trees, grow and reproduce (i.e. change) in a reliably predictable ways, even when presented with the same conditions as other natural entities, like maples, this seems to show that oaks are directed by an internal nature rather than external conditions. Moreover, because oaks grow in remarkably similar ways, this shows that their internal natures provide a much stronger directionality than the direction provided by even the best artisans. Given the powerful directionality of nature, how can its existence be denied?² But another part of Aristotle's motive for positing internal natures is because without them, he thinks knowledge about change is impossible to acquire. To deny the existence of internal natures is like putting oneself in the position of "a man blind from birth [who] would have to make inferences about colours. For such people discussion must be about the words only, and nothing is understood" (*Physics*, 193a8–10). Thus, Aristotle posits internal natures because they seem to be necessary for an understanding most of the phenomena in the natural world.

The assertion that internal natures exist is a prelude to the arguments that Aristotle goes on to offer for teleological explanation in book II of the *Physics*. Yet, there are other areas of the Aristotelian corpus that provide far better instances where

² To be clear, Aristotle does not suggest that internal natures imply anything akin to an intelligent designer. For example, Aristotle says that when spiders weave webs they are clearly goal directed. However, Aristotle says spiders have no intelligence. Therefore, there is no reason to think internal natures that direct entities need to possess, or have been created by, anything like intelligence.

Aristotle actively deploys teleology to explain various phenomena. Some of the best examples are found in his later biological works.³ There, Aristotle encounters a phenomenon that doesn't fit neatly into a teleology that depends on internal goal directedness—spontaneous generation.

3 The problem of spontaneous generation

At its root, the problem with the accounts Aristotle provides of “spontaneous generation” is that some animals appeared to lack the internal goal directedness that makes them develop in a certain way. Rather, it appeared as though certain external forces would come together in just the right ways to generate them. Spontaneous generation takes place when “animals and plants come into being in earth and in liquid because there is water in earth, and air in water, and in all air there is vital heat” (*On the Generation of Animals*, trans. Peck, 762a18–19). The animals that were believed to have been generated in this way were primarily testacea⁴ (*On the Generation of Animals* III.11). Understanding this reasoning requires a quick review of Aristotle's theory of sexual generation.

In sexual reproduction, an animal is created when a father contributes form through his semen and a mother provides material in her menstrual blood. The material must be of the right sort to be shaped by the form. Form is primarily responsible for directing the development of an organism in a way similar to the father who supplied it.⁵ And this is why Aristotle argues that in natural entities the formal cause and the final cause are one and the same. In this way, sexual reproduction is a paradigmatic example of the natural model of teleology (Lennox, 2001, p. 231). Moreover, it is a clean application of Aristotelian hylomorphism. However, in spontaneous generation, the roles that semen and menstrual blood play in sexual reproduction are replaced by a concoction of earth, water and air heated in just the right ways and under the right conditions. This means form is not an ingredient in spontaneous generation. And because there is no form that is passed from one generation to the next, the coming together of the right concoction of natural ingredients does not look like a teleological process at all. Rather, it looks like something that was not intended, i.e. something spontaneous. Thus, there appears to be tension in the account of teleology that we're given in the *Physics* where natural entities necessarily possess internal, formal natures, which move them towards predictable ends and the accounts we get in *On the Generation of Animals* where certain animals come

³ See Lennox (2001) p. 244 for interpretive reasons to think that the *Physics* predates the biological works.

⁴ *Testacea* are invertebrate animals with shell coverings, such as mollusks or shellfish. Given that Aristotle's writing predates (and is partly responsible for) any sort of systematic taxonomy, it's not entirely clear what species he thought were spontaneously generated, though many of the animals he describes are generally translated as 'testacea' in Aristotelian scholarship.

⁵ For a more detailed analysis of Aristotle's theory of sexual reproduction see Lennox (2001) pp. 230–1 and the Stanford Encyclopedia of Philosophy entry “Aristotle's Biology” (2019).

about spontaneously without any obvious teleological trajectory. What is a good Aristotelian supposed to do with this apparent inconsistency?

Hull (1967) argues that Aristotle's accounts of such spontaneous processes undermine his entire teleological project. This is because if Aristotle is able to provide an account for spontaneous generation without relying on a teleology that posits internal natures, then why are such accounts necessary to begin with? Why not just explain nature through spontaneous processes? And Lennox (2001) notes that Balme (1962) raised a similar concern about the apparent inconsistency.

Lennox (2001), however, offers a more nuanced view of the matter. He argues that a close examination of Aristotle's theory of spontaneity shows that it comes close to other theories of causation that Aristotle critiques. According to Lennox, Aristotle does not take spontaneous generation to be a teleological process. The processes that result in spontaneous generation are ones that normally take place teleologically, so if there are occasional cases where organisms are generated spontaneously, that's okay for Aristotle's broader theory. However, if animals are *regularly* generated through spontaneous means, it begins to look more and more like evidence against Aristotle's teleology. Consider the following example to flesh out Lennox's point. Say someone goes to the store to buy milk. In this case, their behavior is explained teleologically in light of their goal to get milk. But let's say they bump into a friend who owes them money while they're out. This incidental occurrence results in the friend paying the debt. One might say that Aristotle's view of spontaneous generation is somewhat like this. Just because an incidental process brought about their being paid, this does not mean they didn't have the goal of getting paid. Sometimes incidental events achieve ends. In this way, spontaneous generation does not have a "form" that is passed along by the male in sexual generation that makes the processes goal directed towards that form. Thus, sometimes goals are reached incidentally. However, for a species to maintain itself spontaneously the sheer frequency at which spontaneous generation takes place would have to be quite high. The higher it gets, the less the process looks to be merely incidental. And this looks less like an occasional instance of spontaneous generation and more like a teleological process without a formal end or internal nature. Such an outcome would pose a problem for Aristotle's broader theory.

Here we might look to Zwier (2018), who views Aristotle's accounts of spontaneous generation as providing an illuminating insight into his methodology. She argues that because Aristotle claimed that there can be no science (i.e. *techne*) of truly spontaneous occurrences, his treatments of spontaneous generation are best understood as examining whether such processes are, in fact, "spontaneous" as had been reported. He was looking for patterns in what he took to be empirical facts to see if a systematic theory, like teleology, could explain them. At the same time, he was open to the conclusion that some natural entities might not be teleological, even if it meant adjusting the applicability of teleology. For Zwier, one might say this is an example of Aristotle the scientist rather than Aristotle the philosopher. In light of these conflicting interpretations, it seems safe to say that the jury is still out on how we should understand this particular Aristotelian puzzle.

Spontaneous generation presents clear cases where Aristotle's theoretical framework for teleology encountered inconsistencies with what he took to be the empirical

facts. We have no interpretive insight to offer on Aristotle's accounts of spontaneous generation. Instead, we draw attention to spontaneous generation because it shows that Aristotle grappled with the problem identified at the end of the last section—there were cases in nature where deferring to internal sources to explain motion and change becomes messy.

