

Freedom and purpose in biology



Daniel W. McShea

Biology Department, Duke University, Box 90338, Durham, NC, USA

ARTICLE INFO

Article history:
Available online 6 January 2016

Keywords:
Evolution
Goal-directedness
Hierarchy
Teleology
Upper direction
ZFEL

ABSTRACT

All seemingly teleological systems share a common hierarchical structure. They consist of a small entity moving or changing within a larger field that directs it from above (what I call “upper direction”). This is true for organisms seeking some external resource, for the organized behavior of cells and other parts in organismal development, and for lineages evolving by natural selection. In all cases, the lower-level entity is partly “free,” tending to wander under the influence of purely local forces, and partly directed by a larger enveloping field. The persistent and plastic behavior that characterizes goal-directedness arises, I argue, at intermediate levels of freedom and upper direction, when the two are in a delicate balance. I tentatively extend the argument to human teleology (wants, purposes).

© 2015 Elsevier Ltd. All rights reserved.

When citing this paper, please use the full journal title *Studies in History and Philosophy of Biological and Biomedical Sciences*

1. Freedom

An American teenager graduates from high school in 1975 and sets out into the world. He finds a job on a steamer bound for Belgium, on arrival hitchhiking to Liege where he joins a farming commune. Tiring of this after a few weeks, he packs up one morning and takes a train to the south of France, where he connects with a group of amateur boat builders. Their goal is to sail the Mediterranean, but progress is slow, and after several months, the boat is still unfinished. Frustrated, he sells his belongings and flies back home. He rests for a week at his parent’s house in New York, then takes off again, this time with the idea that he will walk across the country, from New York to Los Angeles, stopping at roadside diners and meeting “real” people. He makes it as far as his cousin’s house in New Jersey, where he moves in and lives rent-free for the next three years. What is this young man up to? He might say that he is doing his own thing, being true to himself, trying to escape the suffocating constraints of middle-class expectations, of capitalism, of his parents. In the language I shall use here, he is trying to be free.

Let us not worry about whether this is the best way to understand freedom. The experience of that generation suggests that

in some ways it is not, or at least, that it is only part of what freedom means. Still, it is one way to understand the word, one with a recent cultural history and one that is consistent with current folk understandings: freedom in the sense of the absence of organized external constraints. I shall come back to our peripatetic teenager.

The argument of this paper is that entities that are free, that move or change in the absence of large-scale or upper-level forces, tend to wander. They tend to vary. So robust is this expectation that when we do not see wandering and variation, when an entity’s behavior is highly organized and directed, it strongly suggests the existence of upper-level directing forces. I argue that this partitioning of behavior into a free component and an upper directed component has consequences for our understanding of teleology, of goal-directedness and purpose in biology. Both components are essential to teleological behavior. Indeed, I shall argue that teleological behavior arises at a point in the middle, where freedom and direction by higher-level forces are in a delicate balance.

More generally, my mission here is to give a naturalistic account of teleology in biology, an account from the perspective of an engineer. I am addressing the question of how teleological systems are structured physically. How must systems be organized in order to seem to behave teleologically? I do not attempt anything like an analysis of concepts, or of how teleological terms are used. And

E-mail address: dmcshea@duke.edu.

therefore, the argument intersects most of the literature in this area only obliquely. In particular, the large modern-era literature on the naturalization of meaning and function—especially teleosemantics (MacDonald & Papineau, 2006) and treatments of functions as selected effects (Millikan, 1984; Neander, 1991; Wright, 1973)—is not directly relevant. Rather, the line of thought offered here grows out of the pioneering studies of the physical properties of hierarchically structured systems by Simon (1962), Campbell (1958), Wimsatt (1974, 1994), and Salthe (1985, 2009). My concern is not with what goal-directedness means but with how it works, extending the project begun in McShea (2012).

The title, freedom and purpose in biology, is somewhat immodest. I hope not overly so. What I offer is a way to think about causation that applies across a wide swath of biology. It is a perspective that unifies our understanding of how change works in evolution, in the development of organisms from embryo to adult, in the behavior of simple organisms, and even in the behavior of complex ones, including us. Further, it explains teleology, purposefulness. The Darwinian view is generally regarded as sufficient here. But it is not. Natural selection tells us about the process that gives rise to teleology. But it does not explain how it works. It does not point to any particular mechanism behind the mysterious moment-to-moment behavior, the seeking behavior, of seemingly purposeful entities. Selection also cannot explain the apparent goal-directedness of the Darwinian process itself, of selection itself. The viewpoint developed here does all this. The apparently purposeful movements of certain entities and the apparent purposefulness of adaptation emerge as special features of nested systems, as instances of a more general process.

Throughout, I often modify teleological terms with the word “seemingly” or sometimes “apparently.” This is not an eliminativist move. The strange behaviors and capacities that we call teleological are real and have real causes. But almost no one today believes that the future causes the past, that literal “goals” which by definition lie in the future can have any causal influence on the present. This goes for human wanting, thinking, and behaving too. The so-called “goals” that guide our present thinking and behavior are hoped-for hypotheticals, existing in the present and presently intended to have a forward-causal connection to the future. We want and try to make things happen. But, as everyone knows, there is no backward-causal connection. There is nothing literally teleological about my getting in the car to drive to a picnic. The future picnic has exactly zero role in the process and indeed may not be there when I arrive. (If, say, I had the wrong day.) Indeed, maybe it is time to stop implying backward causation in our choice of words. My use of “seemingly” and “apparently” is a gesture in that direction.

2. Nestedness, freedom, and upper direction: three biological examples

I begin with three biological examples, one having to do with the movements of microorganisms, a second having to do with organismal development, and a third with evolution. In all of these systems there is an entity of some kind, one that can move freely on its own but that seems to behave teleologically when immersed within, nested within, a higher-level structure or field of some kind. The point is to illustrate by example what I mean by freedom and nestedness and also to introduce a new term, upper direction. (The next section offers a more formal discussion of upper direction.)

First, imagine a group of ten thousand bacteria, temporarily occupying the same cubic millimeter of water in a pond. I say temporarily because they are free to move and therefore unlikely to stay in that same cubic millimeter for long. Normally, in this species of bacterium, when an individual moves in open water its

flagellum propels it in a series of what are called “straight runs,” interrupted occasionally by tumbles that randomly reorient it. It darts one way, tumbles, zooms off in another direction, tumbles, then zips ahead again, and so on, producing a kind of a random walk.

