Biol Philos (2012) 27:663–684
DOI 10.1007/s10539-012-9326-2

# Upper-directed systems: a new approach to teleology in biology

**Daniel W. McShea**

Received: 23 November 2011 / Accepted: 13 May 2012 / Published online: 6 June 2012


**Abstract** How shall we understand apparently teleological systems? What explains their persistence (returning to past trajectories following errors) and their plasticity (finding the same trajectory from different starting points)? Here I argue that all seemingly goal-directed systems—e.g., a food-seeking organism, human-made devices like thermostats and torpedoes, biological development, human goal seeking, and the evolutionary process itself—share a common organization. Specifically, they consist of an entity that moves within a larger containing structure, one that directs its behavior in a general way without precisely determining it. If so, then teleology lies within the domain of the theory of compositional hierarchies.

**Keywords** Teleology · Goal-directedness · Hierarchy · Purpose · Evolution

A peanut butter sandwich falls into a pond, and the soluble molecular ingredients begin to leach into the water. The result is a concentration gradient, a kind of "field" in which the concentration of a substance, say, the amino acid aspartate, is highest near the sandwich and declines with distance from it. A bacterium that is attracted to aspartate finds itself within this food field. Most objects in the water–dust particles and such—are not affected by the food field, but the effect on the bacterium is profound and highly organized. Normally, when a food gradient is absent, the bacterium's flagellum propels it in a random walk consisting of a series of straight runs interrupted by tumbles that randomly reorient it. But now, enveloped by the aspartate gradient, the increase in concentration when the bacterium happens to move up the gradient induces the flagellar unit to increase the length of straight runs. And when the bacterium's movements happen to carry it down the gradient, the

D. W. McShea (✉)
Biology Department, Duke University, Box 90338, Durham, NC 27708, USA
e-mail: dmcshea@duke.edu

decrease in concentration induces a decrease in the length of straight runs. The result is that, despite many poorly directed straight runs and many random tumbles, the gradient induces in the bacterium a net tendency to move toward regions of higher aspartate concentration.

This story is about a paradigmatic "teleological" behavior, a living organism pursuing food. But I have described it in an unusual way. Ordinarily, in biology, the focus would be on the bacterium's ability to sense its environment, the chemosensory apparatus; on the mechanism of the propulsion unit, the flagellum; and on the signaling pathway that connects them. The bacterium might be said to use information from its environment to direct itself toward higher food concentration. It would be described in the language of active agents, while the gradient would be treated as passive. The bacterium assesses the gradient, propels itself up the gradient, reorients with respect to it, and so on.

But instead I presented the bacterium as just an entity with certain response properties. It was described as nested within an enveloping field, a structure that was larger than it, that contained it, and that acted causally on it, via some unspecified mechanism, directing its behavior. The enveloping structure was the active causal agent. Also, the perspective was hierarchical, and in particular it was a top-down perspective, the view of the entity from the vantage point of the larger structure that directs it.

Why take this perspective? A central question in the literature on teleology has been, what do such systems have in common? What are the key properties shared by behaviors like the bacterium's, and the behaviors of machines like thermostats and homing torpedoes, by physiological feedback systems like those regulating blood pressure and body temperature, and by developmental mechanisms that transform embryos into adults? What gives these systems their teleological flavor?

Various proposals have been made (to be discussed shortly). But here I propose a new starting point, a hierarchical perspective. Specifically, I introduce a class of hierarchical systems with three properties. Then I argue that classic seemingly teleological entities—non-human organisms, certain machines, developmental systems, etc.—are a subclass and that the three properties make sense of two key features of these entities, their persistence and plasticity. I've taken these two terms from Nagel (1979, see also Sommerhoff 1950), but I use them here in a slightly different way. Here, persistence is the tendency for an entity that is following a particular pattern of behavior, a behavioral trajectory, to return to that same trajectory following perturbations that cause it to depart from the trajectory. And plasticity is the tendency for an entity to find a particular trajectory from a variety of different starting points. The bacterium *persists* in a trajectory toward higher aspartate concentrations, despite setbacks. And its trajectory is quite *plastic*, in that it will tend to move toward higher concentration from any of a number of different starting points.

The central claim of the hierarchical view is this: *A critical part of what gives apparently teleological entities their persistence and plasticity is the fact they move or change within a larger containing structure, one that directs their behavior in a general way without precisely determining it. In other words, the claim is that apparently teleological entities are what I call "upper directed."* If this claim is right, it follows that any theory of teleology properly lies within the larger domain of the theory of compositional hierarchies, in the sense of Simon (1962), Campbell (1958), Wimsatt (1974, 1994), Salthe (1985), Valentine and May (1996), and McShea (2001, also McShea and Changizi 2003). Further, the claim is that systems with persistence and plasticity are a subset of upper-directed ones, and among these persistent and plastic systems, it is the more complex ones that we are inclined to call teleological.[1]

The hierarchical perspective has the virtue that it provides a unified account of apparently teleological behavior. In the literature on biological teleology, systems like the bacterium have been treated together with other non-conscious systems like organismal development and apparently goal-directed machines, such as thermostats and torpedoes. But two other sorts of apparently teleological system have seemed different and have been set aside. One is conscious human behavior, which while it can lead to the design of thermostats and torpedoes, does not work like these machines, nor like the bacterium. The other is natural selection, which in certain lights seems to "seek" adaptation but does so via a variation-and-selective retention mechanism that is different from seeking in humans, non-conscious biological systems, and machines. The hierarchical perspective draws these two orphaned systems into the same family with motile bacteria and thermostats, revealing how they can all be understood using the same explanatory logic.

Finally, the hierarchical perspective gives us a way to talk about persistence and plasticity in standard causal terms. There has always been an aura of mystery, of magic, around such systems on account of their seeming future directedness. The three standard terms of discourse (teleology, goal-directedness, purpose) all imply a future object or event (a telos, a goal, an achieved purpose) that is in some sense explanatory of present behavior. We say that the bacterium swims up gradient in order to reach a goal, a food source, as though an object that is only potentially present in the bacterium's future could be the cause of its present behavior. We talk this way even though future causation is impossible. Here I mostly avoid the three standard terms, and when I do use them (to make contact with the existing literature), I modify them with a word like "apparently." This is not an eliminativist position. The phenomena—persistence and plasticity—really do occur. But virtually everyone agrees that no system is literally directed by the future, not even intentional behavior in humans, and I take this fact to heart in my choice of words. As it happens, the hierarchical view makes it easy to avoid future-causation talk, providing a purely present-causal account of persistence and plasticity.

In sum, the hierarchical view offers three things: (1) a present-causal account of persistence and plasticity, an explanation of how these entities do it, one that penetrates the mystery and lightens the temptation to talk in future-causal terms; (2) an account of apparently teleological behavior that unifies all instances of it in biology and technology—unconscious-biological, mechanical, developmental, conscious-psychological, and evolutionary—in a common explanatory scheme;

[1] The present discussion is agnostic about what has been called "downward causation," more specifically about whether higher-level explanations are reducible, either in principle or in fact, to lower-level ones. I coined a new term, upper directedness, in part to distance this discussion from that one.

(3) an account that identifies a scientific context, compositional-hierarchy theory, within which the study of teleology lies, making it more accessible to empirical study.