Now let's return to the NAD. Notice that the analyses of generation above, both sexual and spontaneous, have assumed the *natural model*. Without the existence of internal goal directedness, the natural model doesn't get going. In sexual reproduction, internal goal directedness is easily identified in the form supplied in the semen. There is no obvious internal nature to be found in spontaneous generation. Hence the impasse. So, why didn't Aristotle just invoke the artifact model to explain spontaneous generation? Well, it's supposed to apply to only artifacts, not natural entities. Fair enough. But then we might wonder why he didn't revise the framework of the NAD? There doesn't appear to be an obvious answer to this question, and whatever answer there is, is best left for Aristotle scholars. However, we can now see that the modern view of teleology has been skewed from having been entrenched in more than two millennia of the NAD. It has skewed inward instead of outward when seeking explanations in nature.

In the next section we'll argue that the NAD is a mistake. In fact, we claim the only legitimate teleological explanation comes in a new version of the artifact model. In what follows, we argue that it's a mistake to look for internal sources of direction, full stop. And in doing this we're doing something that others who have sought to do away with the division between nature and artifacts haven't done. We now turn to dispelling the misconception that the natural model, i.e. that internal sources of guidance, have a place in teleology, and in doing so we present what we call an externalist teleology.

4 Field theory: a new artifact model

In the next sections, we argue for two claims in turn. The first is that there are countless cases where natural entities are directed by external sources in a way that mirrors the artifact model. The second claim, and perhaps the more controversial of the two, is that the natural model of teleology is so deeply flawed that it has no applicability. Guidance always comes from the outside, from some larger external entity in which the guided entity is embedded. This means that the only version of teleology that works is the artifact model. And the artifact model applies to far more than just human intentions and artifacts. It applies to all teleological systems in biology. Recognizing this about teleological explanation generally means that teleology, when correctly understood, is not subject to the most serious criticisms raised against it. And it means that outside guidance is neither metaphysically problematic nor a metaphor for intention.

We'll begin by examining how natural entities are directed by external sources. This requires dispelling the idea that the natural model necessarily applies to *all* natural objects and that the artifact model necessarily applies only to artifacts. For us it amounts to claiming that Aristotle was mistaken about the connection between

internality and externality on the one hand, and the causes of change and motion, or goal directedness, in natural and artificial objects on the other.

From a contemporary vantage point it is not difficult to find counterexamples to this Aristotelian premise. Consider the heliotropism exhibited by juvenile sunflowers. These plants move throughout the day in a way that tracks solar light. The process is apparently teleological in that the plants seem to pursue the goal of remaining aimed at the sun. And it is a good example of the artifact model in that the source of direction for the process, the sun's rays, is external to the plants. In Aristotle's terms, sunflowers are "natural" as they were not created by an artisan of any kind. But the source of their guidance is not an internal nature of any kind either. Now internal factors are clearly important. The various mechanisms that detect the sun and produce the differential growth that tilts the plant toward it are internal. At a lower level, the process is regulated by certain genes, which some would construe as an internal "program," recalling Mayr's notion of teleonomy as direction by an internal program. But while the genes and the internal mechanisms they underlie are necessary, they do not supply the plant with the direction that is so critical to its teleological behavior. They do not tell it which way to turn. Indeed, while there is evidence that sunflowers have a circadian rhythm (Atamian et al., 2016), and therefore presumably have some ability to track time throughout the day, without the sun they cease their east to west movement, as has been demonstrated when they are contained in greenhouses with a fixed light source (Lang & Begg, 1979, also see Vandenbrink et al. 2014). For sunflower heliotropism, the external model—that is, the artifact model—works better than Aristotle's natural model, in spite of the fact that sunflowers are not artifacts.

Of course, heliotropism in sunflowers is an example that's conducive to demonstrating the externality of guidance. There are other examples in nature that appear to support a more Aristotelian view, where guidance appears to stem from internal sources. We'll address some of these more challenging cases in section six. For now, the point is that the Aristotelian model is clearly not up to the task of accounting for all the teleological behavior we find in nature. This means, at the very least, some other version of teleology is needed to account for these cases.

For explaining teleology in biology, the external model turns out to be far more general than it may first appear. In organismal tropisms and taxes, movement of the organism is guided by an external factor of some kind, or by what one of us has called an external *field* (McShea, 2012). The notion of a field has a rich twentieth century history, both in systems theory (von Bertalanffy, 1969) and in developmental biology (Weiss, 1969), and is current in that latter field today. In Levin's words:

'Field' denotes both informational and regional relationships. The quintessential property of a field model is non-locality—the idea that the influences coming to bear on any point in the system are not localized to that point and that an understanding of those forces must include information existing at other, distant regions in the system. (Levin, 2012, p. 244)

In the case of the sunflower, it is the sun, or the light field from the sun, that supplies the "information from distant regions." Fields in this sense have been sometimes treated as metaphorical, as when the behavior of some pattern in development is

well described by a set of field equations from physics, but when there is no reason to believe those physical mechanisms are in play (e.g. Thompson, 1942). In other cases, and especially recently, field equations have been developed based on biochemical and biophysical processes known to be actually in play, and in these, fields are understood to be real (e.g., Goodwin, 1994, see Levin, 2012 for references).

For a bacterium seeking food, the field is a chemical gradient created by the food source. For a moth circling a lightbulb, the field is the light field emanating from the bulb. For certain sea turtles navigating to the beach where they hatched, it is the Earth's magnetic field. The same is true for developmental processes. Development is perhaps the archetypical teleological process in biology, a process in which cells and tissues move and differentiate seemingly in pursuit of a predestined, future end-state. As we will explain later, these changes are directed at a large scale by what are called morphogenetic fields.

The conditions for entity or entities constituting a field are fairly permissive. Any physical structure that is larger than, and encompasses smaller entities lying within it, can be a field. "Field" basically describes the larger, outer entity in any physically nested system. Fields can be continuous, as in an electric field or a chemical gradient. But they can also be discontinuous. The plastic wall of a balloon is a field with respect to the gas molecules within it, leaving those molecules free to wander, except right at the wall where it blocks their movement. Fields can be what are sometimes called "boundary conditions." Also fields can be homogeneous, as in the case of the balloon wall, consisting of the same substance achieving containment in the same way everywhere, over its entire inner surface. But they can also be inhomogeneous, as in the case of what we will later call ecological fields. The food sources, climate conditions, predators, competitors, and parasites that govern an organism's reproductive success constitute a field. The system of laws, conventions, and expectations that govern paying for a pack of gum at a gas station is an inhomogeneous field, what we will call a social field. Social fields can consist of a mix of inanimate objects, environmental conditions, and people, complete with their various mental states. Notice that every aspect of the fields described in this paragraph—including ecological and social fields—is physical. Fields are real, not metaphorical. And again the conditions for being a field are not stringent. Fields are everywhere. But where fields become interesting and relevant to teleology is in the cases where they have the capacity to direct the smaller entities nested within them.