Interestingly, its behavior would change if it detected food, if it suddenly found itself inside, say, a gradient of increasing concentration of some food molecule, like the amino acid aspartate, leaching into the pond and diffusing from some distant source. When the bacterium detects an increase in aspartate concentration, a signal cascade within the organism kicks into action, triggering an increase in the length of the straight runs. It still tumbles, and still reorients randomly in each tumble, but the straight runs become longer, with the result that, on average, the bacterium tends to move up the food gradient, closer to the food source. That is what would happen in all of the bacteria in this example if they were embedded within a food gradient. They would all increase the lengths of their straight runs and move on average up the gradient. Notice what has happened here. With the introduction of a food gradient, the structure of the system becomes hierarchical. The bacterium becomes an entity nested within a higher-level structure, the food gradient, and the higher-level structure directs the bacterium toward the food source. We can say that the bacterium is now “upper directed.”

Let us return to the original setup. Suppose no food gradient or any other large-scale external factor affects the movement of the bacteria. And so they wander hither and thither. Straight run, followed by random tumble, followed by straight run, followed by random tumble, going nowhere in particular. In the language I am adopting here, the bacteria are free. They are free in the sense that they are governed entirely by the interaction between themselves and local forces, their movements determined by a combination of their own internal structure and the local external environment. More concretely, the movement of each bacterium during a tumble is determined by a combination of the mechanism governing its flagellar micro-movements and the currents, density, and viscosity of the water immediately around it. Its movement is independent of all of the other bacteria. And it is independent of large-scale influences. There are no factors acting over a large-area—no food gradients and no other large-scale fields, currents, or gradients.

In some cases, including this one, the interaction of an entity with local external forces is complex, so that the entity follows a random walk (technically, since the interaction is understood here to be deterministic, a pseudorandom walk). But freedom need not produce randomness. A bacterium with internal mechanisms programmed to make a 60° turn every centimeter, or to always swim in a direction opposite to the microcurrent in its immediate environment, is also free. Behavior is free whenever, and to the extent that, control is local.

Notice that the freedom of each bacterium makes for highly predictable behavior by the group. We do not know where any individual will go, but the group as a whole will certainly diffuse, expanding outward from the original cubic millimeter that contained them. Another way to say this is that the variance in their positions—whether measured along an east–west axis or a north–south axis or any other axis—will increase. Free movement among members of a group produces an increase in variance in the group as a whole. The group spreads out.

In sum, here is how I propose to describe the situation: in the absence of upper direction, in the absence large-scale external causes, each bacterium is free, and the result is that it wanders and the variance in position among all of the bacteria increases.

A second example. Consider the free movement of cells in a different context, the much larger eukaryotic cells in a

multicellular organism. I have in mind the relatively unconstrained growth of cells in the galls and tumors that afflict many plants and animals, and even the growth of certain cell populations in laboratory culture. In extreme cases, the cells replicate without significant constraint from surrounding tissues or other media, they do not differentiate into distinct cell types, and they are not organized into higher-level structures such as tissues or organs. If we were to plot the locations of a tumor-cell clone on a three-dimensional grid, what we would see is a dispersion of the cells in space, an increase in variance in their positions. In some cases, the cells are also free to vary in shape and in various aspects of their physiology. In other words, they wander, not just in real space but in various abstract spaces that we could construct to describe them. For example, in an abstract shape space defined on one axis by cell length and on another axis by cell width, the expectation is that in a cell culture that started with a single cell, the points representing it and its daughter cells in this shape space would tend to disperse over time. The variance in their shapes would tend to increase. In present terms, this propensity to vary is freedom. The cells are free.

Contrast this situation with the seemingly teleological movement of cells in a developing embryo. Fig. 1A shows a stage in the normal development of a sea urchin larva, an early stage in which the larva consists of a hollow ball of cells. In the next step, a group of cells arises at the south pole of the embryo (Fig. 1A, left) and begins to migrate, eventually coalescing to form a ring of cells lying just below the equator (Fig. 1A, right). The membranes of adjacent cells in the ring then fuse so that the cytoplasm of the cells merges, and it is within this shared cytoplasm that the larval skeleton is laid down (not shown). In a series of experiments in the 1980s and 1990s, the developmental biologist Charles Etensohn (Etensohn, 1990; Etensohn & McClay, 1988) transplanted cells from the south pole to other locations within the developing embryo (Fig. 1B, left) and showed that they migrate to the subequatorial target (Fig. 1B, right), ultimately giving rise to a larval skeleton as in normal embryos. More generally, experiments over a number of decades have shown that a subequatorial ring and a normal-looking skeleton are produced over a wide variety of types of disruption to cell migration. These (and other) findings have suggested that the individual cells are *not* preprogrammed to migrate to any particular target location,

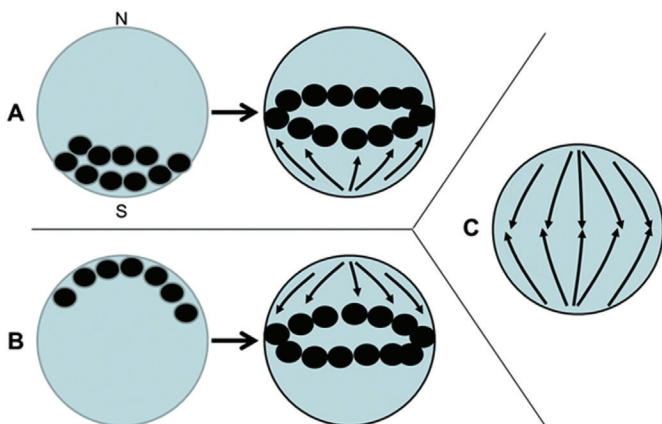


Fig. 1. Early steps in the development of a sea urchin larva. A. Normal development in which specialized cells, arising at the larval south pole migrate up to the equator, where they will later give rise to the larval skeleton. B. A transplantation experiment in which the south-pole cells are removed and injected at the north pole. Despite the displacement, the cells migrate to the equator, giving rise to a normal skeleton. C. A representation of the upper-level field that likely accounts for this movement pattern, a gene-activation field—or other morphogenetic field—that directs cells to the target region from anywhere in the developing embryo.

each cell to a unique address within the embryo. Rather, the cells' distribution seems to be loosely governed by a global field of some kind—perhaps a chemical gradient or a gene-activation pattern—that direct any cells lying within it toward the target area just below the equator. It is because the field is large and widespread that it can provide this general direction to cells no matter what their initial position within the embryo (Fig. 1C).