*Three alerts*. First, this treatment includes no analysis of teleological language. It should be seen as a mainly scientific treatment (with implications for the philosophical discussion), leaning heavily as it does on certain empirical facts, mainly generalizations about the properties of hierarchical systems. Second, the claim that apparently teleological systems are hierarchical rests on my being able to show that such systems in fact have the properties shared by a special class of hierarchical systems. In others words, this is a consistency argument. Nothing like a proof will be offered. Third, the focus here, at least initially, is the present persistent and plastic behavior of entities and the *proximate* causes of those behaviors. This focus has two consequences. First, following Nagel (1979), the causal deep-history of a goal-directed system is not relevant. The proximate mechanisms that govern it may have originated by natural selection—indeed it may be that those mechanisms are adaptive and could *only* have originated by natural selection—but a proximate explanation of how they work in the present does not depend on how they arose, and need not involve selection (except of course for adaptation itself, considered as a teleological process, discussed later). Second, the discussion will not engage the philosophical debate about function (although it has consequences for that debate). A consequence is that structures that are merely functional, but that do not show persistence and plasticity, such as the vertebrate retina, will not initially be of interest. Retinas are the product of a historical process that is persistent and plastic, natural selection, but their present behavior—transducing light into nerve impulses in a mouse in the here and now—is not directed by that historical process from moment to moment.

## Nagel, Mayr, and others

The argument here has a number of connections with past treatments of teleology. In particular, it agrees with and builds on Nagel (1979), Mayr (1992) and others (e.g., Rosenblueth et al. 1943) in taking a systems-theoretic approach to teleology, proposing a common causal structure for systems that seem goal directed. The closest connection is with Nagel, in particular in adopting his suggestion that persistence and plasticity are the key features of seemingly teleological systems, the features that need to be explained.

However, it departs from Nagel in two ways. First, it adopts a view of persistence and plasticity that involves no concept of a goal. In Nagel's understanding, persistence is the tendency of a system to reach a goal state G despite perturbations, and plasticity is its tendency to reach G from any of a number of starting points. Mayr and others, including Schlosser (1998) and Trestman (2010) understand apparently teleological behavior in a similar way, as a relationship between present behavior and a future goal. In the present view, however, persistence is just a tendency to return to a past trajectory. And the explanation for persistence includes no reference to future states, no reference to a G. Second, as will be discussed later, the hierarchical view accepts, along with Nagel and also Sommerhoff (1950), the need for at least two state variables that are in principle independent of each other. The hierarchical view agrees that two independent variables are necessary but adds that they must be hierarchically related. Thus the perspective offered here builds on Nagel's view but claims that his understanding of the physical setup of teleological systems is incomplete.

Mayr (1992) describes certain apparently teleological systems as merely "teleomatic," meaning that they are guided by simple natural laws, like a stone rolling downhill. But the more interesting ones are "teleonomic," meaning that they are guided by an internal program into which the goal of the teleomatic entity has been coded, as in a homing torpedo, a migrating bird, or a developing embryo. In machines, the coding of the goal into the program has been engineered by the designers. In biological systems, it has been coded into the genome by natural selection. The present view differs from Mayr's in two ways. First, as discussed, it invokes no concept of a goal, either coded into a program or otherwise. Second, in the search for causes, it does not look only to the internal structure of the entity. That internal structure could be causally important, but it need not be a program of any kind, and even if it is, it cannot be causally sufficient. In the present view, the causes of persistence and plasticity could be located mainly—and must be located at least partly—in an external field.

Certain elements of the hierarchical view can be found elsewhere in the literature. For example, the hierarchical view is consistent with Braithwaite's (1953) view of teleology insofar as both lean on external factors, what Braithwaite calls "field conditions," rather than goals, to explain seeming teleological behavior.

Finally, it should be said that a fair chunk of the teleology literature is not directly relevant here, the portion concerned with the conditions necessary for the appropriate use of teleological terms, and with how we use and understand those terms. The focus here is on *behavior*, and the aim is to explain the behavior of apparently teleological entities, to explain their persistence and plasticity.

## Hierarchy

Here, hierarchy is just containment, nestedness. A hierarchical system is a larger structure that is composed of smaller entities nested within it. Both structure and entity could be solid objects, as in Russian dolls or Chinese boxes. But hierarchy does not require solidity. In the bacterium story, the upper structure is a field (food particles plus the water around them), not an object. Also, the containment need not be spatial. To start, I will stick with cases in which the upper structure is physical and containment *is* spatial. The bacterium is a physical entity that is spatially nested within a physical food-water structure. But later I will invoke a notion of hierarchy in which the containing upper structure is a "phase space" within which the entity moves.

Unfortunately for present purposes, the word hierarchy has other senses, which could lead to confusion. In particular, the hierarchical relationship of interest here is not a control hierarchy, in which one event causes one or more other events later in time, as when an officer in an army unit gives orders to his subordinates or when one stage in the development of an embryo gives rise to a later stage. In neither of these is there (necessarily) any containment, physical or otherwise.

## A special class of hierarchical systems

I am interested here in hierarchical systems with certain properties. Not all hierarchical systems have these properties, but they are common among such systems, especially in biology. I have extracted these properties from the modest literature on hierarchy theory in biology, mainly Simon (1962), Campbell (1958), Salthe (1985), Wimsatt (1974, 1994), Allen and Starr (1982), and Ahl and Allen (1996). This section describes these properties using non-teleological examples. The next shows how they apply to teleological systems.

*Upper directedness* (*property* 1). Upper direction is any effect of a larger structure on a smaller entity that it contains. Consider an atom of helium in a neutrally buoyant helium-filled balloon. Suppose that we number the gas atoms in the balloon, and our chosen atom is number 42. Viewed hierarchically, atom 42 is an entity contained within an enveloping structure consisting of the plastic of the balloon itself and also the cubic foot or so of gas (i.e., all of the other helium atoms) within the balloon. Now suppose that someone grabs the balloon, moves it, and releases it in a new location, say, centered on a point two inches to the east of its original location. During the time it takes for this two-inch shift to occur, atom 42 follows a zig-zag path within the balloon, changing direction with each collision with another helium atom and also perhaps with the plastic balloon membrane. Given the complexity of the situation, we cannot predict exactly where atom 42 will end up, but we can say that, on average (over many repeats of the shift of the balloon), it will move two inches to the east. Further, we can say that the proximate cause of atom 42's two-inch, on-average shift was the two-inch shift of the balloon that contains it. Or, decomposing the atom's movement into a 2-inches-to-the-east component and all of its other movements, we can say that the 2-inches-to-the-east component was caused by the movement of the upper structure, the balloon as a whole, including the other gas molecules within it. That component of atom 42's movement is upper directed.

Upper directedness is a common feature of entities contained within larger structures. Indeed, on account of containment, it may be difficult for a contained entity to escape the effects of large-scale movements of the containing structure. As the flow rate of a river increases, the average rate of movement of each of its contained water molecules also increases. A change in the heading of a ship affects the average heading of each rat within its holds. The moon affects the average movement of the Earth's oceans, an ocean being an entity contained within the gravitational field of the moon. Still, this is not a property of all hierarchical systems. A shift in the Earth's magnetic field will affect a compass needle but not a feather.

