



Natural Agency: The Case of Bacterial Cognition

ABSTRACT: *I contrast an ecological account of natural agency with the traditional Cartesian conception using recent research in bacterial cognition and cellular decision making as a test case. I argue that the Cartesian conception—namely, the view that agency presupposes cognition—generates a dilemma between mechanism, the view that bacteria are mere automata, and intellectualism, the view that they exhibit full-blown cognition. Unicellular organisms, however, occupy a middle ground between these two extremes. On the one hand, their capacities and activities are too adaptive to count as mere machines. On the other hand, they lack the open-ended responsiveness of cognitive agents to rational norms. An ecological conception of agency as the gross behavioral capacity to respond to affordances, I argue, does not presuppose cognition and allows for degrees of agency along a continuum, from the simplest adaptive agents, such as unicellular organisms, to the most sophisticated cognitive agents. Bacteria, I conclude, are adaptive agents, hence not mere automata, but not cognitive agents.*

KEYWORDS: agency, life, cognition, bacteria, mechanism, intellectualism

Introduction

I contrast an ecological account of natural agency with the traditional Cartesian conception using recent empirical research in bacterial cognition and cellular decision making as a test case. I argue that the Cartesian conception, the view that agency presupposes cognition, forces us to choose between attributing full-blown belief-desire psychology to bacteria or treating them as mere automata. This conceptual scheme, however, is not sufficiently nuanced to capture the middle ground between these two extremes that most organisms, including unicellular ones, occupy. On the one hand, their capacities and activities are too supplely adaptive to count as mere machines. They act purposively in response to the relevance that their conditions of existence have for attaining their lifestyle. On the other hand, they lack the open-ended responsiveness of cognitive agents to rational

I wish to thank audiences at the ISHPSSB meeting in Montpellier, the University of Memphis, the University of Johannesburg, York University, and the University of Toronto. This paper has also benefited from comments by Mark Sainsbury and from conversations with Fernando Flores-Mangas, Rasmus Winther, and Devin Henry. I also wish to thank two anonymous reviewers of this journal and its editor John Heil. The Rotman Institute of Philosophy at Western University provided an ideal intellectual environment to finish the paper. My biggest debt is to Denis Walsh who commented on many drafts and with whom the positive proposal defended here was developed. This research was supported by the Social Sciences and Humanities Research Council of Canada.



norms. The Cartesian conception thus leads into a dilemma between *mechanism*, which fails by underestimating bacterial agency, and *intellectualism*, which fails by overestimating their agency as cognitive.

An ecological conception, I argue, allows us to occupy a *via media* between these two extremes and hence to resolve the dilemma. Agency, on this view, is the gross behavioral capacity of an ecologically embedded purposive system to bias its repertoire in response to what its conditions afford for pursuing its goals (Walsh 2015; Fulda 2016). Since the repertoire of a system can vary in its richness, there can be degrees of agency along a continuum, from the simplest adaptive agents at one extreme, such as bacteria, to the most sophisticated cognitive agents on the other, such as human beings. Unlike the traditional Cartesian conception, the ecological conception implies that agency does not presuppose cognition and that it cannot be defined in noncircular causal-mechanical terms. This account vindicates what we might think of as the commonsense solution to the problem: it seems to exert irreconcilable demands on a theory. An adequate account of agency must show that most organisms including bacteria are adaptive *agents*, hence not mere automata, but not *cognitive agents*.

This dilemma is well known in the philosophy of action. Yet, the standard focus on actions involving belief-desire psychology has promoted an intellectualist conception of agency as a cognitive aspect of the mind. Recently, however, nonhuman animal action has motivated a less intellectualist view of agency as a primitive aspect of the mind, involving noncognitive intentional states, such as trying and perceiving (Steward 2012). Moreover, the behavior of unicellular eukaryotes has motivated a ‘pre-intentional, even pre-representational and pre-perceptual’ conception of primitive agency in terms of biological function (Burge 2010: 327). Here I take a further step and argue that the science of bacterial dynamics motivates the view that agency goes all the way down to prokaryotes, the simplest living systems. Agency thus becomes a primitive aspect of life. Hence, the dilemma is not just a local problem in the philosophy of action but a foundational problem in the philosophy of organismal biology, including microbiology.

First, I introduce the cognitive turn in microbiology and argue that it is predicated on two premises: Bacteria are purposive agents (1.1), and purposive agency is the mark of cognition (1.2). Then I trace the rationale for the second premise in the mechanistic turn led by Descartes (2.1). I argue that it generates a dilemma between mechanism and intellectualism and criticize some attempts to overcome it (2.2). Finally, I propose a novel method of naturalization (3.1) and introduce the ecological account of agency (3.2). I argue that it solves the dilemma by allowing for degrees of agency along a continuum from mere adaptive agents to full-blown cognitive agents (3.3).

1. The Cognitive Turn in Microbiology

1.1 Bacteria as Purposive Agents

Recent empirical research in bacterial cognition and cellular decision making has exposed the remarkable capacities of the simplest living systems to respond

appropriately across a wide range of conditions. Prokaryote cells can attain and maintain functionally relevant states spontaneously, that is, ‘without an associated genetic or environmental difference’ (Balázsi et al. 2011: 910). Genetically identical populations of bacterial cells sharing the same environment can display markedly different phenotypes. This may be the consequence of chance due to internal stochastic fluctuations (Choi et al. 2008) or to external signals triggered by environmental changes or an ‘autonomous decision’ (Norman et al. 2013). Cells are able to integrate a variety of internal and external stochastic signals from past and present conditions (Losick and Desplan 2008) to anticipate future conditions and respond accordingly (Tagkopoulos et al. 2008; Perkins and Swain 2009). Bacterial cells can also detect and repair lesions to their DNA (Shapiro 2007) and sense the density of the population (‘quorum sensing’) in order to form colonies (‘biofilms’) that enable them to respond collectively as a multicellular organism (Ben-Jacob et al. 2006). In short, individually or collectively prokaryote cells display a remarkable supple adaptive responsiveness to their internal and external conditions of existence that allows them to pursue a variety of lifestyles (Locke 2013).

Consider chemotaxis, the capacity of bacteria to direct their movement toward chemical compounds that are nutrients (‘attractants’), such as glucose, and away from chemical compounds that are toxins (‘repellents’), such as phenol, by sampling the concentration of these chemicals in their environment (Bray 2000). Chemotaxis results from alternating two kinds of movements, tumbling at random and swimming in a straight line. If the bacterium senses that it is moving toward the attractant and away from the repellent, it will keep swimming in a straight line for a longer period before it tumbles. But if it senses that it is moving toward the repellent and away from the attractant, it will tumble sooner in order to try a new direction at random. Chemotaxis is thus achieved by modulating the tumbling frequency.

The mechanical explanation of how chemotaxis is causally produced describes the way extracellular chemical signals are detected by transmembrane receptors that transduce the signal to flagellar protein motors (Alon et al. 1999). These motors control the cell’s movement by rotating either counterclockwise for swimming in a straight line or clockwise for tumbling at random in response to feedback. What matters for our purposes, however, is not the mechanical details of chemotaxis but the adaptive, gross organism-environment dynamics that these mechanisms support.