Interestingly, in this system, the variance in the position of the cells can actually increase despite the large-scale directing field, owing to the geometry of the situation. If, say, the cells are injected at the north pole, as in Fig. 1B, so that the initial variance in position is low, the variance will increase as the cells migrate toward the subequatorial ring. Still, even as they approach that destination, the variance in position will be lower than it would have been in the absence of the large-scale field. If no such field were present, presumable the cells would have distributed themselves over the whole surface of the embryo. This decrease in variance, relative to what it would be otherwise, is the result of the direction provided by the field. In my terms, upper direction creates a restriction on freedom. As I shall explain later, it is this upper direction that makes the cells appear to “seek,” that gives the system its teleological feel.

The contrast here is between cells that are relatively free—tumor cells, gall cells, and cells in culture—and those that are nested within and directed by a higher-level structure, a field of some kind within a developing embryo.

A third example, at the scale of lineage evolution. Imagine a population of fruit flies in a laboratory. The flies are kept under near-ideal conditions. They are well fed. Their cages are free of disease. They are allowed to breed freely. And when they breed, all of their offspring are nurtured to adulthood under the same near-ideal conditions. Naturally, in this population, like any in the wild, mutants arise. For example, in one laboratory mutant, called “ocelliless,” the light-sensitive structures normally present at the midline of the head are absent. (They are light-sensitive but they are not eyes.) In another, called “short wing,” the wings are irregularly incised. In another, “comb gap,” the legs are irregularly shortened.

For these laboratory mutants, raised under conditions nearly ideal for their survival and reproduction, natural selection is reduced. Presumably many, if not most, would not survive in the wild, but under reduced selection, we get to see this natural variation, in virtually every dimension that we can measure, in some cases extreme variation.

Actually what we see is variation in two senses, among individuals and among parts within an individual. The variation among individuals is obvious. An initially homogeneous population of normal flies transforms over time into a population that includes the descendants of the mutants. The variation in morphology among individuals increases. If we were to construct an abstract space, a morphospace, with axes defined by various aspects of morphology and each individual represented by a point in the space, what we would see over time is the gradual dispersion of points as mutants accumulated in the population. The lineages are, to some extent, free.

Variation among parts within individuals is a little trickier. To see the sense in which it increases, consider “short wing.” In this mutant, the left and right wings differ from each other more than they do in the wild type. Likewise in “comb gap,” the legs differ in length (and in other features) from each other more than they do in the wild type. In other words, in these mutants, the variance—the degree of differentiation—among the parts of the animal is greater. This within-organism differentiation is complexity, in a technical sense of the word that has become current in biology (Doolittle, 2012; McShea, 1996; Valentine, Collins, & Meyer, 1994). And these

two mutants are morphologically more complex than the wild type.¹

In a recent book, Robert Brandon and I (McShea & Brandon, 2010) argued that in the absence of selection, or under reduced selection, the expectation in evolution is an increase in both diversity—differentiation among individuals—and complexity—differentiation among parts within individuals. And the reason is simply that in biological systems, to put it informally, accidents happen and accumulate, leading to an increase in variance. We call this tendency to increase the zero-force evolutionary law (ZFEL). It is a statistical law in that even in the absence of selection, decreases can occur by chance. For example, the loss of ocelli in ocelliless is a loss of a type of part, and therefore a decrease in the degree of variation among parts, and therefore a decrease in complexity. So decreases can occur, but overall the expectation is gain. The spontaneous tendency for parts to differentiate produces an increase in complexity, on average.²

I hope the analogy is clear between increasing diversity and complexity, on the one hand, and bacterial dispersion and tumor growth, on the other. Free variation in a set of entities causes them to disperse, the variance among them to increase, whether in a physical space or in a morphospace. Again, as in all of the examples in this section, freedom does not imply that causes are absent, only that they are independent and local. In the absence of a food gradient, individual bacterial movements are controlled by the combination of internal factors and local external forces unique to each bacterium. The same goes for a tissue-culture cell in the absence of a developmental gradient. And now, similarly, the morphology of an individual organism's offspring in the absence of selection is governed by whatever local external factors produce the mutation in the individual and the developmental constraints governing how mutation is expressed in that individual, as well as the unique accidents of the individual's life history. Freedom is local causation.

Now suppose we introduce natural selection, but with selection understood in a slightly non-standard way. Think of it as a gradient or a field imposed by ecology. Organisms are immersed in their ecology, just as bacteria are immersed in a food gradient and cells in an embryo are immersed in a developmental field. And it is the forces exerted by the ecological field in which a population is immersed—its predators, parasites, and competitors—that drive a population toward increased fitness. The usual metaphor for an organism evolving in an ecology is a fitness landscape. And the ecological field is indeed something like a fitness landscape, but instead of a passive topography directing movement of a population as hills and valleys direct a river, an ecological field actually powers or drives the movement of the population. Natural selection is upper direction of lineages by ecology.

¹ This usage—complexity as differentiation—is obviously not the colloquial sense of the word. In common usage, complexity is generally taken to include some aspect of sophistication, intricacy, and function, as well as differentiation. Colloquially we say that cars and computers and brains are complex, partly on account of the differentiation among their parts but also on account of what they can *do*. Speaking colloquially, we would not be inclined to call a fly with two different wings, a fly which probably cannot fly very well, more complex than a normal fly. However, complexity in the colloquial sense, incorporating function, has never been operationalized, at least in biology, and so cannot be treated scientifically. Unfortunately, the word and the notion have been tossed around quite a bit in the popular-science literature, leaving many people (including many scientists) with the impression that biologists using the word in the colloquial sense know what they are talking about. They do not.

² And indeed that is what Leonore Fleming found in a study of these mutants (Fleming & McShea, 2013). Among the nearly one thousand mutants considered, the great majority constituted increases in complexity. That is, more of the mutants were like short wing, showing an increase in complexity, than were like ocelliless, showing a loss.