The non-hierarchical alternative to upper directedness is "lateral directedness." A laterally directed entity is one whose behavior is caused by another entity that is about the same size and that does not contain it. The collisions of balls on billiard table are instances of lateral directedness. The turning of the gears in a car transmission is lateral directedness. The booting up of Windows in a laptop computer is lateral directedness. Notice that every case of upper directedness is mediated by a number of cases of lateral directedness, as the on-average 2-inch shift of atom 42 is mediated by the laterally directed impacts of other helium atoms and plastic molecules. Further, it may be that every case of lateral directedness can be construed, with a shift in scale, as an instance of upper directedness, as the movement of a gear of a clock can be described as the effect of an upper structure, the clock as a whole, on that gear. For present purposes, it does not matter that different perspectives yield different causal stories. As we shall see, what matters is that teleological-looking behavior never emerges from a purely lateral-causation perspective. Billiard balls do not seem to behave teleologically. (Or at least, if they do, it is always with respect to some larger perspective, for example, one that includes the pool player, the cue stick, the pockets, and the table.)[2]

*Partial independence of lower entities* (*property* 2). The components of solid objects all show upper directedness. The deck of a ship moves in lockstep with the ship as a whole. A current that drags the ship off course also drags the deck. In solid objects, the correlation between the movements of the whole structure and a contained entity is 1.0, or very close to it. For present purposes, however, entities that are components of solid objects are not very interesting. So the special class of hierarchical systems of interest here includes only those in which upper directedness is intermediate, in which the correlation between structure and contained entity lies somewhere between 0 and 1.0. In such systems, the contained entity retains some freedom of movement. The upper structure affects the behavior of the contained entity but does not precisely determine it. Atom 42 receives upper direction from the balloon as a whole but it also moves to some degree independently. The rats can run about within the ship.

*Stability of the upper structure* (*property* 3). The hierarchical systems of interest here are those in which the upper structure is stable, or changes slowly relative to the movements of the contained entity. We want ships that turn slowly compared to the independent movements of the rats in the hold, and balloons that are moved slowly compared to the movement of their contained gas particles.[3] Notice that it is helpful to stability in these systems if individual entities are unable to affect the upper structure much. If our bacterium in a food field could feed at a rate that substantially reduces the concentration, then its feeding changes the food field, and the changing structure of the food field in turn affects the movements of the

[2] Notice that the way we may best understand the causation may vary with perspectival scale. If we took the whole balloon to be the focal entity, rather than atom 42, and considered the collision between two balloons, the cause of movement in each balloon would be lateral, like that between billiard balls. At the same time, and without contradiction, the movement of each atom within the balloons would be partly upper directed.

[3] Importantly, the relationship of interest here is only that between structure and contained entity, not movement in some larger frame of reference. It does not matter, for example, that both balloon and atom 42 are moving at 65,000 miles per hour around the sun. What matters is that from the perspective of atom 42, the balloon as a whole is moving quite slowly.

bacterium, which in turn affects the food field, etc. It is just such systems that need to be excluded. Our special class of hierarchical systems does not include ships whose movements are affected much by the movement of a single contained rat. Many hierarchical systems achieve this insulation with a size differential. Single gas molecules are too small to affect the movement of a normal-sized balloon. But there are other ways.

## Apparent teleology and the three properties

So far, there has been no teleology. All that the last section did was identify a class of hierarchical systems with three properties. Now, to make the connection with teleology, I will argue that apparently teleological systems in fact have the three properties and that the three properties are what make sense of persistence and plasticity. I do this first for the bacterium, in this section, and then for other non-conscious systems in the following two sections.

What is the proximate cause of the bacterium's movement up the gradient? One seeming possibility is the sandwich. But notice the sandwich is an optional part of this story. The bacterium would move in the same way in a food field that has the same structure but is somehow maintained without any sandwich. Another plausible cause is the evolutionary history that produced the bacterium's internal mechanisms. But this is not a very proximate cause, removed as it is in time. Or one might point to internal mechanisms within the bacterium, the bacterium's signal transduction pathways. And this is surely right, but it is only part of the story. These mechanisms are causally necessary, but they are not sufficient. They do not decide whether or not the bacterium will move in a persistent and plastic way. Bacteria that are not presently immersed in a food gradient have these same mechanisms and do not behave persistently or plastically. So what is left? The only remaining causal factor is the enveloping food field within which the bacterium moves. In other words, the cause of the bacterium's seemingly goal-directed behavior is the upper direction provided by the field (property 1). Given this, the second property follows. Despite the causal influence of the field, the bacterium retains considerable freedom of movement (property 2). For example, the changes of direction during random tumbles are independent of the field and therefore free in this way. Also the bacterium's various movements resulting from buffeting by local microcurrents or collisions with small particles count as free movement. And finally, by assumption, the food field is relatively stable, at least relative to the movement of the bacterium (property 3). The field is not fluctuating rapidly.

How do these three properties give a teleological look to the bacterium's behavior? How do they give rise to persistence and plasticity?

*Upper directedness* (*property* 1). Persistence and plasticity would not be possible if the bacterium's behavior were not at least partly upper directed (property 1). What persistence and plasticity require is that the cause of the entity's behavior be present over a large area, so that most errors and perturbations leave the entity still in a place where it can be directed once again by that cause. What hierarchy delivers, what the presence of an enveloping causal structure confers, is the possibility of receiving the appropriate causal influence from a wide range of locations where errors and perturbations could leave it (persistence) and from a wide range of starting points (plasticity).

In contrast, consider lateral direction. A laterally directed entity is one whose behavior is caused by another that is about the same size and that does not contain it. Can laterally directed entities appear teleological? A collision between two oxygen molecules does not seem teleological. On the other hand, the behavior of some laterally directed entities is quite complex. Imagine a robot bacterium that behaves as follows: swims straight for 2 mm, turning 90 degrees and proceeds for 1 mm in a straight line. And suppose it did this over and over again, executing a series of "knight moves." This behavior may be persistent in some sense, but not in a teleological-looking way. "Relentless" might be a better word. The robot executes the same moves but does not return to the same trajectory. Some behaviors in organisms and even in a few machines seem to be governed this way. Long-chain programmed behaviors in insects seem to work like this. A child's Lego "Mindstorms" robot moves in this preprogrammed, laterally-directed way. Neurons trigger contraction in muscles and muscle groups in this way.[4] None of these look teleological.

Seeming counter-examples come to mind. It might seem that there is some teleology associated with schooling in fish and that the behavior is governed only by lateral causation, by interactions among individual fish, without any direction from a larger containing structure. However, we need to be clear about just what the teleological entity and behavior are. The behavior of the school as a whole is not what seems teleological. The school twists and turns in aquatic space in ways that are wonderful to watch, and may even be adaptive, but it is not the school that shows persistence and plasticity. Rather, it is the behavior of individual fish, each sticking with the school no matter how it twists and turns. And each fish presumably does so using cues from the visual field set up by the rest of the school, or even more narrowly by the small number of other fish that are near it. It is that set of fish that constitutes the containing structure giving upper direction to the individual.