This adaptive dynamics is the mark of goal-directed behavior (Rosenblueth et al. 1943; Bertalanffy 1969; Boorse 1976; Nagel 1977). Goal-directed behavior is characterized by its distinctive modal profile. If some homeostatic stable end-state *G* is a goal, then the system will reliably produce the necessary means to attain and maintain *G* across a wide range of actual and counterfactual circumstances. That is, starting from different initial conditions and across alternative pathways, the system will modify its behavior appropriately in response to perturbations. The simplest living systems thus manifest purposiveness, one of the distinctive if not *the* distinctive feature of life as such. As Sommerhoff puts it:

On the phenomenal level from which all science must proceed, life is nothing if not just this manifestation of apparent purposiveness and

organic order in material systems. In the last analysis, the beast is not distinguishable from its dung save by the end-serving and integrating activities which unite it into an ordered, self-regulating, and single whole, and impart to the individual whole that unique independence from the vicissitudes of the environment and that unique power to hold its own by making internal adjustments, which all living organism possess in some degree. (1950: 6)

Because of this robust counterfactual dependence between means and goals, knowing the bacterium's goal allows us reliably to predict and explain teleologically what it would do under a range of conditions and to evaluate the appropriateness of its response accordingly. Thus, if the goal is to attain glucose, we can predict that the bacterium would modulate its tumbling frequency appropriately and explain why—as opposed to how—it modulates it thusly, namely, because glucose is nutritious and hence has survival value for the bacterium. This teleological epistemic grip that we gain on the bacterium's gross dynamics based on its purposive character is largely insensitive to the mechanisms that causally realize this purposive dynamics.

Chemotaxis requires the coordination of the various subsystems that form the metabolic and sensorimotor organization of the cell. This is a self-regulatory capacity of the cell as a whole. The goal of avoiding toxins or attaining nutrients is thus a goal of the bacterium, not of any of its subsystems. To regulate the tumbling frequency, bacteria exploit a pervasive phenomenon in nature: random walks. These are processes in which moving objects, such as molecules, have an equal probability of moving in any direction away from their starting point. The particles composing a gas can serve as an example; they bounce around in a container and change their direction as they collide. What is distinctive about chemotaxis—and metabolism more generally—is that cells can *bias* these random walks in a direction that is on average adaptive. Unlike the movement of inanimate particles, the movement of bacteria is adaptively rather than randomly self-directed. Moving toward the source of attraction is thus not something that just passively happens to a bacterium as a blind mechanical consequence of genetic and environmental events. Changing the rate of tumble appropriately is rather something that the bacterium *itself* actively *does*. Bacteria are thus *agents* guided by the relevance that the chemical content of their environment has for fulfilling their life cycle.

Agency and mechanism are not exclusive categories. Agents are mechanistic. However, although agency is mechanically realized, agency is not itself a mechanism but a gross dynamical pattern of adaptive, purposive behavior. We have the mechanical concepts to explain how the goal-directed capacities that constitute agency emerge as the effect of the activities of underlying signal transduction pathways, networks, and biochemical mechanisms. But what kind of concepts should we use to make theoretical sense of the agency of these entities as such, that is, as underwriting teleological predictions, explanations, and normative evaluations?

1.2 Purposive Agency as the Mark of Cognition

The standard way to try to make sense of this is to use cognitive concepts. Cognitive agency is after all the paradigm of agency. Cognitive agents have goals—their intentions or desires—and act according to norms—the norms of reason. Hence, goal-directed agency is taken to be the mark of cognition, the intellectual capacity for thought and inference.¹ McClintock's pathbreaking description of genome repair is perhaps the earliest example of this conceptual strategy in cell biology. She begins by characterizing cells as goal-directed agents:

The conclusion seems inescapable that cells are able to sense the presence in their nuclei of ruptured ends of chromosomes, and then to activate a mechanism that will bring together and then unite these ends, one with another. And this will occur regardless of the initial distance in a telophase nucleus that separated the ruptured ends. This ability of a cell to sense these broken ends, to direct them toward each other, and then to unit them so that the union of the two DNA strands is correctly oriented, is a particularly revealing example of the sensitivity of cells to all that is going on within them. (1984: 193)

And she ends by characterizing their agency in cognitive terms:

They make wise decisions and act upon them.... A goal for the future would be to determine the extent of knowledge the cell has of itself and how they utilized this knowledge in a 'thoughtful' manner when challenged. (1984: 193)

Following McClintock's insight, there has been a tendency implicitly to infer cognition from goal-directed agency. Baker and Stock envisage the next task of research on bacterial signal-transduction to be that of understanding 'how [genes and proteins] are connected to form a dynamic, adaptive cell' (2007: 1023). A task they conceive in cognitive terms: 'How is information converted into knowledge, and how is knowledge sorted, evaluated and combined to guide action, morphogenesis and growth?' (2007: 1023). Shapiro talks about the 'purposeful action by cells under challenge' (2007: 809), which he says motivates the 'contemporary view of cells as cognitive entities acting in response to sensory inputs' (2007: 816). On this view, 'bacteria continually pick up and process information about the environment, internal conditions and other cells to decide on appropriate biochemical and biomechanical actions' (2007: 816). Similarly, Ben-Jacob, a leading expert in biofilm formation, and his colleagues, conceive 'the colony's capabilities to perform collective sensing, distributed information

¹ 'Thought' includes the capacity to form cognitive and conative propositional attitudes that can be semantically evaluated for their truth or satisfaction. 'Inference', in turn, includes the capacity for reasoning about what to believe and do guided by the norms of theoretical and practical reason. I will reject the inference from goal-directedness to cognition later.

processing and collective gene-regulation as fundamental cognitive functions' (Ben-Jacob et al. 2006: 504). Likewise, Bray (2012), a leading expert in bacteria motility, advocates the indispensability of cognitive concepts for understanding the 'knowledge' a cell has of itself. Moreover, Perkins and Swain claim that 'human-developed theories of decision-making under uncertainty apply at the cellular level' (2009: 12). Accordingly, they attribute propositional attitudes and inferential capacities to bacteria:

In situations where the signal fluctuates substantially over time, the cell might be expected to continually update its beliefs.... The network generating bacterial chemotaxis performs such real-time inference.... Once a cell has inferred the most probable state of its environment, it needs to decide an appropriate response. (2009: 4, 5)

To be sure, the claim is not that bacteria have human-level cognitive capacities but that the fundamental elements of cognition are present in bacteria. One might argue, however, that this is a mere metaphor that plays no serious theoretical role in the scientific understanding of cellular dynamics. Rather, it is an honorific title meant to highlight the surprisingly resourceful character of cell behavior. Treating cells *as if* they were rational agents making inferences and acting on their beliefs and desires is a mere heuristic tool or at best a predictive strategy for guiding research until a purely mechanical noncognitive account can be given (Dennett 1971). Attributing cognition to cells is thus pragmatically justified but ultimately theoretically dispensable to explain their behavior.