Natural selection has another effect. It reduces the accumulation of variation, slowing the rise in both diversity (variation among individuals in the population) and complexity (variation among parts in an individual). Indeed, it can overwhelm the ZFEL, causing a net reduction in diversity and complexity. Thus, adaptation, in this view, is not only the movement of a population toward a point in morphospace where survival and reproduction are greater—driven by the large-scale forces of the ecological field in which it is immersed—but also a reduction of variation of the population around that point. Upper direction reduces variation. It reduces freedom.

Notice that it is the upper direction that gives adaptive change its teleological appearance. Lineages change somewhat haphazardly on a short timescale, under the influence of, say, drift and mutation, producing much wandering. But there is always a slight bent, a kind of bias in their movements that nudges them toward adaptive design. And the source of that bias is selection, a higher-level field.

I hope it is now reasonably clear how I am using these terms, freedom and upper direction, and also obvious that freedom and increasing variance are the default expectation. It is what we expect to see in the absence of any gradients or fields driving entities in some particular direction, in other words, in the absence of upper direction.

3. Upper direction

How is causation being understood here? How do large-scale structures like gradients and fields cause behavior in the smaller-scale entities that are contained within them? I have in mind a simple model. Imagine a helium-filled balloon hanging weightless and still in the air in the middle of a room. The helium atoms within it zig and zag, this way and that at over a thousand meters per second, bouncing off each other and occasionally bouncing off the plastic molecules that make up the skin of the balloon. I walk over to the balloon, reach up with both hands, and corral it between my palms, holding it steady for a few moments. And then slowly and deliberately, I move it three centimeters to the left. I release it, and it stays. Now it is clear what the effect of this three-centimeter shift has been. The average location of all of the helium atoms has shifted three centimeters to the left. For individual atoms, the movement may not be felt immediately. If we were to number the atoms and focus on, say, atom number 632, we might find that in the time it took to move the balloon, 632 did not experience a three-centimeter leftward shift. Indeed, there is a fair chance it moved to the right. Still, if we follow the trajectory of that molecule over a much longer time we will discover that it has moved, on average, three centimeters to the left.

Notice three things. First the system is described hierarchically, as a system that is nested, a small object within a large object. A physical arrangement like this naturally invites us to take a two-level perspective. There is the focal level, the level of interest, which in this case is the level of a single helium atom. And then there is the next level up, one level of nestedness above the helium atom, the balloon as a whole (including the atoms of gas within it).

Second, one natural way to think about causation in this story is top-down. I moved the whole balloon—plastic, helium atoms, and all—and this caused the on-average movement of atom 632. Of course, this is not the whole causal story behind the trajectory of 632. A complete story would involve the collisions of 632 with other helium atoms and with the plastic molecules of the balloon skin. More formally, we can decompose the causes of the movement of any given helium atom into two components. First, there is the on-average movement to the left caused by the movement of

the balloon as a whole, or what I would call upper direction. And second, there is the rest of its movement, governed by purely local interactions, or what I call freedom.

Third, atom 632 is somewhat free but it could be more free. It would be most free if it were outside of the balloon, bumping up against adjacent molecules of air and having its trajectory completely determined by its own elastic properties and by its collisions with these molecules, and by nothing else. Importantly, freedom and determinism are not opposites here. The movement of the atom is understood to be fully determined. So freedom in the sense intended here is a function of the atom's independence from higher-level causes. If there is no balloon, then it is free. The free-swimming bacterial cells, the tumor and gall cells, and the mutant fruit-flies' variation are also more or less free in this sense. And they become less so to the degree that they come under the influence of imposed food gradients, developmental fields, and ecology, respectively. In other words, freedom here is the degree of independence of a part from the whole.

A delicate balance. A rock that I pick off the beach could be said to be hierarchically structured. It is a whole constituted by molecular components. If I fling it out over the water in just the right way, and if it is flat enough, it will skip over the surface, the rock as a whole carrying all of its lower-level molecular components with it. The movement of each molecule is upper directed by the rock as a whole, each molecule closely following the center of mass of the rock, wherever it goes. This is a case of near-pure upper direction, with lower-level freedom nearly absent. This is true of all solid objects. The opposite is true of the lone helium atom, free in the atmosphere, undergoing collisions with nearby air molecules, completely directed by those purely local interactions. But the helium atoms in a balloon fall somewhere in between, and that is why I picked the example. In the balloon, there is a delicate balance between upper direction and lower-level freedom. Likewise for the bacterium directed by a food field, the cell constrained by a developmental gradient, and a lineage of fruit flies evolving in the wild in the loose grip of an ecology. All are subject to upper direction—and it is that upper direction I have been emphasizing in order to make a contrast with the less directed cases—but in all cases, the directed entity retains considerable freedom.

Reduction versus emergence. I have invented this odd term—upper direction—in order to avoid a confrontation with the reduction versus emergence literature. In particular I am avoiding the term “downward causation,” which figures so prominently in that literature. I am allowed to sidestep that literature, I think, because for my purposes, it does not matter whether higher-level systems and causes are reducible—either in practice or in principle—to lower-level ones. It does not matter whether the properties of the higher-level systems are non-linear or emergent. My own view is thoroughly non-reductionist, but reductionists are welcome to think of upper direction as just a heuristic, a shorthand description of how causation works in a hierarchical system.

4. Teleology

The entities in delicate balance—the bacterium in a food field, the cells in a developmental gradient, and the fruit flies evolving adaptively in the wild—share a common feature. To an observer sitting outside the system, they all look teleological, goal-directed. The bacterium seems to seek the food source. The cells in a normally developing embryo seem to seek the subequatorial target ring. A population of fruit flies evolving in a natural ecological context, under natural selection, seems to seek an adaptive morphology, an adaptive peak. In this section, I argue that the appearance of teleology—of goal-directedness, of purpose—is a

direct consequence of a delicate balance between freedom and upper directedness.³

Decades ago, Ernest Nagel (1979; see also Sommerhoff, 1950) pointed out that goal-directed systems have two salient features, persistence and plasticity. In my terms, persistence is the tendency to return to a certain trajectory following a perturbation or accident. If the cells in the developing embryo are diverted from a path toward the equator, whether by an experimenter or on account of some developmental accident, they return to that path. This diversion and return, error and correction, is persistence. And it is partly on account of persistence, Nagel argues, that certain entities seem goal-directed. In non-thinking systems like bacteria moving up gradients and cells finding target locations in an embryo, persistence looks like seeking, like trying. (In thinking systems, like us, persistence literally is trying. But I will get to that in the next section.)