It might seem that pure lateral direction is still possible for a seemingly teleological entity, if say the entity has programmed into it the entire space of possible locations and an exhaustively complete set of local rules pointing it in the direction appropriate to the appearance of goal directedness. Still, in such systems, while the entity does rely heavily on local rules, inevitably it must also consult an external field, if only to find its location in space. Consider the GPS systems in cars, in which a satellite gives the unit its location in space enabling it to compute the shortest route to a preset destination. The system relies extensive on an internal and preprogrammed set of local rules, but it also necessarily relies on a larger field emanating from the satellite.

Importantly, the field need not be a simple one. A robot with an internal map of a room and an appropriate sensor might orient itself by identifying key landmarks in

[4] It may be that in organisms, unlike machine, pure lateral causation is rare in that most parts are to some degree activated and directed by enveloping fields of various kinds. Consistent with the present view, I am tempted to say that this is what makes organisms so teleology rich and machines so teleology poor.

the room, and then once oriented, propel itself in a persistent and plastic way to any other designated location in the room. One might be tempted to say the interaction between the robot and any given landmark is a case of lateral causation, each landmark interacting in a unique way with the robot. However, the robot's persistence and plasticity, its ability to start from anywhere and recover from errors, requires the entire set of landmarks, the entire spatially distributed set of landmarks, the enveloping landmark "field." The field in this case is not simple. All of these landmarks might be different from each other, and their distribution in space might follow no simple pattern. But together, as an ensemble, they are nevertheless a source of upper direction for the robot.[5]

Thus, again, the teleological look requires upper direction. For an entity to persist, in the face of error and random variation, its source of direction must lie in some structure that is larger than it, that extends to portions of space into which error and random variation could take it. And this is so whether it is moving persistently toward regions of higher food concentration, or toward greater light intensity, or toward the north magnetic pole, and so on.

Importantly, upper direction by itself is not sufficient to produce apparently teleological behavior. The structure of the contained entity is also causally important, obviously. An aspartate field does not provide any upper direction to a floating pine needle. It can only provide upper direction to an entity with the appropriate structure, e.g., a flagellum, signal transduction mechanisms, etc. Thus a complete causal story will generally include the structure and dynamics of the teleological entity as well as the upper direction provided by the containing structure. The claim here is only that upper direction is a necessary component of that story.

*Independence* (*property* 2). Partial independence of the lower entity (property 2) is also part of what makes persistence and plasticity possible. Consider the alternative, an entity whose behavior is completely determined by the larger structure. Suppose the bacterium finds and attaches itself to the sandwich. In effect, the bacterium plus sandwich has become a single solid unit, so that independent behavior on the part of bacterium disappears. And so does our inclination to call the bacterium's movements teleological. Partial independence is critical. It gives the entity the capacity to make errors, or to be deflected in arbitrary ways, errors or deflections to which it responds with upper-directed corrections. And it is the

[5] It might seem that certain coordinated behaviors in organisms—perhaps certain courtship displays—are also counterexamples, in that they appear to be both laterally directed and teleological. They seem laterally directed in that each individual appears to be programmed with an exhaustive set of responses to the likely behavior of a conspecific, with no upper structure providing direction. If that is right, the question then arises whether they are teleological. If the two organisms are thinking and motivated animals, then there is certainly a teleological component to their behavior *as individuals*. (See discussion of human motivation below.) But the issue here is whether their *joint* behavior, the behavior of the pair considered as a unit, meets the requirements for teleology, whether their joint behavior is persistent and plastic. A pair of interactors could be mutually regulating while the pair as a unit moves undirected in the phase space (like a pair of orbiting masses, each regulating the other, while the center of gravity of the pair drifts through space).There is also the issue of whether the trajectory of the pair is plastic. Courtship displays in particular typically have rather strict limits on starting conditions. In sum, for purely laterally caused coordination among individuals, the existence of persistence and plasticity is—I think—an open question.

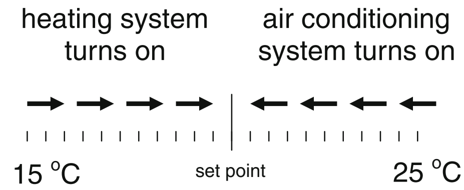


**Fig. 1** Phase space for air temperature in a thermostat-controlled house, showing the vector field controlling temperature (See text)

constant error-or-deflection followed by correction that gives teleological entities their signature persistent behavior.

*Stability* (*property* 3). Finally, the relative stability of the upper structure provides the entity with direction that is reliably present and consistent over multiple attempts. The bacterium could not navigate up gradient if, say, the food source were jiggling violently, creating a field that fluctuated wildly, or if the movements or physiology of the bacterium itself were able to alter the field significantly. On the other hand, if the upper structure changes slowly, the entity can treat it as constant, and use the structure's reliable features to persist and to respond persistently and plastically. Upper-structure stability makes the entity's persistence and plasticity possible.[6]

## A more-general phase space

The view that has been developed so far is incomplete in that it does not account for the seeming teleological behavior of entities that move in "non-spatial spaces." The air temperature in a thermostat-controlled house, for example, has a trajectory, but it does not move in physical space. Rather, it moves in a temperature phase space (Fig. 1). When the air temperature falls below the set point, the heating system comes on, raising the temperature. When the house temperature rises above the set point, the heat turns off and the air conditioning system turns on, lowering the temperature.

Where is the hierarchical structure in this setup? There is a spatial hierarchy, in which the lower-level entity is the air inside the house, and the containing structure is the rest of the house, including the walls and windows, and also the thermostat and the heating and air conditioning systems. But the hierarchy that is directly relevant is that within the phase space, where the enveloping field is the set of directional tendencies shown by the arrows in Fig. 1, and the contained entity is a point representing the air temperature. What makes the system seem teleological is

[6] In some seemingly teleological systems, one of the contained entities may be a control system that directs the containing structure, apparently violating the stability requirement. For example, captains direct ships. How we are to understand this in hierarchical terms depends on which entity we treat as teleological. If the teleological entity is the captain, then for him the ship is a stable upper structure which gives persistence and plasticity to his movements within it, just as it does for a rat. If the teleological entity is the whole ship, including the captain, showing persistence and plasticity in its trajectory across the ocean, then the magnetic field of the Earth or the field of the satellite GPS that the captain uses to steer the ship is the upper structure. And it too is stable. The situation is similar for a cell with its contained DNA, functioning to some degree as a captain-like control system.

that the trajectory of that point is, first of all, plastic. From any starting point on the far left, the enveloping vector field directs the air temperature to the right, and from any starting point on the far right, it directs the point to the left. And of course, it is also persistent. On whichever of the two trajectories air temperature find itself, toward lower or toward higher temperatures, it returns to that trajectory following perturbations.[7]

All three properties of hierarchy are in play here. The reason the entity shows persistence and plasticity is that it lies within a larger causal field (Fig. 1). In other words, it is upper directed (property 1). Of course, upper direction would not be necessary if the air temperature were not free to vary somewhat, independent of the house and its heating and cooling system (property 2). Starting from any point in the temperature phase space, various factors can either raise or lower the temperature, for example, a weather-related rise or fall in the air temperature outside the house. And finally, stability (property 3) is crucial. The assumption here is that the house-heating–cooling-thermostat system is stable, and therefore that the enveloping field in Fig. 1 is stable, on a timescale that is long relative to that on which perturbations affect the interior air temperature. This would not be the case if, for example, a child were changing the thermostat setting every few minutes. In that case, the house-heating–cooling-thermostat system would continue to operate, of course, but the field would be grossly unstable, and the behavior of the air temperature would no longer seem teleological.