However, the fact that cognitive attribution persists despite knowledge of signal-transduction-mechanisms indicates that this latter interpretation fails to capture the distinctive explanatory role of cognition. The mark of theoretical indispensability is explanatory power. Cognition is attributed in order to explain the behavioral dynamics of cells *irrespective* of the biochemical details of its mechanical implementation. As in the folk-psychological case, cognitive attribution explains *why*, that is, for what reason or purpose an agent acted in a particular way regardless of *how* or by which neurophysiological mechanisms the behavior was causally produced. These teleological explanations are rationalizing (Davidson 1963). They explain by making sense of the behavior as rationally justified in light of the agent's intentions and the norms of reason. As such, these explanations have normative implications about what agents ought to do given their intentions and circumstances. These explanations capture a distinctive set of empirically observable and counterfactual-supporting higher-level regularities about the gross behavior of organisms. Cognitive concepts are thus theoretically indispensable to explain the behavior of organisms *as rational agents*, whether human or bacterial. After all, as physical movements, all responses, bacterial or human, are mechanically produced. Therefore, to capture the response as an action due to a decision we need to employ cognitive concepts.

The assumption that goal-directed agency is the mark of cognition is also well entrenched in naturalistic philosophy. For Dretske goal-directed behavior is constituted as such by its intentional etiology:

To describe the hen, for example, as engaging in diversionary tactics (to protect her chicks) is already to describe her behavior in a way that presupposes an intentional structure for the internal source of that behavior. The appropriateness of response, then, insofar, as this is relevant to what the organism believes and intends, is a property the response acquires only in virtue of its production by internal states having content. (1980: 284)

Similarly, for Sterelny goal-directed behavior requires representation and inferential capacities:

There can be no flexible and adaptive response to the world without representation. To learn about the world, and to use what we learn to act in new ways, we must be able to represent the world, our goals, and options. Furthermore, we must make appropriate inferences from those representations. (1990: 21)

Even for Dennett (1971), who rejects representationalism, to interpret a system as an agent requires the attribution of beliefs, desires and rationality. What grounds this assumption? Consider the arguments from intentionality and rationality (see Walsh 2008). A teleological explanation explains the occurrence of some event by showing it to be the means to the attainment of some goal *G*. By the time the means occurs, however, *G* is an unactualized future state of affairs. But how can an unactualized future state of affairs explain the occurrence of the means? Cognitive agency provides a model that avoids commitment to some mysterious form of causation, such as backward causation: It is not the unactualized future state of affairs per se that causes and hence explains the occurrence of the means, but the system's *representation* of *G* that does. Goal-directed action thus requires intentionality, the mental representation of goals by the intentions of the agent.

Alternatively, if an agent has goal *G* and doing *r* is the means to attain *G* under circumstances *C*, then *ceteris paribus* the agent *ought* to do *r*. But some argue that for the means to be something the agent ought to produce, *G* must be a state of affairs that ought to be attained (Bedau 1992). And for *G* to be a state of affairs that ought to be attained, *G* must be *good*. To avoid the antinaturalist commitment to intrinsically evaluable states of affairs, cognitive agency again provides a naturalistic model: It is not *G* per se that is good, but the system's representation of *G* as good or under 'the guise of the good', hence as something to be pursued (Velleman 1992). Goal-directed agency thus requires rationality, the capacity to respond to the norms of reason.

In summary, the conclusion that bacteria are cognitive agents follows from two premises: that bacteria are goal-directed agents, and that goal-directed agency

is the mark of cognition. The former is an empirically observable fact about the dynamics of living systems. The latter seems to be motivated by naturalistic considerations. However, as we will see, it was actually the elimination of life's distinctive purposiveness in favor of a categorical division between inanimate matter and rational mind that rendered all purposive agency cognitive.

2. Cartesian Agency

2.1 The Mechanistic Turn and the Primacy of Cognition

Ironically, the premise that goal-directed agency is the mark of cognition was central to the mechanistic turn in biology led by Descartes in the seventeenth century. Descartes sought to oppose the Aristotelian doctrine that the nature of organisms, understood as their goal-directed capacities to pursue their ways of life, grounded a hierarchy of faculties (Lennox 2010): The nutritive faculty or vegetative soul of plants, which includes reproductive and developmental capacities; the sensitive faculty or perceptual soul of animals, which includes instinct and locomotion, and the intellectual faculty or rational soul of human beings, which includes thought and inference. On this scheme there are different goal-directed capacities along a continuum, from nutrition to thought, but only intellectual goal-directed capacities involve cognition. Growth, for example, is an instance of a goal-directed capacity that is not mediated by the intentional representation of the end-state to be achieved. Yet, growth was conceived to be qualitatively distinct from the mere mechanical change characteristic of inanimate matter. Aristotle thus denied the premise that all goal-directed capacities imply cognition: 'It is absurd to suppose that purpose is not present because we do not observe the agent deliberating' (1984: Physics II, 8). In short, there can be purpose without intention.

Unicellular organisms are capable of nourishment and reproduction; thus, they have a vegetative agency. Some of them have the sensitive and locomotive capacities of animal agency although we now take animals to be multicellular eukaryotes. However, because unicellular organisms fail to exhibit thought and reasoning, they lack intellectual agency. Aristotle may have said of them what he said of plants, namely, that in them 'we find the relation of means to end, though the degree of organization is less' (1984: Physics II, 8). Being embodied souls, bacteria are purposive agents. However, this minimal agency does not involve cognition.

In contrast, Descartes's (1988) mechanistic biology denied any distinctive goal-directed capacities to living beings as such (Grene and Depew 2004). Biological phenomena, such as nutrition, locomotion, growth and reproduction, were just the effect of the local motion between bodily parts, not the manifestation of a vegetative or sensitive agency. Descartes (1988) also famously held that cognition took place in a distinct substance ontologically independent from the physical world and unique to human beings. Nonhuman organisms were then reduced to mere automata, leaving cognition as the only legitimate kind of goal-directed capacity. Goal-directed agency thus became the mark of cognition. In short, no purpose without intention.

To support this conclusion Descartes proposed the argument from flexibility. Cognition (reason) is the goal-directed capacity to behave appropriately across an unbounded or open-ended set of conditions: 'Reason is a universal instrument that is alike available on every occasion' (1988: 53). Machines, in contrast, can only operate across a limited set of prespecified conditions. As nonhuman animals can only operate across a limited set of prespecified conditions, they are mere machines. Cognition and thereby goal-directed agency thus became the exclusive commodity of human beings.

While the mark of biological purposiveness (life) is adjustability to a *wide* range of circumstances, the mark of psychological purposiveness (reason) is adjustability to an *unbounded* or *open-ended* range of circumstances in the qualitative sense of not being predictable from evolutionary history. In the Aristotelian scheme life is a distinctive category bridging the Cartesian gap between inanimate matter and rational mind. Life and mind share a common purposive character, but they differ in the degree and normative character of this purposiveness. In contrast, in the Cartesian scheme there is no room for the purposiveness of life as such, and hence there is no room for noncognitive agency: A system is either unboundedly flexible and hence rational, or it can only respond in a task-specific way and hence is automated. Living organisms are thus forced either into the category of inanimate matter, as in Cartesian biology, or into the category of rational-mindedness, as in current microbiology.