Building on this notion of a field, McShea (2012, 2016) has offered a general account of teleology, or goal directedness, which here we will call *field theory*. The theory borrows from Sommerhoff (1950) and Nagel (1979) the notions of persistence and plasticity. Persistence is the tendency for an entity that is following a particular pattern of behavior, a trajectory, to return to that trajectory following perturbations, while plasticity is the tendency for an entity to find a particular trajectory from a variety of different starting points. A homing torpedo headed for a target ship might be dragged off course by a locally strong ocean current, but so long as it remains within the "sound field" emanating from the target ship, it persists, returning eventually to a path toward the ship. The torpedo's behavior is also plastic in that it can find a trajectory toward the ship from any release point within the sound field. For McShea, persistence and plasticity do not define goal directedness. Rather they

are hallmarks, key features, of goal-directed systems. The mission of a theory of goal directedness is to explain how they are achieved. In this case, the question has already been answered: it is the sound field that make persistence and plasticity possible. The internal mechanics and electronics of the torpedo are critical, of course, but it is the external sound field that provides the guidance. In the language McShea adopts for goal-directed systems generally, the torpedo is the goal-directed entity, and it is “upper directed.” That is, it is directed by—its persistence and plasticity are guided by—an upper-level system, the larger sound field in which it is immersed.

There are three important points to make about upper direction and fields. The first is that the physical arrangement is hierarchical. Entities are nested within fields. The moth is physically contained within the light field emanating from the lightbulb. And there may in some instances be multiple levels, telescoping in an ever-increasing upward hierarchy. A muscle cell in a developing embryo is nested within a muscle tissue and directed by fields generated by that tissue. The muscle tissue in turn is nested within and directed by larger fields generated by, say, the developing leg of the embryo in which the tissue is embedded. And the embryo as a whole is nested within and directed by the maternal environment around it. Second, there is no suggestion that all or even most causal interactions count as upper direction. The alternative, lateral direction, occurs whenever entities of about the same size and causal scale interact with each other, as when one billiard ball collides with and thereby directs the movement of another. Lateral direction is ubiquitous, both in biology and physics. The point is that lateral direction does not produce the persistence and plasticity that are the hallmarks of goal-directed entities. Third, there is nothing spooky about fields. They are real and fully describable in physical terms. The existence of a magnetic field guiding a sea turtle, or a large morphogenetic gradient guiding a cell in an embryo, does not introduce any objectionable entities into our ontology.

A remark is necessary here about the philosophical mission of field theory. It is not to define teleology in a way that an analytical philosopher might want, to offer necessary and sufficient conditions for behavior to be called teleological. (The search for necessary and sufficient conditions was a preoccupation of a small corner of analytic philosophy in the twentieth century and there is some consensus that it failed [Garson, 2016; Lee & McShea, 2020].) Field theory does not claim that all and only teleological systems must be structured hierarchically, as goal-directed entities moving within upper-level fields. McShea (2012) has argued that this could well be the case, but nothing in field theory demands it. Rather, field theory takes an engineering approach in the sense described in Wimsatt (2007). Field theory merely observes what appears to be a regularity about goal-directed entities, that from sea turtles homing to a nesting site to torpedoes tracking target ships, the evidence strongly supports the premise that they are guided by upper-level fields, by structures that are larger than and external to them.

With an outline of field theory in view, it is now possible to see its kinship with Aristotle’s artifact model. It explains the behavior of teleological entities by reference to external fields, rather than internal natures. Of course, field theory offers a very different conception of externality than what Aristotle had in mind when he outlined the artifact model. The external source of motion in Aristotle’s examples are primarily human artisans. But notice that whether the external directing cause

is an artisan or a light field, they all easily fit within idea of externality inherent in the artifact model. To underscore the principle of the NAD once again, for Aristotle what makes something natural or artificial is the location of its source of motion, as opposed to whether the entity that is being directed happens to be what we'd typically classify as natural or artificial. Field theory departs from Aristotle's conception of the artifact model insofar as it recognizes a much larger pool of entities—namely fields—that *externally* direct.

Also notice how easy it is to address the puzzle of spontaneous generation using field theory, which is unencumbered by the NAD. Recall that Aristotle believed spontaneous generation occurred when the right concoction of earth, water and air came together and were heated in the right way to generate certain animals. If such a process were to take place, field theory would merely note that the combination of environmental conditions that directed the critical materials is responsible for spontaneous generation. Without the NAD, without a commitment to internal guiding factors for natural entities, the solution is straightforward.

5 How field theory is different

Having provided an account of field theory, it is now important to situate it within the more recent literature on function and goal directedness. Most of the accounts of function fall under the umbrella of what Allen and Neal (2020) call “teleonaturalism”. Teleonaturalistic accounts aim to legitimize teleological explanations by offering accounts of function that do not have to posit intentional states as they “seek truth conditions for teleological claims in biology that are grounded in non-mental facts about organisms and their traits” (Allen & Neal, 2020). Following Allen and Neal's notion of teleonaturalism, field theory could be deemed teleonaturalistic. However, we want to underscore several important differences between some of the teleonaturalistic accounts and the account offered by field theory.

First, there is a difference between function and goal directedness (see Wright, 1973; Garson, 2016, p. 18). Field theory is concerned with goal directedness as opposed to function. To see this difference, consider a hammer. At first pass, it would appear as though a hammer's function is to secure nails. However, a hammer only has this function if it is part of a larger, hierarchical goal-directed system. The hammer's function is really a teleological system that consists of a hammer in combination with an agent. The agent's intentions provide a field at the upper level within which the hammer is contained. Once in this system, a hammer is then functional as a nail pounding tool. Absent the upper-level field of the agent's intentions, the hammer is function-less. Thus, an account of a hammer's function (i.e. an account of function) is not equivalent to account of the teleological system that gives the hammer function (an account of teleological systems). Field theory is an account of teleological systems, not an account of function like those offered by the selected-effects account or the causal-roles account (see Cummins, 1975; Neander, 1991; Craver, 2001; Garson, 2016), which Allen and Neal group within teleonaturalism. So while field theory might be teleonaturalistic, it is an account of teleological systems, not an account of function.

Our approach is also doing something different from others who have suggested that a strong divide between nature and artifacts is not merited, like Cummins (1975), Wright (1976) and Boorse (1976). For example, Wright's etiological analysis of function (see Wright, 1976) suggests, as we do, that the NAD does not need to be made. However, because Wright argues that teleological explanations are largely metaphorical, his account is quite different from field theory. We do not see teleology as metaphor of any kind. And while Boorse (1976) criticizes Wright's etiological account, he finds biological directedness coming from both external *and* internal sources, like genes.⁶ Again, this sets Boorse apart from what is presented in field theory. Finally, Cummins (1975) similarly does not see the difference between nature and artifacts as important, but because he is concerned with functions rather than goal-directedness, his account is largely adjacent to the central issue of this paper. But, most importantly, while these accounts do not see major differences between artifacts and nature, neither do they focus on whether direction comes from either the inside or the outside. In bringing our inquiry back to Aristotle, we're noting that the original division between nature and artifacts really has to do with the source of directedness, and that is the focus of our interest as well.

Field theory also differs from other attempts to "naturalize" teleology in biology as it is not meant as a mere heuristic. Other accounts argue that many of the mechanistic realities of the biological sciences are not descriptively available to us, and therefore teleological language is a shortcut that enables us to bypass mechanistic puzzles. For example, Mayr (1988, 1992) uses teleological language in this way to describe goal-directed behavior in organisms, arguing that such behavior is guided by "internal programs." And it is telling that he describes the resulting behavior as "teleonomic," in order to distinguish it from true teleology which he takes to be problematic because of its link with Aristotle's final causes. In contrast, while we clearly take Aristotelian internal, final causes to be deeply problematic, field theory understands teleology and the fields that explain it to be very real. Fields are empirically detectable and often directly measurable. There is nothing about the directing force of fields that requires spooky metaphysics. As was argued in section four, fields are multiply realizable and diverse in their composition, and the requirements for being a field are quite permissive.