## Some classic examples considered

The hierarchical perspective accounts for persistent and plastic behavior in all seemingly teleological systems. Consider the following examples:

1. A sound-seeking torpedo, an entity, is launched by a submarine into an underwater sound field, a containing structure generated by a moving target ship. The strength and orientation of the sound field control the movement of the torpedo's rudder (property 1). When errors in sound-field detection and random local currents deflect the torpedo into areas of lower sound intensity (property 2), it corrects itself and returns to a heading up the sound gradient. Unlike the peanut butter sandwich in the bacterium example, the ship is moving so the sound field is not actually stable in space. However, if the torpedo is moving fast enough—much faster than the sound field—the field is stable in a relative sense. Fast-moving objects can treat slow-moving objects as stationary. And if the sound field is essentially stationary relative to the torpedo (property 3), it is continuously able to direct the torpedo's up-gradient trajectory. To the extent that the sound field is not stable in this relative sense—to the extent that it changes on timescales that are short relative to the reaction time of torpedo, perhaps due to sounds from passing pods of

[7] The spatial hierarchy—air within the house, including its windows, heating/cooling systems, thermostat, etc.—is not irrelevant here. It is no accident that the physical structure that makes the state-space structure possible is also hierarchical. As it turns out, in non-spatial phase spaces, physical hierarchical structure is an easy way to achieve upper directedness, although not the only way.

whales—the torpedo's behavior will appear confused and it will not seem to behave teleologically.

2. Conscious organisms present a more complex case. A rabbit pursued through a thicket by a fox is controlled by at least two fields, the sound field generated by the fox's movements and a kind of imagery field, the visual glimpses the rabbit catches of the fox as it zigzags. Quite likely, what the rabbit does is integrate the two fields in an attempt to estimate—on the run—the present location of the fox and to project its future location. The rabbit's brain processes are complex, but from the hierarchical perspective, it is just an entity, one that is located within a fox-generated sound and imagery field, and the effect of this containing structure is to cause the rabbit to behave in ways that tend to take it away from the fox.

What about the three properties of hierarchy? As for the bacterium, the rabbit's apparently teleological behavior requires that the controlling field be much larger than itself, so that it can respond appropriately over a wide range of locations (property 1). And its behavior has a teleological flavor just because it is able to do this despite occasional errors in judgment and random wrong turns (property 2). What about property 3? It would seem to be absent, since the directing field emitted by the fox is moving quickly. And indeed, when property 3 is absent, the rabbit's behavior will appear confused and not very teleological-looking. The rabbit will sometimes run toward the fox, rather than away from it. On the other hand, real rabbits are quite smart, so that property 3 is not entirely absent. The rabbit might know that at this moment the fox is to the left and be able to project that in a few moments it will be to the right. It might hypothesize a kind of four-dimensional map of the sound-and-imagery field, where the fourth dimension is time. In other words, by making predictions, the field can be effectively stabilized and the rabbit's behavior can seem teleological.

3. An organ system that includes the kidneys and the blood vessels provides one of several controls on blood pressure. The containing structure is the whole organ system, the teleological entity is the blood itself, and the relevant variable is its pressure, which changes in a one-dimensional pressure phase space. When the pressure falls too low, special cells in the kidneys detect the drop, triggering the secretion of renin into the bloodstream. Renin causes increased production of angiotensin, which in turn causes the blood vessels to constrict, raising blood pressure. Renin also stimulates secretion of a hormone that causes the kidneys to retain sodium and water. And increased sodium and water in the circulatory system also raises blood pressure. Excessive blood pressure has the reverse effect. Here, a property of the entity, the blood, is controlled from above by a larger structure, the organ system consisting of the kidneys plus the blood vessels of the circulatory system (property 1). Blood pressure is free to wander somewhat (property 2), as when the person relaxes (lowering blood pressure) or when he or she consumes salt (raising it). But the organization of forces in the phase space is such that the there is, at all points in the phase space, a tendency to move toward lower pressure when the starting point is high, and toward higher when the start is low. Finally, the containing structure and the resulting organization of forces in the phase space are quite stable, so that the field does not change much, at least on short and moderate time scales (property 3).

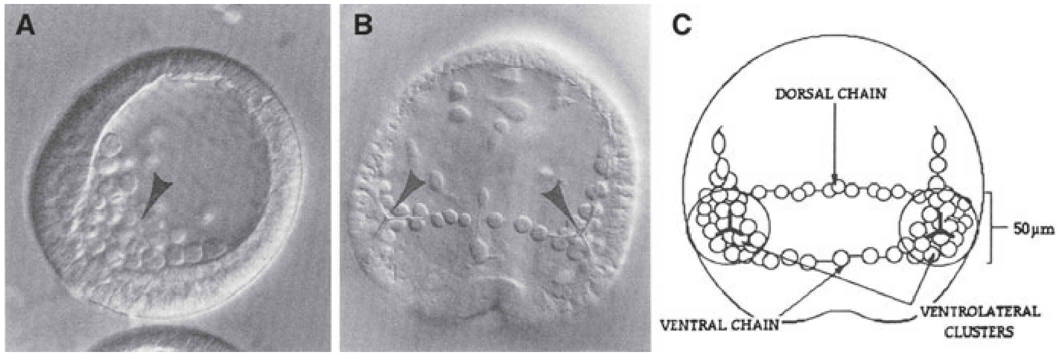


**Fig. 2** Early development of a sea urchin embryo skeleton. **a** Blastula, with primary mesenchyme cells (PMCs) clustered at south pole; **b** PMCs coalesced in ventrolateral pouches; **c** PMCs distributed in subequatorial ring. (From Ettensohn 1990)

4. A classic teleological process in biology is development, for example the transformation of a fertilized egg into a multicellular organism. Many aspects of the process are known to be highly robust, so that a morphologically normal organism can be produced despite large perturbations, what developmental biologists call "regulative development." Sea urchin larvae are model organisms in the study of regulative development, and in particular the regulative development of the larval skeleton has been of interest. Figure 2 shows a normal course of skeletal development, beginning with an early stage in which the larva consists of a hollow ball of cells, the blastula (Fig. 2a). In the next step, a group of cells called primary mesenchyme cells (PMCs) at the south pole of the embryo (arrow in Fig. 2a) move into the internal cavity and begin to migrate, eventually coalescing to form two lateral pouches of cells, the ventrolateral clusters (arrows in Fig. 2b), joined by a ring lying of cells lying just below the equator (Fig. 2c). Boundaries between adjacent cells in the clusters then fuse so that their cytoplasm becomes continuous, and it is within this shared cytoplasm that the larval skeleton is laid down. Experiments over a number of decades have shown that a normal-looking skeleton is produced over a considerable range of initial conditions, overcoming extreme variation in the number of PMCs (including a starting value of zero!) and a wide variety of types of disruption in PMC migration. In other words, development of the larval skeleton looks highly teleological.