This research has shown, against Descartes's conclusion, that the responsiveness of even the simplest living systems is remarkably flexible. This responsiveness is not unbounded, however, and thus, strictly speaking, it does not satisfy Descartes's stringent condition for cognition. Yet, microbiologists take this responsiveness to be evidence of cognition. After all, it is too supplely adaptive for a machine. Bacteria can respond appropriately to *novel* conditions, that is, conditions that have no evolutionary or developmental counterpart. Microbiologists are thus implicitly taking the mark of goal-directedness—responsiveness across a wide range of conditions—to be the mark of cognition—responsiveness across an open-ended range of conditions. In doing so, they accept the Cartesian primacy of cognition even if they reject Descartes's conclusion. So while Descartes denied goal-directed capacities to nonhuman organisms on the grounds that they lacked cognition and hence concluded that they were mere automata, contemporary microbiologists acknowledge these organisms' goal-directed capacities and instead conclude that they are cognitive agents. Both arguments, ironically, are predicated on the premise that goal-directed agency is the mark of cognition.

One might object that while Descartes characterizes 'central' cognitive capacities (reason) as domain-general, recent studies in cognitive science suggest that they might be domain-specific (Carruthers 2006). However, the debate between domain-specificity and domain-generality is about the *architecture* of the mind. The issue there is not whether rational thought and action exhibit the remarkable flexibility that Descartes points out—that is just an observable fact—but whether the subpersonal mechanisms that implement this flexibility are modular. This is an open empirical question. It is worth mentioning that although computers

can outperform human beings in domain-specific tasks, such as games, ‘central’ cognitive capacities, such as common sense and abductive inference, remain recalcitrant to an algorithmic treatment (Fodor 2000).

Another objection is that although the deterministic machines of Descartes’s days lacked flexibility, computers can learn and hence adapt. However, this is a quantitative sense of learning (Marcus 2012). A program can identify statistical patterns in data structures, such as recognizing syllables or images, and it can generalize to new instances within the same domain. Organisms, in contrast, do not just identify correlations but interpret them causally, and they can transfer learning across different domains in response to context-dependent relevance. This is not to deny the possibility, in principle, of a robot indistinguishable from a living organism. Still, the fact remains that no robot has been produced ‘that could be confused with a living organism for more than an instant’ (Brooks 2001: 409), let alone with a cognitive agent.

2.2 The Dilemma between Mechanism and Intellectualism

The problem with the priority of cognition—no purpose without intention—is that it generates a dilemma between *mechanism*, the view that most organisms, including bacteria, are mere automata, and *intellectualism*, the view that most organisms, including bacteria, are full-blown cognitive agents. Mechanism preserves the intuition that there is a fundamental difference between bacteria and full-blown cognitive agents. Cognitive agents are responsive to the norms of reason rather than just to the biological demands of survival and reproduction. This difference marks an important discontinuity between living and cognitive agents that any account of the nature of agency should explain. However, by making this a difference in kind, mechanism fails to preserve the intuition that like cognitive agents, bacteria are purposive agents. After all, they can respond appropriately even to novel conditions. This marks an important continuity between living and cognitive beings that an account of natural agency should also explain. In turn, intellectualism preserves the intuition that bacteria and human beings are alike in pursuing goals, and hence it vindicates the continuity between living and cognitive systems as purposive agents. By taking this continuity to be cognitive, however, intellectualism fails to support the intuition that human beings differ from bacteria in being responsive to rational norms. Unlike the Aristotelian scheme then, the Cartesian scheme is not sufficiently nuanced to capture the continuity and discontinuity between life and mind.

Mechanism has been defended on the grounds that language is necessary for having concepts and beliefs and hence for responding to rational norms (Davidson 1982). An alternative defense holds that unicellular organisms merely transduce but do not semantically represent stimuli and hence lack the capacity to respond inferentially (Fodor 1986). Despite their differences, in both cases unicellular organisms are considered to be mere automata by virtue of failing to satisfy some cognitive criterion. In contrast, biologists Kirschner and colleagues reject mechanism on noncognitive purposive grounds:

Bacterial chemotaxis also appears superficially to be a simple signal-response machine, where an attractant or repellent is perceived by receptors on the bacterial surface to generate a signal that is converted to directed movement. We could imagine all sorts of linkages that would control a motor or a steering mechanism to guide the bacterium by chemical signals. In fact, bacterial chemotaxis is based on the modulation of random movement by ligand binding, resulting in a biased random walk. The specific path any bacterium takes is not directly informed by the binding of the ligand, nor does the individual bacterium at any moment sense a spatial gradient (Berg, 1988). This is quite different from any machine of human design! Biological systems look even less like machines when one considers they can generate order from disorder and can arrive at functional states and responses over a range of starting points, size of components, and size of final product. (2000: 79)

Unlike machines that are the product of design, organisms are embodied, self-producing, self-maintaining, and self-organizing complex adaptive systems that synthesize the very materials out of which they are made. The goals of a machine are determined by the intentions of its designer and hence are extrinsic to the system. The goals of an organism, on the other hand, are determined by the organism itself and hence are intrinsic to the system (Jonas 1966; Nicholson 2013). Mechanism thus fails to accommodate the fact that organisms are the agents of their own adaptive responsiveness, rather than the mere passive vehicles for the replication of genetic information. Even the mechanical description of a robot indistinguishable from an organism would explain how its agency is causally possible, but not its constitutive character as an agent, which grounds the distinctive teleological role that its goals play in understanding its gross dynamics. Bray gives voice to the poverty of mechanism:

What words can we use to describe this cellular information? Contemporary biology embraces reductionism and eschews vitalism. It has been inordinately successful in revealing the structures and function of biological molecules, often at an atomic scale. But it has left us with an extreme, almost puritanical rejection of any account of biological processes that goes beyond physics or chemistry. (2012: 196)

Faced with the poverty of mechanism, Bray is forced toward intellectualism: 'The most natural way to describe this behavior is simply to say: yes, the bacterium knows about temperature and what it means for its survival' (2012: 196).

Intellectualism faces the converse problem of overestimating bacterial agency. However flexible bacterial responses are, they fall short of the unbounded responsiveness of cognitive agents. Bacteria can act in pursuit of biological goals. But their purely adaptive dynamics indicate no responsiveness to rational norms and no inferential capacities to think or deliberate about how to act. Cognition,

furthermore, allows organisms to abstract and generalize beyond the concrete and particular conditions of their immediate environment. Although bacteria are able to anticipate future conditions, these conditions remain circumscribed to their immediate environment. Attributing cognition to bacteria is thus an excessive stretch that blurs important theoretical distinctions.

While mechanism fails to capture bacteria as agents, thus eliminating their distinctive purposive character, intellectualism succeeds but by overestimating their agency as cognitive. There seems to be no intermediate position in this conceptual scheme for most organisms, including unicellular organisms, to occupy. Davidson articulates this tension:

We have many vocabularies for describing nature when we regard it as mindless, and we have a mentalistic vocabulary for describing thought and intentional action: what we lack is a way of describing what is in between. This is particularly evident when we speak of the ‘intentions’ and ‘desires’ of simple animals; we have no better way to explain what they do. (1999: 11)

One strategy is to revise the concept of cognition. On the ‘biogenic’ approach, the dilemma is a consequence of the narrow anthropocentric conception of cognition as reasoning, which applies only to an elite group of organisms and leaves the rest as mere automata. In contrast, a broader biological conception of cognition as sensorimotor coordination can be applied all the way down (van Duijn et al. 2006; Lyon 2015). As Godfrey-Smith puts it: ‘all known (metabolically) living systems engage in some cognitive or proto-cognitive processes’ (forthcoming: 8). Being minimally cognitive, bacteria are thus neither mere automata nor full-blown cognitive agents.