With these distinguishing features of field theory now in view, we'd like to turn to the cases that appear to be the most difficult for an externalist teleology. In the next section, we address the cases where today's science has seemingly capitulated to the temptation to invoke internal natures—genes and intentions. We hope to show that field theory is actually needed to account for the goal-directed phenomena they are said to guide.

⁶ Interestingly, Boorse (1976) does note that things like ecology can direct certain biological entities. Something like ecology is external in our view. Insofar as Boorse gestures at these types of external sources of direction, we're sympathetic to his view.

6 The hard cases

We are now ready to argue for the more contentious claim that teleological entities are never directed internally. In other words, we will argue that all teleological behavior can be explained with reference only to external upper-level fields, that teleology never requires positing metaphysically suspect internal natures like those described by Aristotle, nor does it require ascribing intentions to nature, of the sort that Reiss and others (rightly) find objectionable. In doing so, we hope to show how teleology can be made “safe,” so to speak, for modern biology. But to be convincing on this point means having to overcome the highly intuitive cases where there is a strong *appearance* of an internal nature. Showing this by pointing at heliotropism in sunflowers was easy. Let’s turn to the hard cases.

6.1 Genes

6.1.1 Genes in development

The obvious place in biology to look for an internal source of direction is the genes. They are not just inside us, they are buried down deep inside individual cells. And therefore since genes are generally understood to be the entities that guide organismic development and physiology, they look like good candidates for internal natures. Mayr’s (1988) teleomatic account of genes, where genes represent something like an internal program that directs the development of the organism, is a view along these lines. Boorse (1976) also points to genes as supplying something like guiding force.

In this subsection we will show why viewing genes in this way is a mistake. We must recognize right at the start that making this case will be difficult, so deeply ingrained is the view that genes are mechanically central in development and guide it. But our view becomes instantly more plausible if we are allowed to separate those two kinds of cause, mechanism and guidance. The parts and electronics of a homing torpedo are central mechanical causes in homing, but they are not the source of any information about the location of the target ship, and in that sense, they do not guide it. The many neurons in the brain of a moth are critical to the mechanics of its light-seeking behavior, but those neurons by themselves have no information to offer about the location of the light bulb, and in that sense, they give no guidance. As will be seen, our view fully understands the importance of both internal and external causes, but it also recognizes and takes seriously an asymmetry. The guidance, the direction, that is the signature of goal-directed systems, must have a source with access to information about the goal. As will be seen, in developing embryos, the guidance does not come from an external point source as it does for the torpedo and moth, but from large-scale biophysical structures, fields that are internal to the embryo but enveloping its cells and the genes—enveloping and therefore external to them.

Consider the development of an embryo of a large, multicellular organism. For certain cell and tissue movements and changes, the source of the required guidance

is unambiguously external, as when an embryo's transformations are guided by uterine factors, gene products secreted by the mother. But even later in development, where the embryo appears to guide itself, changes are actually guided by morphogenetic fields, by chemical gradients arising from the simultaneous activation of genes in certain regions of the embryo. In sea urchin eggs, the cells that give rise to the skeleton seem to be guided to their proper destinations by a morphogenetic field that spans much of the embryo (Ettensohn, 1990). In some cases, electrical fields have been implicated. [See review and citations in Levin (2012)]. These fields are internal to the embryo, to be sure, but they are external to the suites of cells and tissues they guide, larger than and thoroughly enveloping them, and therefore external to the genes. In most cases, the precise nature of the field is unknown, but field origins and structure have long been a central focus of developmental biology (Weiss, 1969).

It is these fields that guide the embryo. What of the genes? They are the ultimate source of the fields, manufacturing the chemical substances that self organize into fields. But manufacturing substances is not guidance. Nor could it be. The genes, considered as a set of molecules, lying buried within the cytoplasm of single cells, simply do not have the informational resources to guide large numbers of cells, to run the show in the development of a large organism. The target morphology of the organism is not present, or represented, anywhere in those molecules. They contain no blueprint, no map of the whole (Nijhout 1990; Fox Keller 2010). Genes are switches, of a sort (Goodwin, 1994; Kauffman, 1993), turning on and off—or continuously regulating—the production of proteins. And no set of lower-level switches—however many copies of that set exist in other cells—can “know” the morphology of the whole, or any large part of it, and cannot guide the transformation of that whole toward a particular outcome, toward a goal. More concretely, the genome of a cell that is part of a developing vertebral column contains no blueprint, no map, of the future column. If such a “map” exists it can only be present in something large, on a scale commensurate with the column, in other words, one or more vertebral-column-sized morphogenetic gradients, that is, “fields” of some kind. In other words, guidance at the scale of the whole must come from something large, something external to the developing column, if it is to have the information necessary to specify what goes where, to recognize errors at the large scale when they occur, and to correct them. The internal process of gene “switching” is a powerful mechanism, causally central to development, but it does not *guide* development.

Of course, none of this entails that it is metaphysically impossible for genes—or any tiny internal entities—to guide development. Our point is that they do not, in fact, play that role. Consider for a moment what would be necessary if genes *were* directly responsible for guiding development without giving any consideration to larger-scale structures, that is, to fields. The genome of each cell would need to be pre-programmed with instructions for what cell type to transition to and when to make that transition, as well as where to move to, what connections to make with other cells, and when to make them. And the programming would have to operate with micrometer-scale and microsecond-scale precision. From an engineering perspective such a system would be extremely fragile and error intolerant, as a single error would result in a cascade of misaligned developmental progress that would doom an embryo. It would be a design disaster. Imagine a parallel case: individuals

in a national economy. In order for a national economy to function properly, and to do so in the absence of any larger-scale forces—that is, market forces such as those imposed by pricing, interest rates, information on job and resource availability, and so on—every individual would have to be programmed with a complete set of instructions for where to live, what job to do, what to buy, and when to buy it. Or imagine a system of self-driving cars in which each car had to come equipped with a full set of instructions for how to get from A to B, with no GPS and no roadmap, but instead with every turn of the steering wheel pre-programmed. Worse, pre-programmed for all possible trips, for all A and all B. While such a picture is not metaphysically impossible, it's poor architecture from an engineering standpoint, and it isn't the way biological systems work.

As we said earlier, our intent is not to downplay the mechanical centrality of the genes. They generate essential signals and substances, often at precise moments, at specific rates, and they are so deeply foundational that small defects in their function can derail the entire process. And obviously, they are difference makers (see Waters 2007). It is mainly the genes that decide whether a given embryo will develop into a monkey or a bat. But for development, for the generation of form at the scale of, say heads and legs, the source of information can only be something on the scale of the heads and legs. The genes are too small, not to mention poorly positioned (inside the developing structures), to offer any guidance. The same goes for behavior and physiology. The same goes for smaller, less complex organisms. Generally speaking, external fields guide. Internal entities provide mechanism.