How does it work? The developmental biologist Charles Ettensohn and colleagues performed a series of manipulations in the 1980's and 1990's, investigating how urchin PMCs find their target locations (Ettensohn 1990; Ettensohn and McClay 1988). Ettensohn took advantage of the fact that PMCs can be removed microsurgically from blastulae and replaced with PMCs from other embryos, and yet development will proceed normally. In one experiment, he removed PMCs from the ventrolateral clusters of one embryo and implanted them in an early-stage embryo from which the native PMCs had been removed. He observed that the implanted PMCs did not migrate exclusively to the ventrolateral-cluster site of the host embryo, but instead distributed themselves among all normally-PMC-occupied sites, including the dorsal chain. Likewise dorsal chain cells implanted into an emptied early-stage embryo did not move exclusively to the site of a new dorsal chain, but distributed themselves over all normally-PMC-occupied sites. These (and other) findings suggested that the initial population of PMCs does not consist of specialized lateral-cluster cells and specialized dorsal-ring cells, each preprogrammed to migrate toward a particular target location on the ring. Rather, the migrating PMCs are all more or less the same, interchangeable, and their eventual distribution is directed by some global field.

If Ettensohn is right, and an upper-level organizing field is present, the movement of cells within the field is a case of upper direction (property 1), combined with considerable freedom of movement of individual PMCs (property 2). And it is the stability, of the upper-level field—at least on the timescale of PMC migration—that endows the system with its robustness (property 3). The point here is not that all of development must be organized this way. Many steps in development must rely mainly on movements programmed directly into the individual lower-level entities, individual cells. Rather, it is that where the behavior of entities appears to be teleological, the system seems to be hierarchically structured.

## Teleology in evolution

A lineage evolving by natural selection also has the look of an entity persisting in a behavior, a trajectory toward adaptation. One standard view of adaptation is implicitly hierarchical: a species moves uphill on a fitness landscape. This section is an invitation to talk about this process in the language of hierarchy, a language that I think better reveals the causal structure underlying the apparent teleology. In hierarchical terms, a fitness landscape is an ecological context. And an evolving species is an object contained within that context. Upper direction is the effect that the ecological context has on the trajectory of the evolving species.[8]

Consider a Mesozoic lineage of turtles with just the beginnings of a neck-folding capability that would allow the animal to pull the head under the shell. In an ecological context that included turtle predators, suppose that the adaptive arrows pointed uphill toward improved neck-folding capability, and that they did so over a broad range of possible local environments. This is a nice example, because the paleontological evidence indicates that neck folding evolved five times independently, at five different times on five continents (Rosenzweig and McCord 1991). In other words, the process was plastic, with multiple turtle lineages followed similar trajectories from five different starting points.[9] The process was also persistent. Each lineage must have experienced deflections from its trajectory toward improved

[8] Ecological context is relational, of course. The ecological context for an evolving honey bee is different from the ecological context for the flowers evolving in the same area. For present purposes, the ecological context is the set of selective factors that are relevant to the organism in question.

[9] Normally in thinking about plasticity, we think of repeated runs of the system from different starting points within the same enveloping field, which in the turtle examples would mean the same environment. But in this example, there are five somewhat-different environments. However, the assumption is that those environments were sufficiently similar—in terms of selection pressures exerted by predators—that the five runs count as a kind of plasticity.

neck folding (at a minimum, short-timescale deflections due to mutation and drift), deflections from which it recovered.

Where is the hierarchical structure here? A predator-filled environment is the enveloping structure and the source of the upper direction that each evolving turtle-lineages experienced (property 1). Despite this upper direction, there was considerable room for variation—from mutation and drift—in the evolutionary trajectory of each lower-level entity, each turtle lineage (property 2). Finally, in order for selection for neck folding to be effective, the predation component of the ecological context must in each case have been stable, relative to the rate of change of each of the lineages (property 3).

This is an example of parallel evolution, multiple related lineages evolving in the same direction under similar selection pressures. The story would have been different for a single lineage, but still would have demanded hierarchical structure to the extent that the evolutionary change seemed teleological. A bird lineage evolving in an environment with only large seeds experiences selection for a larger, thicker bill to crack the seeds. The ecological context represented by the large seeds is the source of upper direction (property 1). Upper direction is consistent with considerable random variation in bill size due to mutation and drift (property 2). And increase in bill size over a number of generations requires stability of the ecological context in time (property 3).

A word is needed here about function. Sentences like, "The function of the heart is to pump blood," and "The heart beats in order to pump blood" have seemed to imply that a heart beating now is teleological, that it is caused, in some sense, by its consequences. It is this obvious nonsequitur—a present action caused by a future state—that has driven the rich discussion of function in the philosophy of biology (beginning with Wright 1973; Cummins 1975). However, the hierarchical view does not encounter this nonsequitur, because it is only about present proximate causes. In other words, it addresses a different question. Instead of asking how to account for function, it is concerned only with the causes of persistence and plasticity, wherever they occur.

Now persistence and plasticity actually are present in beating hearts, as mechanisms that allow them to correct for irregularities. Further, persistence and plasticity probably played a huge role in the evolution of the heart, certainly persistence. [10] In the course of its evolution, the evolving heart must have suffered many perturbations each generation (e.g., from mutation), but its evolution persisted. In standard evolutionary terms, the heart was subject to stabilizing—as well as directional—selection. In hierarchical terms, we might say that the persistence in evolution of a heart was upper directed by an ecological context in which a heart that pumps blood was essential to survival. However, for present

[10] Interestingly, two levels of hierarchy are involved in many biological systems. The evolution of persistent and plastic physiological systems is itself a persistent and plastic process. In other words, we might say that a "primary" hierarchically structured adaptation process—in the present case, the evolution of the heart—produced a "secondary" hierarchically structured physiological system—the heart's ability to correct for irregularities. We can take this terminology a step further. A torpedo shows persistence and plasticity but it is also a device that was designed by secondary hierarchically structured systems, the minds of its designers, making it a "tertiary" hierarchically structured system.

purposes, none of this matters. In a discussion of function, the phenomenon to be explained is the present blood-pumping activity of the heart, not its error correction capability and not its evolution. And in that present activity, the beating and pumping of blood, there is no persistence and plasticity. And thus there is nothing for a hierarchical account of teleology to explain.

In sum, the hierarchical view offers a causal analysis of persistence and plasticity. But functionality need not involve either of these. Thus the project here is orthogonal to the project of naturalizing function, whether understood from a causal-role (present-causal) or historical (past-causal) perspective (Perlman 2004).

## Human teleological behavior: a plausible speculation

"I bought a Maserati because I wanted one." We will not be concerned for the moment with why I wanted a Maserati, only with the much more proximate cause of my buying one, namely my brain state or states corresponding to wanting one. And we will stipulate that my wanting one was the only cause of my buying one. (No drugs, posthypnotic suggestions, subconscious motives, and such, were involved.) What is wanting? What is desiring? Neuroscience can tell us very little, at present. But if all seemingly teleological systems are best understood in hierarchical terms, as I argue, then we can say something about the causal structure of want-driven behavior. We can say that wanting must be a kind of causal field, one that extends over some behavioral phase space, with the effect of driving the individual to behave in certain general ways from various starting points in that phase space.

One virtue of this view is that it accommodates the folk psychological view of how human behavior works. We think of our behavior as flexible. It is not preprogrammed. There does not seem to be an exhaustive set of situation-specific local rules connecting particular inputs with particular motor sequences. We are not robots. Instead, wants arise, dictating to us the general *sort* of action we should take but without dictating any precise sequence of behaviors to execute. Wanting-behaving, in other words, has the signature properties of an upper directed system. We can think of wanting as a kind of field, one that acts causally but not deterministically on the motor-control system within it.