This revision eliminates the distinction between sensorimotor and cognitive capacities and even between metabolic and sensorimotor capacities. There are indeed organisms that are borderline cases between plants and animals, such as sea anemones, and capacities that are borderline between metabolism and sensorimotor coupling, just as there might be borderline cases between sophisticated action-perception cycles and cognition. However, the absence of a natural joint does not imply that there is no theoretical difference. The elitist conception of cognition figures in rationalizing explanations of behavior, but explanations in terms of sensorimotor coordination are mechanical. As such, they lack the characteristic normative implications of rationalizing explanations. Sensorimotor mechanisms explain the occurrence of a behavior as a physical movement, not as a purposive action. The elitist conception, it seems, plays an indispensable role in understanding the behavior of organisms as done by an agent for a reason. Sensorimotor coordination is after all an adaptive capacity and, arguably, all adaptive capacities serve biological goals. However, cognitive capacities are qualitatively different in that they serve goals irrespective of their biological value. Much cognitive behavior is after all maladaptive. The minimal cognition approach, however, takes all organismic capacities to be located on a single dimension such that cognition

differs from other adaptive capacities only in degree. Thus Godfrey-Smith defines ‘proto-cognition’ as the capacity of simple organisms to

sense and respond to events, both internal and external, in a way that helps keep them alive, implementing a distinction between states and outcomes that are sought and maintained and other states and outcomes that are avoided. (12)

There is however nothing distinctively cognitive about this capacity. Neither inferential capacities nor propositional attitudes are involved. But neither is protocognition a mere blind mechanical response. To say that states and outcomes are ‘sought’ or ‘avoided’ implies purposiveness. Staying alive is a goal, and responding to sensed events is a means to it. Citing the goal explains why the events were sought or avoided, namely, because they are good or bad, respectively, for staying alive. To call this capacity ‘minimal cognition’ or ‘protocognition’ rather than ‘proto-agency’ or ‘minimal agency’ betrays a commitment to the Cartesian primacy of cognition.

A recent attempt to avoid this commitment appeals to the distinctive ‘autonomous’ organization of living systems (Barandiaran et al. 2009; Moreno and Mossio 2015; see also Juarrero 1999). An autonomous system, such as a metabolic system, is a nonlinear, far from equilibrium system that maintains its internal organization by adaptively regulating its dynamical coupling with the environment. The goals are the stability states that the system must attain in order to maintain itself. Since these states must occur for the system to exist, the system *ought* to produce them. These states thus specify the norms the system must conform to in order to exist. Because goals are not specified psychologically in terms of intentions or desires and since the norms they generate are not rules of rationality, this account is not committed to the primacy of cognition. As autonomous systems, bacteria are neither mere automata nor full-blown cognitive agents but ‘minimal-agents’.

This approach makes a fundamental contribution to the naturalization of agency by locating it in the causal structure of the physical world. However, by defining agency in a noncircular way in terms of the causal-mechanical organization that realizes it, agency is identified with its causal basis. In particular, goals are defined as a subset of the effects that a self-maintaining system produces, and hence they explain behavior as effects not as goals. Teleology is thus reduced to mechanical explanation by self-maintenance. But the mechanical explanation, we have seen, cannot capture the constitutive purposive and normative character of agency as such. Agency is thus located in the causal structure of the world at the cost of eliminating its constitutive teleological role. Therefore, although the autonomous systems account avoids intellectualism, it collapses into a sophisticated form of mechanism. Hence, it cannot solve the dilemma.

To capture the constitutive teleological-normative character of agency, ‘enactivism’ complements this causal account with a phenomenological account of cognition as ‘sense-making’ or ‘conduct in relation to meaning and norms that the system itself enacts or brings forth on the basis of its autonomy’ (Thompson 2007:

126; Di Paolo 2005). Sense-making is not defined in causal-mechanical terms, and it does not involve belief-desire psychology. Accordingly, as (minimal) sense-makers, bacteria are neither mere automata nor full-blown cognitive agents.

However, enactivism is committed to the thesis that life equals cognition. But if cognition is identical with or constitutive of life and life is intrinsically purposive, then cognition is constitutive of life's intrinsic purposiveness. Thus, despite its anti-Cartesian character, enactivism is committed to the primacy of cognition. Furthermore, there is again nothing distinctively cognitive about sense-making. Cognition, unlike the generic adaptive responsiveness of organisms, involves *propositional* meaning and *rational* norms. Finally, the irreducible teleological-normative character of sense-making is partly transcendently constituted: The subjective experience of our own embodied agency is an a priori epistemic condition of possibility for accessing life's purposiveness (Jonas 1966). This supraempirical, first-person methodology is incompatible with *scientific* naturalism.

Replacing the belief-desire conception of cognition with a more basic adaptive capacity accounts for the life-mind continuity without succumbing to intellectualism. However, the belief-desire conception remains indispensable to explain the behavior of rational agents capable of thought and inference. Revisionism thus fails to account for the life-mind discontinuity. This suggests two desiderata for solving the dilemma: Reject the primacy of cognition—no purpose without intention—and reject the demand for a mechanical definition—to be is to be causally realized.

3. Ecological Agency

3.1 Two Methods of Naturalization

The standard strategy of naturalization exemplified by the autonomous systems approach seeks to provide noncircular, causal-mechanical conditions for something to be an agent. This strategy is predicated on the strong reductionist assumption that to say what some phenomenon is—its principle of individuation—just is to say how it is causally realized or physically constituted. But this assumption is incompatible with indisputably natural phenomena familiar from the sciences of complex systems dynamics.

Consider viscosity, the disposition of a fluid to resist gradual deformation under shear stress. This disposition is causally realized by the properties of the individual molecules, such as shape, momentum, and charge, that physically constitute the fluid. However, the role the concept of viscosity plays in fluid dynamics is not specified in terms of the properties of its molecular realizer but in terms of other concepts at the same gross dynamical scale, such as density or surface tension. In fact, models of viscous behavior assume that fluids are materially constituted as a continuum, despite their actual discrete microphysical constitution. The scientific role of viscosity is thus largely insensitive to the underlying molecular dynamics that causally realize it (Batterman 2005; Strevens 2005). As the continuum idealization indicates, viscosity *as viscosity*—i.e., as set of behaviors—plays an indispensable

role in the prediction and explanation of the gross dynamics of fluids. As Goldenfeld and Kadanoff put it:

In fluid dynamics the large-scale structure is independent of the detailed description of the motion of the small scales.... To get these gross features, one should most often use a more phenomenological and aggregated description, aimed specifically at the higher level... by trying to separate universal scaling features from specific features. (1999: 87–88)

Similarly, thermodynamic behavior is ‘universal’ in the sense that it can be instantiated across systems with very different microphysical constitutions (Batterman 2005). Temperature, for example, is realized by the mean molecular kinetic energy only in gases, by the black body distribution of electromagnetic waves in a vacuum, and in plasma there is temperature even if there are no constituent molecules. A gas made up of molecules and radiation made of wavelengths of light can have the same temperature (Sklar 2015). We cannot individuate temperature just by specifying the microphysical conditions that causally realize it. We need to appeal to the gross behavior of the system as a whole. This universality enables us to explain and predict when systems will be in thermal equilibrium regardless of their composition and size. Thermodynamic behavior is particularly insensitive to the mechanical details of its implementation at critical points at phase transition. This insensitivity is represented as a mathematical singularity in statistical mechanics.