6.1.2 Genes in evolution

There is a second argument that undermines the genes' role in directing development. But before giving it, we need to say something about the multiple levels of nestedness, and the multiple levels of teleology, that are present in certain teleological systems. As a result of this nestedness, there is no contradiction in saying that an entity is guided simultaneously at multiple levels. Thus, an individual could be guided from above by, say, the forces exerted by her social circle—a friends and family “field”—and that the same individual could also be guided from a much higher level by the larger-scale forces of the national-economy field in which her social circle is immersed. Both are external to her and direct her behavior to some extent, albeit at very different scales.

The situation for genes is similar. Genes do not participate only in development, behavior, and physiology but also in evolution, providing the necessary ancestor-descendant continuity. The genome of a given organism is part of an ancestor-descendant series of populations of individuals, a genetic lineage, or more precisely a species individual (see Ghiselin, 1974, Hull, 1976). Species, though spatially discontinuous, exist at a scale that is vastly larger than individual organisms, encompassing all of the individuals in the species over the (potentially) multimillion-year lifetime of that species. And therefore, at the same time as the activity of the genes of an organism is guided by upper-level fields in development, like morphogenetic fields, both genes and fields are guided on a much longer timescale by natural selection, which is to say by the ecological setting in which the species

finds itself. The species ecology is the larger-scale system in which the species is immersed, here construed as an upper-level field. This ecological field is obviously larger than, and external to the species, and is therefore (hugely) larger than and external to the genes. Any attempt to isolate genes apart from genetic lineages of which they are parts treats them as independent of the lineage, which as one of us has noted, they are not (see Babcock, 2020). Thus, genes lie at the bottom of a telescoping hierarchical series—genes within organisms within species individuals within ecological settings—and external upper-level fields guide them simultaneously at two levels, in development and in evolution.

It is worth pointing out that, for simplicity of explication, our argument in this section actually skipped a number of hierarchical levels. In particular, in developing multicellular organisms, there is at least one level intermediate between the genes and the organism as a whole, the level of the cell. And to some extent, what morphogenetic fields guide is not the genes, but the cells within them, many of them little teleological entities in their own right, with some capacity for homeostasis, self organization, and regeneration. Looking down from the organism level, morphogenetic fields direct the cell, and the cell's role is that of mechanism. Looking further down, the cell generates smaller upper-level fields that direct its parts, including its genes, and the role of these parts and genes is that of mechanism. The point is that the hierarchical decomposition of the system we gave above was incomplete. And a full description would include a fair number of additional levels between ecology and the genes. Also, we will add, for many systems, especially complex biological systems, boundaries between levels will be indistinct, field relationships will be unclear, and hierarchical decomposition will not be straightforward (Wimsatt, 1994). Complete treatments of teleology in such systems await further development of theory and methods in the study of hierarchy and fields.

Before leaving the issue of levels, we want to mention an example that involves more than one level and also highlights the importance of correctly identifying the teleological entity to be explained. Imagine an organism that is positively phototropic in one phase of its development and negatively phototropic in another. And suppose that the switch from positive to negative is accomplished by a genetic switch of some kind. In observing the organism, we see that it switches from moving toward the light to moving away from it, and knowing that the underlying cause is a change of state of a gene, we might be tempted to say the organism's teleological movements are directed by the gene and therefore the guidance is internal. What this analysis fails to recognize is that there are likely two teleological systems here. There is the movement of the organism toward or away from light, which is clearly directed externally, by the light field. And then, independent of that, there is the change in the state of the gene. Suppose that this latter change is also goal directed, that it is persistent and plastic, meaning that the change in state is monitored and controlled so that if the switch fails to occur as scheduled, the system tries again to make it occur, and if the switch is successful, the system stabilizes it in its new state. If so, then our argument above is that this gene-switch is likely under the control of something larger and external to it, a morphogenetic field of some kind. In any case, the critical point here is that there are not one but two teleological systems, the organism as a whole whose phototropism—whether positive or negative—is directed by

something external to it (a light source), and the gene switch, whose position is also directed something external to it (a morphogenetic field).

Finally, it might seem that the foregoing discussion vindicates Aristotle. Our discussion has placed the source of developmental goal directedness in morphogenetic fields, which while external to genes are nevertheless still internal to the embryo. If genes are poor candidates for internal natures then perhaps morphogenetic fields are good ones? Not so fast. Like genes, morphogenetic fields are guided externally, by natural selection. Thus, while it is true that the proximate guidance for embryos comes from structures that are internal to the embryo, morphogenetic fields, ultimate guidance—the guidance that creates and shapes the morphogenetic fields and gives them their causal powers—comes from outside, from ecology. As the great 20th century evolutionary biologist, Leigh Van Valen, wrote, “evolution is the control of development by ecology” (Van Valen, 1973, p. 488). In our terms, ecology is the source of teleology in development.

6.2 Intentions

Human intentions are what Reiss (2009) identifies as the source of many of our actions and he sees this as obviously true. Most intuitions about behaviors having an internal origin likely stem from intentions. An intention to buy bread clearly explains someone’s behavior to go to a store, pick up a loaf of bread and bring it to the checkout. According to Reiss, this explanation is teleological and requires no further justification because intentional states are things we all experience. Likewise for Aristotle. While Aristotle seldom explicitly makes connections between artisans, their intentions and the artifacts they build, his theories in other areas make it fairly easy to determine his likely account (see Leunissen, 2010). These are found in his detailed account of the psychology of plants, animals and humans in *De Anima*, and from them we could infer that in Aristotle’s view a builder builds a house because of certain perceptions she has in her soul that have to do with the forms of houses. Notice the causal chain of artifacts leads back to the internal soul of the artisan, and we arrive at an internal nature once again.

In this subsection, our task is to demonstrate that this seemingly simple explanation is actually far more difficult to understand than it first appears. We will propose that the behavior exhibited by teleological entities—behaviors like human bread-buying or house building—are driven by external fields, that intentions *are* external fields. More precisely, intentions are external to the conscious processes they guide, and further they are guided by even larger external fields. In this way, we block critiques of teleology that claim it requires metaphysically dubious mental actions of the mind.

Reiss’ and Aristotle’s implicit claim that intentions are internal almost has to be right. Intentions are in the mind, and the mind is internal. But the mind has its own internal structure, and one of the separations we routinely observe is the separation between intentions and the non-affective mind, between them and the first-person “I,” what might be called consciousness. The intuition is that this must all be internal. But just relying on intuition is insufficient as it presents a conflicting picture.

Many common idioms seem to place intentions outside the actor, e.g. “the spirit moved me,” “I was swept away by my emotions,” or “my urges got the better of me.” For present purposes what matters is that the relationship between intentions and consciousness that our intuitions deliver is all over the map. And, of course, these usages are metaphorical and the metaphors may not accurately reflect the structure of the mind, so let us not make too much of them. And the seeming obviousness of the relationship between intentions and our conscious decisions to pursue some end turns out to be far from clear. So merely gesturing at intentions is hardly an explanation of behavior. Some account of intentions, consciousness, and action has to be given.