Let us consider the three properties of apparently teleological systems in the context of wanting the Maserati. The mental state, or series of states, corresponding to wanting the car drives persistent and plastic behavior: looking up the dealerships, investigating the various models, studying my finances, etc. All of these can occur in a huge number of different ways, all different from each other in detail. And any one of them would count as behavior driven by my desire for the car. Also, this desire-for-the-car allows for considerable error-correction and deviation-correction, on a variety of timescales. If an unexpected obligation upsets my plan to buy the car today, I can adjust my schedule to buy it tomorrow, or next week. If I check my bank account and discover I don't have enough money to buy the car this year, I can save my money in the hope of being able to buy it next year. In other words, we think of wanting-to-buy-the-Maserati as a consistent brain state, or series of states, that extends over and governs a broad range of different motor activities, a state or series that is stable on timescales that are long relative to the motor activities involved in making a purchase. If it is not, if I change my mind frequently about wanting this car, from minute to minute or even week to week, then my behavior will appear erratic, and not very teleological.

In hierarchical terms, my buying behavior is directed by wanting, which by hypothesis here we take to be a case of upper direction (property 1). The details of the motor activity involved in buying are highly flexible (property 2). And the teleological aspect of the behavior requires that the wanting "field" be stable on timescales that are long relative to the behavior (property 3).

This hypothesized hierarchical structure is not inconsistent with anything known about affective states, or what are variously called the emotions, motivations, wants, desires, feelings. And it is consistent with two central aspects of the folk-psychological view of deliberate behavior. The first is that deliberate behavior—including both action and speech—is driven and directed by affective states. Value-neutral beliefs, facts, knowledge, etc. are deeply involved in evoking these affective states, of course, and these beliefs establish constraints on behavior, but the common understanding is that motive force and direction arises only from the affective state itself. Second, wants drive behavior in a way that leaves room for variation. When it comes to action and speech, we have options. My wants urge me to escape some awkward situation, but do not tell me how to escape it. The hypothesis here is that the one-to-many relationship between a want and the many behaviors that will satisfy it—i.e., the things I say and do to get it—is captured and accounted for by a hierarchical structure.

What follows from this speculation? We cannot say much at this point about how this hierarchical structure might be physically organized. Want-generating structures could be physically large, encompassing a large number of brain areas, with behavior-generating structures much smaller and physically contained within them. Or want generation could be more localized but with causal influence piped over a larger domain that includes behavior generation. What we can say is that, whatever the physical organization, the domain of influence of wants will be large in some phase space, extending over all of the regions of that space in which conscious behavior-generation occurs.

## A possible distinguishing criterion: complexity/mysteriousness

What distinguishes apparently goal-directed systems from other upper-directed systems showing persistence and plasticity? My tentative answer is that apparently goal-directed systems are more mysterious, perhaps more complex. To put it another way, in persistent and plastic systems that do not seem goal directed, the causal structure is typically simple and obvious. Consider Nagel's (1979) example, a ball in a bowl. A ball resting at the bottom of a bowl is moved away from that point by a perturbing force and then released, so that it rolls back toward the bottom. The structure is hierarchical. The ball is immersed in a gravitational field, which provides upper direction, a force tending to move the ball down toward the bottom from any point within the bowl. But we are not moved to say that the ball seems goal directed, and the reason is that the causes of persistence and plasticity are simple, transparent, familiar. In contrast, in the systems we think of as goal directed—the bacterium, the homing torpedo, the rabbit evading a fox, the developing organism—the causal structure is complex, even mysterious, to most of us. Suppose we found a way to cancel gravity locally and suppose the ball continued to behave as before. Its behavior would now be quite mysterious, and we would now be more inclined, would we not, to describe that behavior as teleological?

Nagel's handling of the problem of distinguishing seemingly goal-directed and non-goal-directed systems—the problem of demarcation—is partly consistent with this proposal and partly not. He has two points to make about demarcation. The first is that the ball in bowl is not goal directed because his systems-theory approach requires that the two relevant variables be in-principle independent. For the ball in the bowl, the two variables are the perturbing force, which displaces the ball from equilibrium at the bottom of the bowl, and the restoring force, gravity, which drives it back toward the bottom. And they are not independent, he says, because the magnitudes of the perturbing and restoring forces are guaranteed by the law of gravity to be equal. In contrast, Nagel explains, a steam engine in which the speed of the engine is controlled by a governor, the engine and the governor are in principle independent of each other.

The hierarchical view agrees that independence is important. There is reason to wonder whether it is really absent for the ball in the bowl in that the sources—if not the magnitudes—of the perturbing and restoring forces are in principle independent. The restoring force is gravity, but the perturbing force could be anything at all, perhaps a flick of my finger, filliping the ball up the side of the bowl. In any case, independence is critical, and it is present by assumption in property 2 of the relevant class of hierarchical systems. Both the ball in the bowl and the governor-regulated steam engine have this property. The ball is immersed in and directed by the gravitational field it lies within, but it can move independently of the field. In the steam engine, the speed of the engine is directed by the structure of the containing system, which includes the governor, yet that speed can change independently (see the thermostat example). So the two views agree on independence, but the hierarchical view tells us something more, namely that in a seemingly goal-directed system, the independence is a special sort. It is the partial independence of an entity from the containing structure that directs it.

Nagel's second point has to do with the source of the causal relationship between the two variables. As he explains, the variables are independent in principle but in fact there must be a causal connection between them in order to produce persistence and plasticity. For the ball in the bowl, that connection arises directly from the laws of nature. In the steam engine, it is imposed by the laws of nature acting through the organization of the system's parts. And this is the difference that lies at the heart of Nagel's view. What makes a persistent and plastic system teleological, he thinks, is that persistence and plasticity arise from the internal structure and organization of the system, not just from natural laws.

One could argue that there is no need to choose between Nagel's view on demarcation and the one I propose, because the two identify roughly the same systems. Causes based in natural law are usually simpler than those based in system structure, and teleological-seeming systems tend on the whole to be rather complex. Still, cases can be imagined that open up some space between the two views. For example, a particle that moves on a trajectory that never perfectly repeats but stays within a small region of space, say on a strange attractor, might count as teleological according to the complexity criterion—its persistent behavior seeming quite mysterious. But using Nagel's criterion, it might not. For example, it might be moving in a law-like way, following force fields with a simple mathematical description.

Nagel recognized that, in his approach, what counts as goal directed will depend on what counts as lawlike in a given scientific context, and therefore on the state of scientific knowledge. The present suggestion that goal-directedness is a function of the perceived complexity of the system is likewise relative to our knowledge. Systems can appear complex or not, depending on how much we know and understand about them. The relativity of goal-directedness to our knowledge is appropriate, given that we are talking about the *appearance* of goal-directedness, about *seeming* future causation. In other words, seeming goal directedness necessarily has an epistemic, even psychological, aspect. And both views acknowledge that. However, the complexity criterion has the advantage that it is applicable even in pre- or non-scientific contexts. Our intuitions about goal directedness are prior to any understanding of natural law. Children use the language of teleology before they know anything of natural law, as did everyone historically before notions of natural law were introduced. The complexity criterion suggests that what our pre-law intuitions may be responding to is the complexity of the system.