Viscosity and temperature are thus natural as such by virtue of the indispensable theoretical role they play irrespective of their causal realizers. Hence, the standard strategy of naturalization is committed to an unnecessarily stringent view of naturalism. From the fact that having an autonomous organization is a causally necessary and/or sufficient condition for instantiating agency, it does not follow that the concept of agency is the concept of having such a causal-mechanical organization. After all, the teleological explanation and normative evaluation of an organism’s gross dynamics is insensitive to the mechanisms that underpin it. This suggests an alternative two-stage strategy of naturalization: First, locate the phenomenon in the causal structure of the physical world; second, provide a constitutive account of the nature of that phenomenon in terms of the set of behaviors that characterize its gross dynamics.

3.2 Agency as an Ecological Concept

I take the autonomous systems approach to offer an account of the first stage. As for the second stage, I follow a recent approach in the philosophy of biology—‘situated adaptationism’—that seeks to understand the nature of organisms as natural agents and its significance for evolution (Walsh 2015; I develop this account further in Fulda 2016). Agency, on this view, is essentially an ecological rather than a mechanical phenomenon. It is defined by the way the organism is dynamically *embedded* in the world, not by its inner workings. This embeddedness

is not a mere external relation of reciprocal causation. It is rather a relation of reciprocal constitution between the organism and its ecological niche or ‘affordance landscape’, the set of surrounding conditions that are relevant for pursuing its lifestyle. Agency thus emerges as agency only when we consider the organism and its surroundings as a single, coupled dynamical system. More precisely, agency is the gross behavioral capacity of an organism to bias its repertoire in response to what its conditions afford for attaining its goals. This set of ecological concepts—goal, repertoire, and affordance—form a tight interdefined theoretical cluster that specifies the ecological profile of the system, just as the concepts of temperature, entropy, volume, and pressure specify the thermodynamic profile of a system. I take each in turn.

A system S has a goal G if S has the *repertoire* to attain and maintain G across a wide range of actual and counterfactual conditions C by responding appropriately to what C affords. If G is S ’s goal and r is a means to the attainment of G in conditions C , then *ceteris paribus*, S would bring about r in C for a suitably large range of values for G , r , and C . And if r is detrimental to the attainment of G in C , then S typically would not bring about r . A system S has a goal G , then, if it supports the following counterfactuals: If the goal G were to change to G^* while conditions C remain constant, then S would modify its behavior r accordingly toward G^* by producing r^* . And if the conditions C change to C^* while the goal G remains constant, then S would modify its behavior r accordingly to maintain its trajectory toward G by producing r^* . Call this modal profile *hypothetical invariance*. A state of affairs G is a goal only if there is a system S that tends to attain and maintain G in a hypothetically invariant way. Goals, in turn, explain their means not by causing them but by rendering their occurrence hypothetically necessary for their attainment (Walsh 2008). Against the argument from intentionality, then, there is no worry about mysterious causes or a need to appeal to mental representations: There can be purpose without intention.

Hypothetical invariance implies a *repertoire*. That is, a set of responses $R \{r_1, r_2, \dots, r_n\}$ that S could produce under current conditions and the capacity to select the subset that is appropriate given what C affords. Accordingly, a repertoire is not just a set of movements or responses, but a set of things that the system can do. A given response from a repertoire must be selected by the system *because* it is conducive to G in C . Contra intellectualism, the capacity to select an item from the repertoire because it is appropriate need not amount to a decision in the sense of deliberation. The self-regulatory capacities of the system as a whole *bias* its repertoire toward achieving the goals by exploiting or mitigating the system’s affordances.

In turn, a repertoire implies a set of conditions that facilitate or impede the organism in attaining and maintaining its goals. These conditions are affordances, opportunities for action that an aspect of a system’s physical surroundings provides (or impedes) relative to the system’s capacities and goals (Gibson 1979). A hole in the wall affords escaping to a mouse but not to a horse, and water affords drinking to most animals, swimming to a fish, and a supportive surface to water striders. Affordances capture the constitutive interdependence between organism

and environment, that is, the fact that they imply each other. If the bacterium has the goal of getting nourishment and the capacity to digest glucose, then the presence of glucose affords nourishment to the bacterium. Hence, *ceteris paribus* the bacterium will respond to glucose as edible by swimming toward it. The bacterium swims toward the glucose, then, because it is edible. Being edible is not an intrinsic property of the physicochemical structure of the glucose molecule. It is a relation between the glucose molecule and the bacterium jointly constituted by the physicochemical structure of the molecule and the metabolic capacities of the bacterium (Thompson 2007). Affordances are thus emergent properties of the whole organism-environment system and are qualitatively distinct from their constituent parts (Chemero 2003).

Affordances also have normative implications. For an organism to experience its ecological setting as an affordances landscape is for the elements of that setting to have value or significance for the attainment of the organism's goals. The fact that glucose affords food to the bacterium implies a *ceteris paribus* that getting it is good for pursuing its lifestyle. Its ecological value in turn imposes a requirement on the bacterium to implement the appropriate item from its repertoire, namely, swimming in the direction of the glucose, the nourishment. Against the argument from normativity, then, normativity does not require cognition. As Gibson puts it: 'The meaning and value of a thing consist of what it affords' (1982: 407).

The ecological approach satisfies the two desiderata for solving the dilemma: First, agency does not presuppose cognition. This avoids the intellectualist horn. Second, agency cannot be given a causal-mechanical definition. It is defined ecologically in terms of goals, repertoire, and affordances. This avoids the mechanistic horn. The combination of scientific naturalism without reductionism and normativity without intellectualism makes organismal life the *tertium quid* between inanimate matter and rational mind that solves the dilemma.

One might object that the conclusion that agency does not presuppose cognition depends on setting an implausibly low bar for agency and an implausibly high bar for cognition. A more plausible view would be that simple agency involves simple cognition while complex agency involves complex cognition. But however simple we take cognition to be, it is an *intellectual* capacity. Having propositional attitudes and inferential capacities, however rudimentary, implies a minimal sensitivity and responsiveness to rational norms. The evidence, however, indicates that bacteria and most organisms are sensitive and responsive only to the most basic biological norms. Hence, to attribute even the simplest cognition to bacteria is to intellectualize them, however minimally. In contrast, the view that agency does not presuppose cognition allows us to throw away the bath water of intellectualism without the purposive character of bacterial responsiveness.²

² Does a thermostat count as an agent on this view? The thermostat lacks a repertoire as such, that is, a range of possible responses to various circumstances. It has only one response—expanding—to a very circumscribed set of conditions—a specific change in temperature. Furthermore, a thermostat's 'goals' are determined extrinsically by its designer, not intrinsically by the way the thermostat *itself* is embedded.

3.3 A Continuum of Natural Agency

An agent is any system that can respond to affordances as affordances. But the repertoire of an agent can vary in its richness along a continuum, from the simplest biological agents, such as unicellular organisms, to the most complex cognitive agents, such as human beings. The broader the repertoire of a system, the broader the array of affordances it can respond to and hence the wider the range of goals it can pursue. As we move along the continuum, the ecological dynamics of organisms become increasingly more complex until belief-desire psychology becomes necessary. The ecological approach thus allows for degrees of agency across a broad spectrum of ecological profiles (Fulda 2016). Cognitive agency is thus not agency *tout court* but only a highly sophisticated instance.