For present purposes, we adopt a view that dates back at least to Hume (1978), see also Shaw (1998). For Hume, intentions are “passions,” or what today we would call affective states. Intentions, emotions, and urges are all affective states, in the same category with other affective states like wants, preferences, and cares. They are valenced states, in the sense they are directional, they guide us, they incline us toward one sort of behavior or line of thought rather than another. In contrast, the conscious mind is where what Hume calls the “intellect” operates (see McShea, 2017). Consciousness allows us to move around premises, construct possibilities, create scenarios, apply reason, and draw conclusions about hypothetical alternative outcomes by mixing and matching them. When thinking about buying bread, one might entertain dozens of different premises about which breads to buy, where to buy them, where to find the car keys to get to the store, etc. All of this happens at the level of consciousness. But, as Hume astutely observed, the various possibilities tossed around in our consciousness are impotent when it comes to causing action (Shaw, 1998). No matter how sound a conscious judgment about how to buy bread might be, without the additional ingredient of a desire to buy it—a want, a preference, an inclination, in other words an *intention*—no behavior will ensue. In contrast to intentions, which are thoroughly affective, the activity of consciousness is *non-affective*. It is the intention to buy bread that teleologically directs one’s thinking about the steps necessary to buy it, and that pushes one along some pathway toward bread-buying behavior. Hume’s account of the intentions and consciousness is at least consistent with the dominant themes in our intuitions, and while some might be leery of this view, we don’t mind staking a position on Hume’s laurels—one could do much worse.

The proposal here is that consciousness and behavior are embedded within intentions—and indeed within affect generally—that they occupy some space internal to the intentions, guided by them as a charged particle is guided by an electric field it moves within. For conscious thought and behavior, guidance, direction, and motive force come from above, from the larger field, from the intentions. It is this nested, hierarchical relationship that creates the persistence and plasticity we see in motivated behavior. In the grip of a desire to buy bread, consciousness flits about: “Bread ... I’ll go to the bakery ... I need my car keys ... where are my keys... I hate those keys... they’re too big... maybe I shouldn’t have bought the car with the big key fob ... I’ll drive by the dealership on my way to the bakery ... there’ll be less traffic that way anyway ... maybe I need two loaves ... etc.” Other intentions might intrude. One’s phone rings, and the desire for bread is temporarily backgrounded during the

ensuing 20 min conversation with a friend. But after hanging up, the desire returns. Throughout the many mental excursions, the intention to get the bread is a constant, a steady hum in the background, nudging one's thinking back to a bread-getting theme no matter where it wanders. Thoughts are rapid, fleeting. Intentions are slow, persistent. The charged particle may start to wander, say under the influence of other forces, but wherever it goes, there is the electrical field nudging it back to its original trajectory. Likewise wherever the mind goes, there is the intention, nudging it back to an intention-satisfying trajectory. Relative to the conscious mind, intentions are a kind of field, directing it from some location external to consciousness, somewhere "out there."

This view of intentional states has three virtues beyond its Humean roots. First, it is consistent with what we do know from psychology and neuroscience on the relationship between intentions and thought or action. In psychology, a fair amount is known about the various triggers of intentions, both conscious and unconscious, and the various factors that bias intentions, but only very recently has there been much focus on actual mechanism by which intentions drive conscious thought and action (Brass et al., 2013). In neuroscience, there has been considerable work in recent decades on emotion (Barrett, 2015; Damasio, 2003; LeDoux, 1996) but little is known about the central issue here, the physical structure of the neural mechanisms involved in the connection between thought and action and the softer affective process like motivation, volition, wanting, and intentionality (which for present purposes we treat as closely related). Only very recently has neuroscience taken on this problem (Brass et al., 2013). One common model of volition invokes chains of causation that begin with pre-frontal brain areas and end with certain motor areas (Haggard, 2008). Others start the process at lower levels, taking the seat of affect to be certain brainstem areas, with the causal arrow propagating upward to higher cortical areas. Such models are called hierarchical in the literature (Berridge, 2004), because there is a flow from physically older and anatomically lower down brain areas to those younger and higher up, but in our scheme, they are non-hierarchical, in that the structure involves no physical nestedness, entities within larger fields. On the other hand, some doubt the validity of these models, seeing the structure of affective processes as much more complex (Berridge, 2004) and distributed widely in a neuroanatomical network (Davidson, 2003; Symmonds & Dolan, 2012). That view is consistent with what we propose, that affective structures are large and enveloping, relative to the structures responsible for thought and action. In any case, from a perspective outside the field, it seems safe to say that little is known. Indeed, Roskies (2010) remarks that current neuroscience has not had much impact on our concept of volition.

Second, the view is consistent with certain folk psychological observations about intentions, in particular that they resist precise description. It is difficult to describe the feeling of intending something. Intentions appear to us only indistinctly, vaguely. Our view of them is only partial, as a cloud appears when we fly through it, at once everywhere and nowhere in particular. Finally, the view is consistent with some of the central observations in hierarchy theory, the study of the physical properties of nested systems (Campbell, 1958; Salthe, 1985; Simon, 1962; Wimsatt, 1974, 1994). In most natural hierarchical systems, lower-level entities move and change rapidly,

relative to the larger system in which they are embedded, just as thoughts move quickly relative to intentions. Upper levels act downwardly on the lower-level entities affecting their behavior, but only on average, leaving these entities considerable freedom of movement, as an intention guides thinking in certain directions without precisely determining it. And then, upper levels are resistant to influence from below, lower level entities having insufficient mass and power to move the upper, generally speaking, just as intentions resist alteration by mere thoughts. It is only with difficulty, or by deploying certain psychological tricks, that we can decide to not intend something, if we can do it at all.

Again, a word is needed about the word “internal.” Intentions seem to be internal to us, occurring within our minds. Some might say that they are larger still, that like emotions they are embodied, involving more than the mind, perhaps requiring the participation of the entire body. But either way, and consistent with the proposal here, they are external to conscious thought. But, if one doesn’t quite agree with this view of intentions, the burden of proof now rests on those who would locate them internal to consciousness.

6.2.1 Telescoping levels

We are not done. We have argued so far that intentions are external to the intellect, to consciousness, i.e., to non-affective processes. But they would still seem to be in the mind and therefore internal to us, consistent with Aristotle’s implicit suggestion that they are internal natures and Reiss’s suggestion that they arise within us, that they are self-contained, so to speak. Our reply is the same as it was for genes, that hierarchies telescope upward. For intentions, it is that they have a source of guidance larger than themselves, larger than the mind. We argue that this guidance arises from the environment, especially social fields, and also from—as for the genes—the evolutionary ecological fields in which we are all embedded.

Walking down the street, the smell from a nearby bakery wafts in our direction, evoking a desire for and the intention to buy bread. Or perhaps that intention is evoked by a newspaper ad, a buy-one-get-one sale at the bakery, noticed in the paper that morning. Both count as higher-level fields in that they are larger than the shopper, their reach extending over a time and space range that is far greater than any single individual. The pleasant bakery smells wafting through the streets influence the intentions of many potential customers and the ad for the sale nudges the intentions of many throughout the newspaper’s readership. Sometimes intentions do appear to vault into existence from nowhere, to spring up from some deep, ineffable psychological well: *I simply want the bread*, with no further explanation needed. But this appearance is often illusory. To see the upper-level fields that routinely evoke intentions, and to appreciate their effectiveness, one need look no further than the astonishing predictive power of modern advertising. Using data collected from social media profiles, advertisers know that exposing users to certain kinds of external stimuli will elicit in a certain percentage of users an intentional stance toward the advertiser’s product. They will want it. Of course, the effect is always on-average. Like the charged particle in the electric field, we retain considerable freedom and no advertiser can predict the buying behavior of any given user. We respond to other

forces, some internal and some external. What these advertisement “fields” do is nudge our intentions, persistently, on long time scales, toward the advertised object. Other social fields work the same way. At a small scale, we respond to an invitation to a family holiday dinner with an intention to attend. At a larger scale, people in a neighborhood keeping their yards tidy sets up a social pressure to which some will likely respond with intentions to mow their own (however grudgingly). And at a still larger scale, a lowering of the national home-mortgage interest rate is a change in an external field that guides some people toward an intention to buy a home.