## Against "teleology" and "goal directedness"

Upper direction offers a way to talk about persistence and plasticity that removes the temptation to adopt the language of goal-directedness, whether used literally or metaphorically. In the hierarchical view, we are even not tempted to say that a goal or any object in the bacterium's future directs its behavior. The direction clearly comes from a present food field. Nor are we tempted to say that bacterial persistence and plasticity is like or analogous to "real" teleological behavior in humans, because persistent and plastic behavior in humans is also not literally goal directed. Rather, it is directed by present affective responses to presently imagined objects and scenarios. The hierarchical perspective also eliminates certain puzzles that arise in contemporary discussions of teleology. For example, the question of whether teleology has an axiological component (Bedau 1992) does not arise. Also eliminated are the puzzles arising from "goal failure" (Scheffler 1959) and incapacity to achieve goals (Ehring 1984). How can my trip to the refrigerator to get that last piece of cheesecake be called goal directed when the goal is not even present there, because—unknown to me—someone has eaten it? The answer, of course, is that my trip is not literally goal directed, whether or not the cheesecake is there. It is directed by my present desire for cheesecake evoked by my present (false) belief that it is in the fridge. And how can my cat, pawing at the refrigerator door to get the same cheesecake, be said to be goal directed when it does not have the capacity to open the door and achieve the goal? The answer is that it is not goal directed. No "goal state G" plays any part in directing—or even in explaining—its behavior. Rather, the cat's behavior is directed by its present desire for the cheesecake, along with a present (false) belief that the cheesecake is in there and a present (false) belief that it has the capacity to get it. Importantly, I do not mean to imply here that anyone posing these puzzles is implicitly accepting future causation. Virtually no one today does. Rather, the point is that in the hierarchical perspective these puzzles do not even arise as challenges to intuition, as problems to be explained.

It should be clear now that the decision here to avoid terms like teleology and goal directedness (except with modifiers like "seeming" or "apparent") does not imply an eliminitivist position. Persistence and plasticity are real. All that has been eliminated is a loose way of speaking which implies that the future can cause the past.

## The magic revealed

How to deal with the persistent intuition that certain entities behave as though enchanted, acting with impossible knowledge of the future? Here I argue that one way to see behind the magic is to take the perspective of the upper-level structure, rather than the contained entity. Consider the bacterium moving up the food gradient. From the view of the contained entity, from below, the bacterium is almost magically changing its flagellar activation pattern via some very-interesting internal signal-transduction pathways, in response to changes in the local concentration gradient. While the food field looks like a background condition, constant and boring, and not the sort of thing that looks causally important. But from the view of the upper-level structure, from above, the field appears as a thing with causal powers, its broad extent and near constancy making the persistence and plasticity of the contained bacterium possible. From below, all we see is mechanism, the machine-like responses of blood pressure to the activities of certain organs or the predictable electronic cascade in the central processing unit within a torpedo. But from above we see that the contained entities have options, that there is room for error, for variation, and that it is variation that makes persistence and plasticity possible. From below, we see only cognition-related areas of the brain with linkages to muscle groups and speech centers. From above, we see what we believe to be actually going on: my car-buying behavior is driven by a pervasive, pre-existing large-scale structure, a want or a desire.

In a sense, what the hierarchical view does is substitute space for time. When we see an entity persisting on some trajectory, we reflexively expand our temporal view to include later states of the system, goals. In this expanded temporal context, it is tempting to think of those later states as causal. The hierarchical view also invokes an expanded context, but one that is spatially broader (whether in real space or phase space) rather than temporally broader. Seemingly teleological entities operate in spatially broader contexts that are quite stable on the timescale of interest. Indeed, this may be why they seem to transcend time. From the perspective of the behaving entity, those contexts are there at the beginning and they are there at the end. Of course, they do not transcend time at all. They are reliably present simply because they are stable. And the source of their special powers, of their ability to direct the seemingly teleological entity no matter where it goes, is not some occult ability to reach backward in time. It is the fact that they envelope and contain the entity. From the entity's perspective, the containing structure—the source of its direction—is everywhere.

**Acknowledgments** I thank Robert Brandon and Carl Simpson for invaluable comments/suggestions on an earlier version of the manuscript. Also many thanks to Kriti Sharma, Greg Wray, Dave McCandlish, Dave Raup, and the Duke Philosophy of Biology Discussion Group.

## References

Ahl V, Allen TFH (1996) Hierarchy theory: a vision, vocabulary, and epistemology. Columbia University Press, New York
Allen TFH, Starr B (1982) Hierarchy: perspectives for ecological complexity. University of Chicago Press, Chicago
Bedau M (1992) Where is the good in teleology? Philos Phenomenol Res 52:781–806
Braithwaite R (1953) Scientific explanation. Cambridge University Press, Cambridge
Campbell DT (1958) Common fate, similarity, and other indices of the status of aggregates of persons as social entities. Behav Sci 3:14–25
Cummins R (1975) Functional analysis. J Philos 72:741–765
Ehring D (1984) The system-property theory of goal-directed processes. Philos Soc Sci 14:497–504
Ettensohn CA (1990) The regulation of primary mesenchyme cell patterning. Dev Biol 140:261–271
Ettensohn CA, McClay DR (1988) The regulation of primary mesenchyme cell migration in the sea urchin embryo: transplantations of cells and latex beads. Dev Biol 117:380–391
Mayr E (1992) The idea of teleology. J Hist Ideas 53:117–135
McShea DW (2001) The hierarchical structure of organisms: a scale and documentation of a trend in the maximum. Paleobiology 27:405–423
McShea DW, Changizi MA (2003) Three puzzles in hierarchical evolution. Integr Comp Biol 43:74–81
Nagel E (1979) Teleology revisited and other essays in the philosophy and history of science. Columbia University Press, New York
Perlman M (2004) The modern philosophical resurrection of teleology. Monist 87:3–51
Rosenblueth A, Wiener N, Bigelow J (1943) Behavior, purpose and teleology. Philos Sci 10:18–24
Rosenzweig ML, McCord RD (1991) Incumbent replacement: evidence for long-term evolutionary progress. Paleobiology 17:202–213
Salthe SN (1985) Evolving hierarchical systems. Columbia University Press, New York
Scheffler I (1959) Thoughts on teleology. Br J Philos Sci 9:265–284
Schlosser G (1998) Self-reproduction and functionality: a systems theoretical approach to teleological explanation. Synthese 116:303–354
Simon HA (1962) The architecture of complexity. Proc Am Philos Soc 106:467–482
Sommerhoff G (1950) Analytical biology. Oxford University Press, London
Trestman M (2010) Goal directedness, behavior, and evolution: a philosophical investigation. Graduate dissertation, University of California, Davis
Valentine JW, May CL (1996) Hierarchies in biology and paleontology. Paleobiology 22:23–33
Wimsatt WC (1974) Complexity and organization. In: Schaffner KF, Cohen RS (eds) Philosophy of science association 1972. D. Reidel, Dordrecht, pp 67–86
Wimsatt WC (1994) The ontology of complex systems: levels of organization, perspectives, and causal thickets. Can J Philos 20(Suppl):207–274
Wright L (1973) Functions. Philos Rev 82:139–168