Unicellular organisms have comparatively the most limited repertoires. Chemotaxis involves only two modes of locomotion, swimming in a straight line and tumbling at random. Because of their limited repertoire, bacteria can respond only to a comparatively limited range of affordances, and hence they can attain their goals only across a very limited set of conditions. Furthermore, the range of goals that they can pursue is limited to the most basic biological demands of their life cycle. Still, their activities are supple and adaptive enough to warrant treating them as responses to affordances, that is, as hypothetically invariant and hence as agential. This agency is surely minimal, but is agency nonetheless.

As we move along the continuum, new repertoires emerge, and hence new and more complex forms of agency appear. The emergence of new mechanisms and new environmental conditions expand what a system can do and hence enable new repertoires, which in turn disclose new affordances and allow for a new range of goals that open new lifestyles to pursue. For example, recent experimental evidence shows that slime mold, an eukaryotic single-celled organism, tends to avoid retracing its own paths by using ‘externalized spatial memory’ to navigate its surroundings (Reid et al. 2012). By leaving behind a trail of slime in its path the slime mold can later detect where it has already been. The repertoire of the slime mold includes responses that are unavailable to bacteria, such as moving in various directions at the same time. This repertoire partly constitutes a new affordance landscape that allows the organism to pursue a new range of goals, such as following or avoiding the trails left by other organisms.

The emergence of multicellularity represents a crucial transition in the sophistication of agency, as exemplified by the diversity of plant and animal lifestyles. And so does the emergence of basic mental capacities, such as perception, in the sophistication of animal agency. The repertoire of arthropods, which includes jumping, running, hiding, and hunting, is comparatively much richer than that of unicellular organisms, but not as rich as the repertoire of mammals, which includes playing and extended parental care.

The emergence of cognitive capacities, in turn, allows animals to pursue a new range of goals well beyond those imposed by mere biological demands, such as cognitive and conative goals, and hence they can respond to the norms of reason, such as epistemic and ethical norms. This new affordance landscape makes available a novel normative dimension for evaluating purposive behavior that requires the

conceptual apparatus of belief-desire psychology. After all, we attribute desires and intentions to cognitive agents irrespective of whether their satisfaction promotes or impedes their biological well-being. We also explain their behavior in terms of these intentions and judge them to be rationally justified or not according to whether they contribute to the fulfillment of those goals. Accordingly, intentional attributions are made true not by discrete facts about the inner workings of the agent but by relational facts about the overall dynamics of the whole ecological system. Cognitive agency is thus a highly sophisticated kind of natural agency, but still continuous with the minimal agency of unicellular organisms.

The ecological approach thus offers the *via media* that Davidson thought was unavailable, but it also does something subtler than this. The ecological approach transforms the problematic Cartesian dilemma into an Aristotelian continuum with mechanism and intellectualism at its poles. On this graded picture, bacteria are very close to the mechanistic pole while a squirrel is very close to the intellectualist pole. There is thus a vast in-between space for most organisms to occupy. And because the concept of agency applies to patterns of behavior rather than to the mechanisms that underpin them, there is no problem with borderline cases. Arguably, bacteria are automata approaching minimal agency while squirrels are adaptive agents approaching cognitive agency. This continuum makes room for a potential measure of the distance from each of the poles. In principle, we can locate the agency of any organism on the basis of this measure.

The revisionist approaches also conceive the difference between bacteria and human beings as a difference in degree along a continuum. However, they take a one-dimensional view of the continuum. Each new element in the repertoire of an agent is just the amplification or sophistication of a previous one, its precursor. But on the multidimensional view of the ecological approach there can be differences of kind along the continuum. As the agency of organisms increases in its richness, it takes on qualitatively new elements to its repertoire, such as cognition. The ecological approach thus accommodates both the continuity and the discontinuity between life and mind. Surely, there are degrees of cognition. However, it does not follow that there are precursors of cognition in previous forms of noncognitive adaptive agency. For example, as the metabolic rates of organisms increase, new elements of the repertoire can emerge, such as new modes of locomotion. These do not have precursors in simpler forms of agency. They are qualitatively new elements of the repertoire.

4. Conclusion

The ecological account vindicates what we might think of as the commonsense solution to the dilemma: it seems to exert irreconcilable demands on a theory. The fact that bacterial dynamics can be described ecologically shows, against mechanism, that bacteria are not mere automata but adaptive *agents*. They are responsive to what their conditions afford for attaining their biological goals. However, the fact that their repertoire is very limited shows, against intellectualism, that bacteria are not *cognitive* agents. In contrast, the Cartesian conception offers

us a dichotomy between mechanism, which entails the complete lack of agency, and intellectualism, which entails full-blown cognitive agency, with no gradations in between. However, on the ecological approach, mechanism and intellectualism occupy either pole of a spectrum. The recent cognitive turn in microbiology has made a fundamental empirical contribution to our understanding of organisms, including prokaryotes, as agents. But the Cartesian conception that underlies the cognitive assumption obfuscates the fundamental neo-Aristotelian insight that agency, not cognition, is a primitive feature of organic life.

FERMÍN C. FULDA
 ROTMAN INSTITUTE OF PHILOSOPHY,
 WESTERN UNIVERSITY
fuldafermin@gmail.com

References

- Alon, U., M. G. Surette, N. Barkai, and S. Leibler. (1999) 'Robustness in Bacterial Chemotaxis'. *Nature*, 97, 168–71.
- Aristotle. (1984) *The Complete Works of Aristotle*. Edited by J. Barnes. 2 vols. Princeton, NJ: Princeton University Press.
- Balázs, G., A. V. Oudenaarden, and J. J. Collins. (2011) 'Cellular Decision Making and Biological Noise: From Microbes to Mammals'. *Cell*, 144, 910–25.
- Baker, M. D., and J. B. Stock. (2007) 'Signal Transduction: Networks and Integrated Circuits in Bacterial Cognition'. *Current Biology*, 17, 1021–24.
- Barandiaran, X., E. Di Paolo, and M. Rohde. (2009) 'Defining Agency: Individuality, Normativity, Asymmetry and Spatio-temporality in Action'. *Adaptive Behavior*, 17, 367–86.
- Batterman, R. (2005) 'Critical Phenomena and Breaking Drops: Infinite Idealizations in Physics'. *Studies in History and Philosophy of Modern Physics*, 36, 225–44.
- Bedau, M. (1992) 'Where's the Good in Teleology?' *Philosophy and Phenomenological Research*, 52, 781–805.
- Ben-Jacob, E., Y. Shapira, and A. I. Tauber. (2006) 'Seeking the Foundations of Cognition in Bacteria: From Schrödinger's Negative Entropy to Latent Information'. *Physica A: Statistical Mechanics and Its Applications*, 359, 495–524.
- Bertalanffy, L. (1969) *General Systems Theory*. New York: Braziller.
- Boorse, C. (1976) 'Wright on Functions'. *Philosophical Review*, 85, 70–86.
- Bray, D. (2000) *Cell Movements: From Molecules to Motility*. New York: Garland Science.
- Bray, D. (2012) 'The Cell as a Thermostat: How Much Does It Know?'. *Advances in Experimental Medicine and Biology*, 736, 193–99.
- Brooks, R. (2001) 'The Relationship between Matter and Life'. *Nature*, 409, 409–11.
- Burge, T. (2010) *The Origins of Objectivity*. New York: Oxford University Press.
- Carruthers, P. (2006) *The Architecture of the Mind*. New York: Oxford University Press.
- Chemero, A. (2003) 'An Outline of a Theory of Affordances'. *Ecological Psychology*, 15, 181–95.
- Choi, P. J., L. Cai, K. Frieda, and S. X. Xie. (2008) 'A Stochastic Single-molecule Event Triggers Phenotype Switching of a Bacterial Cell'. *Science*, 322, 442–46.
- Davidson, D. (1963) 'Actions, Reasons and Causes'. *Journal of Philosophy*, 60, 685–700.
- Davidson, D. (1982) 'Rational Animals'. *Dialectica*, 36, 317–28.
- Davidson, D. (1999) 'The Emergence of Thought'. *Erkenntnis*, 51, 511–21.
- Dennett, D. (1971) 'Intentional Systems'. *Journal of Philosophy*, 68, 87–106.
- Descartes, R. (1988) *The Philosophical Writings of Descartes*. Edited by John Cottingham, Robert Stoothoff, and Dugald Murdoch. 3 vols. New York: Cambridge University Press.