Still there seem to be cases where intentions aren’t guided by environmental fields like bakery smells or social fields like large-scale advertising. For example, the fear of spiders seems to be innate in some people, given that it has been observed in children as young as 6 months old (Hoehl et al., 2017). In this case, intention takes the form of a fear. And it seems clear that the fear does not arise solely from an external stimulus, i.e. spiders, as there is no reason why spiders should elicit fear from someone who has no prior experience of spiders. Such fear is only explainable by reference to evolved traits programmed genetically, and therefore seemingly internal to the child. Our response here should be quite predictable. An innate fear of spiders has two sources, one in development (explaining its origin on the scale of a lifetime) and another in evolution (explaining it on the multimillion-year time scale of the species individual). The one in development arose in the structuring of the child’s brain by various morphogenetic fields, shaping the placement and pattern of connection of nerve cells in the development of the individual. Insofar as these complex neural interactions extend throughout a child’s developmental trajectory, they are external to the fear. The one in evolution arose in the structuring of the genome by the ecological field in which the species evolved, one in which an innate fear of spiders may have been favored. Thus the ultimate guiding force behind a six-month-old’s affective stance towards spiders, his fear, is not something deep inside, internal to the mind of the child. Rather it arises from external structures, the brain-generating fields in development and the much larger and historically deep ecological fields that drive evolution.

Intentions, then, are not as Aristotle and Reiss would have it, as stand-alone points of origin for thoughts and behavior. Intentions, like all other teleological entities, are directed in a predictable but non-deterministic manner by fields that are external to them. The inclination to treat intentions as internal to ourselves likely stems from the idea that we are unrestrained and free to set our own goals and direct our own lives. And while this is true to an extent because our intentions are our own, this ignores the evidence that strongly suggests they are directed in all sorts of predictable ways.

7 Conclusion

The NAD, which has its roots in Aristotelian teleology, is at the center of many of the criticisms of the use of teleology in biology. However, as we have argued, the NAD is a not a legitimate conceptual division, nor is it necessary. Teleology in all of its manifestations can be explained with an expanded version of Aristotle’s artifact

model, one that invokes purely external forces, specifically, upper-level fields. And while we have argued against the existence of an Aristotelian-type internal natures, we think Aristotle was quite right when he said it may well be impossible to gain knowledge of nature without understanding the sources of the directionality (*Physics* 193a). It's just that teleology does not require positing any mysterious internal natures that somehow guide objects to distant future ends. Even in the cases of genes and intentions, where many are most inclined to posit something that might look like an internal nature, external fields are able to account for goal directedness. Indeed, in some lights, they seem essential. For genes, information about how to guide macroscale development is unavailable to them without external fields, like morphogenetic fields. For consciousness, decisions about what to think and what to do are impossible to make without looking to external affective states—wants, preferences, intentions—and other fields at higher levels in the hierarchy, like culture and selection. And these external fields are not only metaphysically unproblematic, they are scientifically plausible. In this way, field theory is able to deliver the epistemic benefits that Aristotle saw in genuine teleological explanations.

In addition, moving away from the search for internal guiding forces encourages us to explore the unnecessary conceptual separation of artifacts and nature that has loomed large in, and in some ways distorted, discussion in other areas. It encourages us to examine, for example, whether it is legitimate to invoke a strong division between natural and artificial intelligence, between natural and artificial life, and perhaps even between natural and artificial environments. What we should be focused on is understanding and explicating the upper-level fields that guide goal-directed entities.

Finally, an externalist teleology opens the door to a unification. Field theory unifies disparate goal-directed phenomena—ranging from developmental systems to human intentions to goal-directed artifacts to natural selection and evolution—under a single explanatory rubric. Indeed, in the absence of any other unifying scheme, field theory seems to be the only approach that both takes these goal-directed phenomena seriously, takes them to be real, and finds the shared sense in them. An externalist teleology may be not just safe for biology, but essential.

Acknowledgements We would like to thank the John Templeton Foundation for its generous support (grant #61408, "Increasing Complexity: The First Rule of Evolution?"), which helped make the work that contributed to this paper possible. We would also like to thank Mark Brennan and the anonymous referees who provided valuable feedback on earlier versions of this paper.

References

- Allen, C., & Neal, J. (2020). Teleological notions in biology. In E. N. Zalta (Ed.), *The stanford encyclopedia of philosophy* (Spring 2020 Edition). <https://plato.stanford.edu/archives/spr2020/entries/teleology-biology>.
- Atamian, H. S., Creux, N. M., Brown, E. A., Garner, A. G., Blackman, B. K., & Harmer, S. L. (2016). Circadian regulations of sunflower heliotropism, floral orientation, and pollinator visits. *Science*, 353, 587–590.
- Babcock, G. (2020). Are synthetic genomes parts of a genetic lineage? *The British Journal for the Philosophy of Science*. <https://doi.org/10.1093/bjps/axz046>.