Plasticity is similar. Again in my terms, it is the ability to adopt a certain trajectory from alternative starting points. A bacterium entering a food field at any point, anywhere in the entire field, will adopt an up-gradient trajectory. Likewise, in evolution, a lineage moving on a smooth fitness landscape will tend to climb a fitness peak regardless of which direction it approaches from. And many lineages, approaching and tending to move toward the peak from many different starting points (i.e., what we call convergence in evolutionary biology), will seem to be behaving purposefully.

How are persistence and plasticity possible? How do seemingly goal-directed systems do it? My answer is upper direction. Consider persistence. The cells might wander freely to some degree, but they find their targets from anywhere in the embryo because the field that provides upper direction is pervasive. The field can direct them to the target ring from anywhere in the embryo that the cells might wander because the field exists everywhere in the embryo. Likewise for plasticity. Cells can be set on a proper trajectory from any of a variety of alternative starting points, because the field that provides direction is large and exists wherever focal entities happen to start.

The same is true for persistence and plasticity in abstract spaces, such as the morphospaces within which evolving populations move. Ecology might direct a population of birds toward, say, medium-sized beaks so that they can crack a range of seed sizes (medium-sized seeds but also some larger and some smaller). And it does so whether the starting population has large beaks or small. That is plasticity, driven from above by ecology. And there is also persistence. Chance deviations from a trajectory toward medium beak size are corrected, with the result that over many generations, morphology moves on average toward a medium-sized beak. This persistence and plasticity is central to our understanding of adaptation. It is also what creates an air of mystery around the process, what gives evolution its teleological feel. We feel the urge to say that beak size changed “in order to” deal with medium-sized seeds, as though beak evolution has been directed in part by an end result, an ability to function in a useful way. Of course, no one today believes it is an end result that does the directing. So what *does* do the directing? The answer is: something physically large, present now, and continuously present over a long period of time: a large and

³ We would not ordinarily think of the molecules in the balloon as goal-directed, but it would not be hard to make them look that way. One way is simply to pretend that we are observers watching the balloon from a distance, observers that cannot see the person moving the balloon, cannot see the plastic membrane of the balloon, but can see the molecules themselves, individually, each a small dot. What we would see is an on-average movement to the left of a cloud of little dots, each dot bizarrely and counter-intuitively staying within a fixed radius of the center as the cloud moved. For some observers at least, the possibility that some teleological process was at work here would doubtless cross their minds.

stable ecological field. In effect, the widespread availability of medium-sized seeds constitutes a large, stable “resource field” that directs bird beaks—persistently and plastically—toward the appropriate size.

The same is true, incidentally, for artificial goal-directed systems. A homing torpedo, for example, motors through the water, directed by a large sound field emanating from a target ship. The torpedo is a lower-level entity, and the sound field is the higher-level field in which the torpedo is immersed. The torpedo’s behavior is persistent in that it corrects for deviations. Drawn off course by, let us say, the sounds from a passing pod of whales, the torpedo finds itself still within the sound field from the target ship and returns to a trajectory toward it. The torpedo’s behavior is also plastic in that it can adopt a trajectory toward the target ship starting from any point within the sound field, which is present over a large area.

This is why upper direction matters for teleology. It provides the direction over a large area that makes persistence and plasticity possible. But why does lower-level freedom matter? The answer is that in order to behave persistently, lower-level entities need to have enough freedom to be able to deviate, to *make* errors. Consider a laptop that is booting up. This is a very rigidly controlled system in that each flip-flop in the chip of the laptop is a slave to the flip-flops that are connected to it as inputs. There is, ordinarily, no possibility of error. And systems whose components do not make errors—do not have enough freedom to make errors—never have the chance to show any persistence. And thus they do not look goal-directed. That is why machines generally do not look goal-directed. (Torpedoes are among the exceptions.)⁴

This is the heart of my proposal: what creates the conditions for apparent teleological behavior is hierarchical structure, nestedness, a smaller entity moving or changing somewhat freely within a larger directing field, an entity with just the right balance between freedom and upper direction. Freedom is determination only by local causes, only by the entity’s unique properties and the unique local environment in which it finds itself. And it is freedom that makes for the appearance of error, the deviations we associate with persistence. Conversely, the field directs the entity over a large portion of the space in which it moves or changes. And it is the field that powers the returns, the corrections that make persistence and plasticity possible. The field can do this because it is—from the perspective of an entity moving within it—everywhere. In effect, teleology is explained spatially, not temporally.

Two conditions. For the appearance of goal-directedness, two further conditions generally need to be met. One is that the upper-level field needs to be relatively stable. If the food gradient for the bacterium is fluctuating wildly, perhaps due to water currents, then the bacterium’s movement will be erratic and will not look goal-directed. If ecological factors change too quickly, perhaps with the seed supply changing from year to year from large to small to large to medium, and so on, then there will be no organized movement of the population toward an adaptive peak. Upper levels can change, but to support persistent and plastic behavior by the lower-level entities immersed in them, they need to change relatively slowly.

The other condition is what might be called obscurity or opacity. A marble released at the inside lip of a bowl will roll toward the bottom, then up the other side, then back down toward the bottom,

and so on. Its behavior is persistent and plastic. It moves, on average, toward the bottom of the bowl from any starting point on the sides, and if perturbed (say with a fillip delivered by my finger) it returns to a trajectory toward the bottom. I might be tempted to say that the ball is seeking the bottom of the bowl, but in doing so I would be aware that I am using the word metaphorically. The ball in the bowl does not really seem goal-directed, not in the same sense that the cells in the developing embryo seem goal-directed. And the reason, I venture to guess, is that the system is too simple. It is hierarchically structured, and there is a delicate balance between freedom and upper direction, but we can understand the ball’s behavior and forces at work (gravity and friction) at a glance. The complexity of the mechanism in the embryonic cells makes the causal structure somewhat obscure, lending it that air of mystery which—for whatever psychological reason—leads us to see it as teleological. Another way to say this: the ball in the bowl actually is teleological, to some degree, on account of the hierarchical structure of causation. It is just too simple for us to dignify it with that label.