- Di Paolo, E. (2005) 'Autopoiesis, Adaptativity, Teleology, Agency'. *Phenomenology and the Cognitive Sciences*, 4, 429–52.
- Dretske, F. (1980) 'The Intentionality of Cognitive States'. *Midwest Studies in Philosophy*, 5, 281–94.
- Fodor, J. (1986) 'Why Paramecia Don't Have Mental Representations'. *Midwest Studies in Philosophy*, 10, 3–23.
- Fodor, J. (2000) *The Mind Doesn't Work that Way*. Cambridge, MA: MIT Press.
- Fulda, F. (2016) 'Natural Agency: An Ecological Approach'. PhD. diss. University of Toronto.
- Gibson, J. J. (1979) *The Ecological Approach to Visual Perception*. Boston: Houghton Mifflin.
- Gibson, J. J. (1982) *Reasons for Realism: Selected Essays of James J. Gibson*. Edited by E. Reed and R. Jones. New Jersey: Lawrence Erlbaum.
- Godfrey-Smith, P. (Forthcoming) 'Mind, Matter, and Metabolism'. *Journal of Philosophy*.
- Goldenfeld, N., and L. P. Kadanoff. (1999) 'Simple Lessons from Complexity'. *Science*, 284, 87–89.
- Grene, M., and D. Depew. (2004) *The Philosophy of Biology: An Episodic History*. New York: Cambridge University Press.
- Jonas, H. (1966) *The Phenomenon of Life: Toward a Philosophical Biology*. New York: Harper & Row.
- Juarrero, A. (1999) *Dynamics in Action: Intentional Behavior as a Complex System*. Cambridge, MA: MIT Press.
- Kirschner, M., J. Gerhart, and T. Mitchison. (2000) 'Molecular Vitalism'. *Cell*, 100, 79–88.
- Lennox, J. (2010) 'Bios and Explanatory Unity in Aristotle's Biology.' In D. Charles (ed.), *Definition in Greek Philosophy* (Oxford: Oxford University Press), 329–55.
- Losick, R., and C. Desplan. (2008) 'Stochasticity and Cell Fate'. *Science*, 320, 65–68.
- Locke, J. C. (2013) 'Systems Biology: How Bacteria Choose a Lifestyle'. *Nature*, 503, 476–77.
- Lyon, P. (2015) 'The Cognitive Cell: Bacterial Behavior Reconsidered'. *Frontiers in Microbiology*, 6, 264.
- Marcus, G. (2012) 'Is "Deep Learning" a Revolution in Artificial Intelligence?'. *The New Yorker*, November 25.
- McClintock, B. (1984) 'The Significance of Responses of the Genome to Challenge'. *Science*, 226, 792–801.
- Moreno, A., and M. Mossio. (2015) *Biological Autonomy: A Philosophical and Theoretical Enquiry*. Netherlands: Springer.
- Nagel, E. (1977) 'Teleology Revisited'. *Journal of Philosophy*, 84, 261–301.
- Nicholson, D. (2013) 'Organisms ≠ Machines'. *Studies in History and Philosophy of Biological and Biomedical Science*, 44, 669–78.
- Norman, M., D. L. Nathan, J. Paulsson, and R. Losick. (2013) 'Memory and Modularity in Cell-fate Decision-making'. *Nature*, 503, 482–86.
- Perkins, T. J., and P. S. Swain. (2009) 'Strategies for Cellular Decision-making'. *Molecular Systems Biology*, 5, 1–15.
- Reid, C. R., T. Latty, A. Dussutour, and M. Beekman. (2012) 'Slime Mould Uses an Externalized Spatial "Memory" to Navigate in Complex Environments'. *Proceedings of the National Academy of Sciences of the USA*, 109, 17490–94.
- Rosenblueth, A., N. Wiener, and J. Bigelow. (1943) 'Behavior, Purpose and Teleology'. *Philosophy of Science*, 10, 18–24.
- Shapiro, J. A. (2007) 'Bacteria are Small But Not Stupid: Cognition, Natural Genetic Engineering and Socio-bacteriology'. *Studies in the History and Philosophy of Biology and Biomedical Sciences*, 38, 807–19.
- Sklar, Lawrence. (2015) 'Philosophy of Statistical Mechanics'. In Edward N. Zalta (ed.), *The Stanford Encyclopedia of Philosophy*. Available at: <http://plato.stanford.edu/archives/fall2015/entries/statphys-statmech/>.
- Sommerhoff, G. (1950) *Analytic Biology*. Cambridge, UK: Cambridge University Press.
- Sterelny, K. (1990) *The Representational Theory of Mind: An Introduction*. Oxford: Wiley-Blackwell.
- Steward, H. (2012) *A Metaphysics for Freedom*. Oxford: Oxford University Press.

- Strevens, M. (2005) 'How are the Sciences of Complex Systems Possible?'. *Philosophy of Science*, 72, 531–56.
- Tagkopoulos, I., Y. Liu, and S. Tavazoie. (2008) 'Predictive Behavior within Microbial Genetic Networks'. *Science*, 320, 1313–17.
- Thompson, E. (2007) *Mind in Life*. Cambridge, MA: Harvard University Press.
- van Duijn, M., F. Keijzer, and D. Franken. (2006) 'Principles of Minimal Cognition: Casting Cognition as Sensorimotor Coordination'. *Adaptive Behavior*, 14, 157–70.
- Velleman, D. (1992) 'The Guise of the Good'. *Nous*, 26, 3–26.
- Walsh, D. M. (2008) 'Teleology'. In Michael Ruse (ed.), *The Oxford Handbook of Philosophy of Biology* (Oxford: Oxford University Press), 113–36.
- Walsh, D. M. (2015) *Organisms, Agency, and Evolution*. Cambridge, UK: Cambridge University Press.