- Balme, D. M. (1962). Development of biology in Aristotle and Theophrastus: Theory of spontaneous generation. *Phronesis*, 7, 91–104.
- Barrett, L. F. (2015). Ten common misconceptions about psychological construction theories of emotion. In L. F. Barrett & J. A. Russell (Eds.), *The psychological construction of emotion*. (pp. 45–79). The Guilford Press.
- Basl, J., & Sandler, R. (2013). Three puzzles regarding the moral status of synthetic organisms. In G. E. Kaebnick, & T. H. Murray (Eds.), *Synthetic biology and morality: Artificial life and the bounds of nature* (pp. 89–106). MIT Press.
- Berridge, K. C. (2004). Motivation concepts in behavioral neuroscience. *Physiology and Behavior*, 81, 179–209.
- Boorse, C. (1976). Wright on functions. *The Philosophical Review*, 85(1), 70–86.
- Brass, M., Lynn, M. T., Demanet, J., & Rigoni, D. (2013). Imaging volition: What the brain can tell us about the will. *Experimental Brain Research*, 229, 301–312.
- Campbell, D. T. (1958). Common fate, similarity, and other indices of the status of aggregates of persons as social entities. *Behavioral Sciences*, 3, 14–25.
- Charlton, W. (2006). *Aristotle physics books I and II*. Clarendon Press.
- Craver, C. F. (2001). Role functions, mechanisms, and hierarchy. *Philosophy of Science*, 68(1), 53–74.
- Cummins, R. (1975). Functional analysis. *The Journal of Philosophy*, 72(20), 741–765.
- Damasio, A. (2003). *Looking for Spinoza*. Harcourt Books.
- Davidson, R. J. (2003). Seven sins in the study of emotion: Correctives from affective neuroscience. *Brain and Cognition*, 52(1), 129–132.
- Ettensohn, C. A. (1990). The regulation of primary mesenchyme cell patterning. *Developmental Biology*, 140, 261–271.
- Forster, E. S. (1944). Aristotle's De Generatione-Aristotle: Generation of Animals. With an English translation by AL Peck. (Loeb Classical Library.) Pp. lxxviii+ 608. London: Heinemann, 1943. Cloth, 10s. (leather 12s. 6d.) net. *The Classical Review*, 58(2), 54–55.
- Fox Keller, E. (2010). *The mirage of a space between nature and nurture*. Duke University Press.
- Garson, J. (2008). Function and teleology. In A. Plutynski & S. Sarkar (Eds.), *A companion to the philosophy of biology*. (pp. 525–549). Blackwell.
- Garson, J. (2016). *A critical overview of biological functions*. Springer International Publishing.
- Ghiselin, M. T. (1974). A radical solution to the species problem. *Systematic Zoology*, 23, 536–544.
- Goodwin, B. (1994). *How the leopard changed its spots: The evolution of complexity*. Princeton.
- Haggard, P. (2008). "Human volition: Towards a neuroscience of will. *Nature Reviews: Neuroscience*, 9, 934–946.
- Hoehl, S., Hellmer, K., Johansson, M., & Gredebäck, G. (2017). Itsy bitsy spider...: Infants react with increased arousal to spiders and snakes. *Frontiers in psychology*, 8, 1710.
- Hull, D. L. (1967). *The conflict between spontaneous generation and Aristotle's metaphysics*. (pp. 245–250). Université Laval.
- Hull, D. L. (1976). Are species really individuals? *Systematic Zoology*, 25, 174–191.
- Hume, D. (1978 [1740]). *A Treatise of Human Nature*. Edited by L. A. Selby-Bigge, 2nd ed. Oxford: Oxford University Press.
- Kauffman, S. (1993). *Origins of order: Self-organization and selection in evolution*. Oxford University Press.
- Lang, A. R. G., & Begg, J. E. (1979). Movements of *Helianthus Annuus* leaves and heads. *Journal of Applied Ecology*, 16(1), 299–305.
- LeDoux, J. (1996). *The emotional brain*. Simon and Schuster.
- Lee, J. G., & McShea, D. W. (2020). Operationalizing goal directedness: An empirical route to advancing a philosophical discussion. *Philosophy, Theory and Practice in Biology*. <https://doi.org/10.3998/ptpbio.16039257.0012.005>.
- Lennox, J. (2001). *Aristotle's philosophy of biology: Studies in the origins of life sciences*. Cambridge University Press.
- Lennox, J. (2019). Aristotle's Biology. In E. N. Zalta (Ed.), *The stanford encyclopedia of philosophy* (Winter 2019 edn). <https://plato.stanford.edu/archives/win2019/entries/aristotle-biology>.
- Leunissen, M. (2010). *Explanation and teleology in Aristotle's science of nature*. Cambridge University Press.
- Levin, M. (2012). Morphogenetic fields in embryogenesis, regeneration, and cancer: Non-local control of complex patterning. *Bio Systems*, 109(3), 243–261.
- Lewens, T. (2004). *Organisms and artifacts: Design in nature and elsewhere*. MIT Press.
- Matthen, M. (1997). Teleology and the product analogy. *Australasian Journal of Philosophy*, 75(1), 21–37.

- Mayr, E. (1988). The multiple meanings of teleological. In E. Mayr (Ed.), *Towards a new philosophy of biology*. (pp. 38–66). Harvard University Press.
- Mayr, E. (1992). The idea of teleology. *Journal of the History of Ideas*, 53, 117–135.
- McShea, D. W. (2012). Upper-directed systems: A new approach to teleology in biology. *Biology and Philosophy*, 27, 663–684.
- McShea, D. W., et al. (2016). Hierarchy: The source of teleology in evolution. In N. Eldredge (Ed.), *Evolutionary theory: A hierarchical perspective*. (pp. 86–102). University of Chicago Press.
- McShea, D. W. (2017). Logic, passion, and the problem of convergence. *Journal of the Royal Society Interface Focus*, 7, 20160122.
- Nagel, E. (1979). *Teleology revisited and other essays in the philosophy and history of science*. Columbia University Press.
- Neander, K. (1991). Functions as selected effects: The conceptual analyst's defense. *Philosophy of Science*, 58(2), 168–184.
- Nijhout, F. (1990). Metaphors and role of genes in development. *BioEssays*, 12, 441–445.
- Reiss, J. (2009). *Not by design: Retiring Darwin's watchmaker*. UC Press.
- Rose, D., Schaffer, J., & Tobia, K. (2018). Folk teleology drives persistence judgements. *Synthese*, 1–19.
- Roskies, A. L. (2010). How does neuroscience affect our conception of volition? *Annual Review of Neuroscience*, 33, 109–130.
- Salthe, S. N. (1985). *Evolving hierarchical systems*. Columbia University Press.
- Shaw, D. J. (1998). *Reason and feeling in Hume's action theory and moral philosophy*. The Edwin Mellen Press.
- Simon, H. A. (1962). The architecture of complexity. *Proceedings of the American Philosophical Society*, 106, 467–482.
- Sommerhoff, G. (1950). *Analytical biology*. Oxford University Press.
- Symmonds, M., & Dolan, R. J. (2012). The neurobiology of preferences. In R. Dolan & T. Sharot (Eds.), *Neuroscience of preferences and choice*. (pp. 3–31). Elsevier.
- Thompson, D. W. (1942). *On growth and form*. University Press and Macmillan.
- Van Valen, L. (1973). Festschrift. (Book reviews: Evolutionary biology. Vol. 6). *Science*, 180(4085), 488.
- Vandenbrink, J. P., Brown, E., Harmer, S., & Blackman, B. (2014). Turning heads: The biology of solar tracking in sunflower. *Plant Science*, 224C, 20–26.
- von Bertalanffy L. (1969). *General system theory: Foundations, development, applications*. G. Braziller; New York.
- Waters, C. K. (2007). Causes that make a difference. *The Journal of Philosophy*, 104(11), 551–579.
- Weiss, P. A. (1969). The living system: Determinism stratified. *Studium Generale*, 22, 361–400.
- Wimsatt, W. C. (1974). Complexity and organization. In K. F. Schaffner & R. S. Cohen (Eds.), *Philosophy of Science Association 1972*. (pp. 67–86). Dordrecht, Netherlands: D. Reidel.
- Wimsatt, W. C. (1994). The ontology of complex systems: Levels of organization, perspectives, and causal thickets. *Canadian Journal of Philosophy*, 20(supp), 207–274.
- Wimsatt, W. C. (2007). *Re-Engineering philosophy for limited beings: Piecewise approximations to reality*. Harvard University Press.
- Wright, L. (1973). Functions. *The Philosophical Review*, 82(2), 139–168.
- Wright, L. (1976). *Teleological explanations: An Etiological analysis of goals and functions*. University of California Press.
- Zwier, K. (2018). Methodology in Aristotle's theory of spontaneous generation. *Journal of History of Biology*, 51(2), 355–386.

Authors and Affiliations

Gunnar Babcock¹  · **Daniel W. McShea¹**

✉ Gunnar Babcock
gunnar.babcock@duke.edu

Daniel W. McShea
dmshea@duke.edu

¹ Department of Biology, Duke University, Durham, NC, USA