Lower-level mechanisms. Standard thinking about teleological systems generally emphasizes the causal role of lower-level mechanisms. When explaining a bacterium swimming up a food gradient or a torpedo homing in on a target ship, we focus on the intricate and impressive internal machinery of the cell and the torpedo. This focus is not wholly inappropriate. These mechanisms are important in the seeming goal-directedness of the behavior. But they are also causally insufficient. The larger food and sound fields are also part of the causal story. Indeed, as [Salthe \(1985\)](#) has pointed out, a full understanding of hierarchically structured systems ordinarily requires a three-level perspective: the focal level, one level down (the level of mechanism), and one level up (what I am calling the upper-level field). I acknowledge the importance of mechanism in explaining particular instances of teleological behavior but am deliberately de-emphasizing it here in order to make a broader point about teleological systems generally.

5. Human wanting

Beyond simple tropisms, development, and evolution, there is something to be said here about human wanting and behavior. What I have to say is tentative, speculative, but a hierarchically structured mechanism seems to fit the human case so neatly that I cannot pass up the opportunity to speculate.

Human behavior is the teleological phenomenon we are closest to. We think of our behavior as driven by wants, preferences, motivations, and passions, whether conscious or unconscious. And to the extent that wants are conscious, we seem to be aware of them directly, without mediation by the senses and with little or no analysis. I swing the tennis racquet because I want to hit the ball. I plan to buy the car because I want to own it. Our access to our wants seems to be direct, immediate, intimate. Despite this intimacy, they may be—among all the biological entities and all the machines that seem to seek—the least understood. And that combination of obscurity and intimacy has orphaned human wanting, in a sense. We may not know what wanting is, but we think we know that it is something that happens in a mind. We think we know that only minds can truly want anything. If bacteria, evolution, and homing torpedoes behave in ways that mimic wanting, we reason, it must be wanting in a sense that is merely analogous to what minds do. We imagine that only human minds—and perhaps certain other animal minds, we might concede in generous moments—can be truly goal-directed. The goal-directedness of non-thinking systems is only metaphorical, we say. Their teleology, we suppose, is “as-if.”

Let us put aside the fact that any claim about human goal-directedness is, if taken literally, wrong. (Because nothing in

⁴ Notice that the word “error” is being used in a perspective-dependent way here. A deviation from the path toward the goal is an error only from the perspective of the higher-level field. The cell that wanders away from the path that is encouraged by the field has made an error. However, from the perspective of the lower-level entity, the perspective of the cell itself, these so-called errors are expressions of the cell’s freedom, of its independence from upper direction.

nature—not wanting, not anticipating, not planning—is caused by a goal or anything else that is literally in the future.) Instead let us consider the possibility that even the underlying insight is wrong. Consider the possibility that human wanting is not especially special. Suppose the seeming goal-directedness of wanting, arises from the same causal structure as every other kind of seeming goal-directedness: an entity moving or changing within a field, an entity in a delicate balance between freedom and upper direction. Suppose that human wanting has the same causal structure as simple bacterial tropisms, development, and evolution, even the same as apparently goal-directed human artifacts, such as homing torpedoes. The suggestion has at least a superficial plausibility, if only because behavior under the influence of wants shares the trademark persistence and plasticity that characterize other goal-directed systems. I want to go to the post office to mail a letter. As I leave the house, I see the newspaper in the driveway and stop to bring it into the house. As I start to get into the car, I get a phone call, and I pause to talk, starting up the car and continuing on my way when the call is done. *En route* to the post office, I encounter a road that has been closed for construction and I take a detour, eventually arriving at the post office. In other words, I get there despite the diversions. This is persistence. The want, the motivation, is an ever-present drive of some kind, a mental gradient that directs my behavior—my thoughts and physical movements—from above. My thinking and behavior are free to some degree, free to vary under the influence of momentary and idle observations and thoughts, as well as the diversionary wants that they evoke. But to the extent that I am on some mission, under the influence of some single want or some coalition of wants, I am not free. I am directed.

My suggestion here is that wants are upper-level directives, each a higher-level field of some kind (see [McShea, 2012, 2013](#) for further discussion). And whatever entities in the brain generate and control thought and behavior *lie within*—and are to some extent directed by—that field. Since I am speculating here, there is no reason not to broaden the scope of the claim. The suggestion is that all affective phenomena—feelings, passions, wants, preferences, cares, and so on—are upper-level directives, and that all thought and behavior—including speech—is directed from above in this way. This suggestion contradicts nothing that we know about how affective systems are built and how they work. And it is consistent with folk psychology. It is also consistent with certain (admittedly impressionistic) observations about affect. From the perspective of the conscious mind, affective states do seem “large.” Like large things generally, when they change, they change slowly. Affective states change slowly compared to the movement of thought and muscle, compared to the “movement” of the conscious self. Further, like many larger things, they are resistant, sometimes impervious, to efforts to modify them from below, from within. We cannot ordinarily change a want or feeling by an act of conscious will. Nor can we easily summon wants or feelings. We cannot simply decide to want or feel something. Nor can we easily dismiss one that appears unsummoned. To a first approximation, affective states just happen. And finally, consistent with all this, wants and feelings are extra-logical, operating somehow outside of the rules of conscious rational thought. Following a venerable line of argument in philosophy, no want or feeling follows rationally from any set of facts about the world, from any thoughts or observations that we make ([Hume \[1740\] 1978](#)). Wants, feelings, preferences, and cares do follow thoughts and observations in time, of course. And there is a connection between these affective states and conscious thought. But it is not a rational connection. Lightning follows thunder as a matter of physics, not logic. And sympathy follows from the sight of a crying baby as a matter of brain structure, not logic ([McShea, 2013](#)).

Thus, I argue, affective phenomena seem large, because they are in fact large, larger than consciousness. They are not something consciousness is merely influenced by. They are fields in which consciousness is immersed. Affect directs consciousness, and it does so from above.

Some backtracking is needed here. It is true that this viewpoint contradicts nothing we know about how affect works, but it also tells us very little about how wanting is instantiated in brains. It does not tell us anything, for example, about what sort of field is involved. An electric field? A neural activation field? A gene-activation field? I have no idea. Also, academic rigor demands a number of caveats as well as consideration of several possible counterarguments. Sadly I do not have space here to present and reply to them, although I will pause for one. The dimensionality of the causes involved obviously differ among the various systems I have considered here. Human wanting clearly is more complicated in some difficult-to-specify sense than a homing torpedo. The point here is only that despite these obvious big differences, they could nevertheless share a common basic nested structure. And if so, human wanting is an orphan no longer. It has found a natural home within the family of other nested, seemingly goal-directed systems.

It is unclear how this move I have made would affect standard usages of teleological terms. If bacterial tropisms, development, evolutionary adaptation, and human wanting have the same hierarchically nested structure—an entity moving or changing within a field—perhaps we should use the same teleological terms—wanting, seeking, and goal-directed—for all of them. Or maybe we need some new words.

6. Evolution, the teenager, purpose everywhere

Evolution. The discussion here has taken me pretty far from my home territory. Evolution is the only topic on which I can speak with even a pretense of authority. So let me try to say something pretentiously authoritative about evolutionary freedom and purpose, in my sense of those terms. If freedom is change in the absence of external organized forces, in the absence of upper direction, then a free lineage is one that wanders purely under the influence of local circumstances. For a small population, on a timescale of several generations, absence of upper direction is absence of selection, and therefore wandering is drift.

That is microevolution. Macroevolution is different. The reason is that “upper” and “local” are scale-relative terms. For a lineage changing on a timescale of millions or tens of millions of years, “local” forces include all of the short-timescale *selection pressures*, for example, pressures that favor larger or smaller body size, immunity to this or that parasite, slower or faster metabolism, and so on, all changing this way and that on the millions- or tens-of-millions-of-years timescale. In contrast, upper direction refers to selection at a much larger scale. At the large scale, when upper direction is absent, what is absent is not selection per se, but long-timescale selection, selective forces that persist over tens of millions of years. Consider beak size in a bird lineage. In macroevolutionary terms, if there is no long-term selection for, say, large beak size, the resulting trajectory would not be drift, but it would still be a random walk, with beak size sometimes selected for increase and sometimes selected for decrease. Absent upper direction, and subject only to the vagaries of short-timescale selection, the lineage wanders.

That was the evolution of single lineage. What about a group of related lineages, what biologists call a clade? Freedom for a clade means change in each lineage is the product of purely local forces. Each could be drifting. Or each could be under the control of a different set of selection pressures. In either case, the expectation is

divergence. Lineages in the clade should become ever more different from each other with time. And when new species arise, they too should tend to become different. Freedom means diversity. And that is what the zero-force evolutionary law (ZFEL) says: in the absence of large-scale forces acting simultaneously on multiple lineages, diversity is expected to increase.

The same goes for complexity in macroevolution. I need only a single lineage to make this point. When large-scale selection is reduced or absent, the parts in an organism that are free to vary independently will be selected independently, at least to some extent. Thus, a tooth at the back of a tooth row will be subject to different selection pressures than a tooth at the front of the tooth row. And as a result, the teeth will tend to become different each other. Likewise, segments in a worm will tend to differentiate. The wings of a fly will tend to differentiate, left becoming different from right. Even parts that are already quite different initially—the bones of a skull—will tend to become even more different from each other. In microevolution, we think of this independent variation as drift. But in macroevolution, it can also be the result of selection that is “local,” acting on each part independently. In either case, it is the ZFEL, and the consequence is differentiation of parts, an increase in complexity.

In sum, for both diversity and complexity, the fundamental expectation in the absence of large-scale forces is increase. The ZFEL is freedom.

But suppose upper direction is present, in other words, suppose large-scale selection is present. If it is not too strong, not strong enough to overcome the ZFEL, not enough to wipe out freedom, the result for a single lineage is wandering with an overall direction, wandering with a kind of purpose, in other words, adaptation. For multiple lineages, at the macroevolutionary scale, the result may still be growth in diversity, but constrained growth. And if accompanied by adaptation in each lineage, the result is flourishing ecologies, constituted by mutually adapted organisms. At the level of parts, the result might also be constrained increase, producing organisms that are constituted by multiple, different, interacting, mutually adapted parts. (Perhaps this is “complexity,” not just in the technical sense of differentiated parts, but in the sense of adaptation, the sort of “complexity” that we currently have no idea how to operationalize: complex eyes, and brains, and metabolisms, and immune systems, and so on).

Kant said there would never be a Newton for a blade of grass, that science would never explain the teleology embedded in living things. But Kant was at least partly wrong. Darwin was that Newton in that natural selection explains where teleology comes from. It gives a mechanism that explains its origin and maintenance. It explains why there are plants that track the movement of the sun across the sky. It can even explain how a particular species of bacteria evolved to be able to home in on food, and so on.

But origin and maintenance are only one piece of the teleology puzzle. There is also the question of moment-to-moment mechanism. Natural selection does not direct the bacterium while it is swimming, nor development as the organism grows. It is not selection that gives these movements and changes their teleological feel. What does? Nagel’s and Sommerhoff’s answer that persistence and plasticity are crucial seems plausible but we still need to know how *those* are achieved. And then what about natural selection itself? What explains the teleological feel of the Darwinian process, its capacity to persistently and plastically search out and find adaptive design? We cannot invoke selection to explain how selection itself does it.

My proposal is that the answer to the question of mechanism is hierarchical structure, physical nestedness. What explains persistence and plasticity in all of these processes (including natural selection), what gives them their teleological feel, is a delicate balance

between lower-level freedom and upper-level direction. For simple organismal tropisms, freedom is locally determined movements in balance with upper-level, enveloping, directing fields. For development, freedom is locally determined movements of cells, in balance with upper-level, enveloping, directing fields. For evolution, freedom is locally driven changes in phenotype in balance with upper-level, enveloping, directing fields, in other words, the ecology within which the organism evolves. At the macroevolutionary scale, freedom is both the wandering and divergence of lineages and the wandering and differentiation of parts, driven by a combination of drift and lower-level selection acting differently on different individual lineages and also on different individual parts. It is the spontaneous tendency to diversify and complexify that is identified in the ZFEL. Offsetting this lower-level freedom and in delicate balance with it are the macro-scale ecological forces (mainly selection) that drive large-scale trends.

In other words, the explanatory scheme offered here is more general than natural selection, it is one in which adaptation is just a special case. In this scheme, the mysterious apparent purposefulness of adaptation emerges as a feature of nested systems generally, as an instance of a more general teleological-looking process.

The teenager. I return now to the teenager, roaming the world in search of himself, devoting himself in an unfocused way to farming, and then to building a boat, and then to meeting “real people.” This is freedom. This is how people behave when they are responding only to internal whims and local external causes, when they have no greater mission than to respond to the opportunities, exigencies, or urges of the moment. They are carried by the winds, a combination of the winds of their own immediate preferences and the winds of their immediate circumstances. This is what freedom looks like when the delicate balance is broken, when upper direction is absent or much reduced. Freedom is variation. It is wandering.

Upper direction in people takes the form of at least two kinds of “fields.” First, there are social groups: families, friend groups, institutions, companies, clubs, and other purposeful organizations. And second there is individual character, a personal sense of purpose. Our teenager lacks both. He is not immersed in, and does not feel the force of, either of these fields. At least, not yet. That is why his wanderings look so pointless, so purposeless. Moment to moment they are not purposeless. He is doing what he wants. But on a longer timescale, they are. Some people do very well living this way. But for others, some larger sense of purpose is necessary. If our teenager is one of them, we catch ourselves hoping that he finds some upper direction. If he cannot find it within himself, we hope he finds it out there, in the world of purposeful groups.⁵

Purpose everywhere? The argument developed here is an attempt to naturalize purpose. It reveals goal-directedness, teleology, purpose to be a property of a particular physical organization, entities moving or changing within a higher-level field that partly directs them, entities in a delicate balance between freedom

⁵ In all of the earlier examples, the upper-level directing field is external to the goal-directed entity. The food gradient is external to the bacterium. Ecology is external to the lineage. In the human case, however, the upper-level fields corresponding to deeply embedded character elements—those corresponding to long-term purposes, when they exist—are internal. They exist within a mind, not out in the world, not external to the individual. For the view offered here, this is not troublesome. They may nevertheless be larger than and envelope the moment-to-moment wants that drive free thought and behavior. Whatever systems produce our free wanting, thinking, and behaving must be immersed in these fields, and directed by them to some extent, from above. In other words, the entire hierarchical structure that underlies goal-directed behavior can exist entirely within a mind. (Still, as already pointed out, some of the deepest and more powerful elements of individual character arise during a lifetime from immersion in systems that are in fact external, especially social groups: families, institutions, etc.)

and upper directedness. Now one might accept the basic argument but decline to apply these terms—teleology, purpose, goal-directedness—unless I concede that they are being used metaphorically. Indeed, I appreciate the discomfort associated with calling bacterial movement or evolutionary change teleological or purposeful. Personally, I am happy to use these terms non-metaphorically (with the no-backward-causation caveat discussed earlier) for any entity that has found the delicate balance between freedom and upper direction. But of course, choice of words here is a matter of taste.

Whatever words we use, the understanding developed here raises the possibility that goal-directedness/teleology/purpose could be rather common in the universe. I have talked mainly about biological systems, but there is no reason to expect it only or even mainly in living things. We should expect to see it anywhere that freedom and upper direction have found that delicate balance. That includes certain human artifacts, such as the homing torpedoes discussed earlier, but also thermostats and other regulative devices. Some degree of it can perhaps be found in simple chemical systems. It is probably also present in certain large-scale physical systems, perhaps local weather cells entrained by large-scale atmospheric structures, local storms moving within hurricanes. Some degree of it is possible at the scale of subatomic entities moving within electric fields and also at the scale of galaxies moving within the gravitational fields of supergalactic clusters. And beyond?

References

- 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.
- Doolittle, W. F. (2012). A ratchet for protein complexity. *Nature*, 481, 270–271.
- Ettensohn, C. A. (1990). The regulation of primary mesenchyme cell patterning. *Developmental Biology*, 140, 261–271.
- Ettensohn, C. A., & McClay, D. R. (1988). The regulation of primary mesenchyme cell migration in the sea urchin embryo: Transplantations of cells and latex beads. *Developmental Biology*, 117, 380–391.
- Fleming, L., & McShea, D. W. (2013). *Drosophila* mutants suggest a strong drive toward complexity in evolution. *Development and Evolution*, 15, 53–62.
- Hume, D. (1740[1978]). In L. A. Selby-Bigge (Ed.), *A treatise of human nature* (2nd ed.). Oxford: Oxford University Press.
- MacDonald, G., & Papineau, D. (2006). *Teleosemantics*. Oxford: Oxford University Press.
- McShea, D. W. (1996). Metazoan complexity and evolution: Is there a trend? *Evolution*, 50, 477–492.
- McShea, D. W. (2012). “Upper-directed systems: A new approach to teleology in biology. *Biology and Philosophy*, 27, 663–684.
- McShea, D. W. (2013). Machine wanting. *Studies in the History and Philosophy of Biological and Biomedical Sciences*, 44, 679–687.
- McShea, D. W., & Brandon, R. N. (2010). *Biology's first law*. Chicago: University of Chicago Press.
- Millikan, R. G. (1984). *Language, thought, and other biological categories*. Cambridge: The MIT Press.
- Nagel, E. (1979). *Teleology revisited and other essays in the philosophy and history of science*. New York: Columbia University Press.
- Neander, K. (1991). The teleological notion of ‘Function’. *Australasian Journal of Philosophy*, 69, 454–468.
- Salthe, S. N. (1985). *Evolving hierarchical systems*. New York: Columbia University Press.
- Salthe, S. N. (2009). A hierarchical framework for levels of reality: Understanding through representation. *Axiomathes*, 19, 87–99.
- Simon, H. A. (1962). The architecture of complexity. *Proceedings of the American Philosophical Society*, 106, 467–482.
- Sommerhoff, G. (1950). *Analytical biology*. London: Oxford University Press.
- Valentine, J. W., Collins, A. G., & Meyer, C. P. (1994). Morphological complexity increase in metazoans. *Paleobiology*, 20, 131–142.
- Wimsatt, W. C. (1974). Complexity and organization. In K. F. Schaffner, & R. S. Cohen (Eds.), *Philosophy of science association 1972* (pp. 67–86). Dordrecht: D. Reidel.
- Wimsatt, W. C. 1994. “The ontology of complex systems: Levels of organization, perspectives, and causal thickets.” In “Biology and Society: Reflections on Methodology,” supplement, *Canadian Journal of Philosophy* 20 (S1):S207–S274.
- Wright, L. (1973). Functions. *Philosophical Review*, 82, 139–168.