Teleology, Intentionality and Acting for Reasons

In this paper I would like to contrast two radically different approaches to the evident linkage between an agent acting for a reason and that agent possessing intentional states. Both of these strategies have distinguished lineages. One of these views traces back to Aristotle, the other traces back to Kant. These approaches to intentionality share in common a commitment to two central premises. First, both the Aristotelian and the Kantian strategy hold that the intentional states of agents are the states of those agents that can serve essential roles in rational explanations of the actions of those agents. Second, they both hold that the content of an intentional state is fixed by the specific role of that state in providing those rational explanations. The strategies fundamentally differ, however, in their analyses of the ultimate character of reason giving explanations. For the Aristotelian, reasons are unintelligible outside of teleological contexts; ultimately nothing can be acting for a reason unless that agent also acts so as to achieve goals. For the Kantian, reasons are unintelligible apart from the ability to represent a principle or rule; ultimately nothing can be acting for a reason unless that agent is capable of acting because of the acknowledgment of a representation of a law.

After articulating both of these general approaches to the relation between practical reason and intentionality I will argue that, in a crucial respect, the Aristotelian tradition is right. Nothing could act for a reason in the way that Kant analyzes unless it could also act for reasons in the way in which Aristotle analyzes. But it is possible for an agent to act for a reason in an Aristotelian sense without being capable of acting for a reason in the Kantian way. And, because this is true, it is possible for non-linguistic organisms to have intentional states. Further, it is a necessary condition on linguistic creatures such as ourselves having the type of intentionality that we possess, that we also have intentional states that are rooted in the teleology of our animal natures. These are several of the claims of my recent book, *Rational Animals: The Teleological Roots of Intentionality*.

ARISTOTELIAN TELEOLOGY

"And if [an organism has the faculty] of sense-perception, then that [organism has the faculty] of desire also; for desire comprises wanting, passion, wishing…"[1]

How is a lion different from an iron bar? This question is not as easy to answer as it might at first glance seem to be. For example, one might be tempted to say that the lion perceives her environment, and an iron bar doesn't. But what is it to perceive? I suppose that in a general sort of way, perception is the ability of an entity to detect differences in an environment in the process of responding differentially to those differences. But the iron bar is differentially responsive to aspects of its environment, just as the lion is. In fact, the bar is more reliably responsive to certain aspects of the environment than the lion. The bar will rust just in case there is water vapor present in the environment, but a lion won't always respond to the presence of a buffalo by crouching, or moving to the left, or moving downwind from the buffalo, or moving towards the buffalo, or … So why do we say that the lion can perceive and the iron bar can't?

Surely part of the answer to this question is hinted at in the quote from Aristotle. Differential responsiveness, by itself, is no indication of perception. Rather, such responsiveness indicates the presence of perception only in a context provided by the entity that responds having a goal or *telos* that is facilitated by the response. Iron bars have no such goal to their 'endeavors'. They simply respond to their environment in a way that can be straightforwardly explained in Hempelian fashion by appealing to laws that govern the behavior of iron. Lions, on the other hand, can be counted on to respond to their environment in ways that are, in some sense to be explicated, instrumental to reaching certain end states. A lion, for example, *when it is hunting*, often responds to the presence of a buffalo by taking various actions (moving downwind from the buffalo, crouching in the grass, etc.) that result in concealment from the buffalo. That is, when the lion is pursuing the buffalo in order to kill it in order to eat it, it moves, relative to the position of the buffalo, in such a way that the buffalo can't perceive the lion; i.e., in such a way that the buffalo can't detect the position of the lion so that it can respond in a way that will make it less likely that the lion will achieve *its* goal. But, of course, when the lion is sleeping, or

[1] Aristotle, trans. D. H. Hamlyn, *On the Soul*, in J.L. Ackrill, ed. *A New Aristotle Reader*, (Princeton: Princeton University Press, 1987), p. 169.

merely sleepy, or confronted with a whole herd of buffalo, or simply not hungry it can't be counted on to do any of this. So how the lion behaves depends to a considerable extent on the goal to be facilitated by that behavior.

And that tells us something about what the faculty of perception consists in. Not any old response to any kind of environmental feature signals perception. The ability to perceive is the ability to gather *information* regarding a certain kind of feature of an environment, so that the entity is able to respond to those features in a certain kind of fashion. When there is perception the entity responds differentially to features in the environment that are in some way relevant to the future occurrence of some state of affairs, so that the entity can respond appropriately to those environmental factors, that is, respond in ways that increase the probability that that end state occurs. We say that the lion perceives the position of the buffalo, for example, only because what the lion does, after the perception, makes it more likely that, given the actual position of the buffalo, the buffalo will wind up in the belly of the lion. Absent such a linkage with instrumentally appropriate behavior, we would be hard put to determine any instance of an effect of the environment on an entity as a perception.

The characterization I have given here of acts that occur in order to achieve some goal, that such acts raise the probability of the kind of possible state specified as the goal occurring at some future time, is far too simple to be accurate, of course. In fact, much of what agents do in order to achieve an end is counter-productive to reaching that goal. In *Rational Animals* I articulate a detailed analysis of the 'in order to' locution that treats the goal directedness of a behavior as a complex, holistic relation among a set of behaviors, an agent of those behaviors, and the environment in which those behaviors occur. On the view I articulate there, it is a logical requirement on an agent acting in order to achieve an end that most of what an agent does over time is successful at achieving its proximate goal. This, I believe, is the original principle of charity that under-girds the more familiar semantic forms of that principle. But it doesn't follow from this principle that everything that an agent does in order to reach a goal succeeds, that each kind of thing that an agent does in order to achieve some kind of goal succeeds most of the time,

or that everything an agent does to achieve a goal raises the probability that that goal will be achieved.[2]

The teleological structure in which perception is embedded is even more complicated than this simplified example suggests, of course. That the lion behaves in ways that, given the actual environment, increase the probability that the buffalo is ingested, is itself intelligible because ingesting buffalo, or eating them, is both conducive to the continuance of and part of an ongoing, repeating, pattern of events that constitute the life history of the lion. Lions not only respond differentially to their environment so as increase the probability that buffalo parts wind up in their stomachs. It is also the case that unless the right kind of stuff winds up in their stomachs they cease doing a whole variety of things necessary for that stuff winding up in their stomachs. That is, they die. In general, it is only things that behave in ways that, in their actual environments, keep such unlikely repeating patterns of events happening that we say are 'alive', and it is only living things that we say might have the capacity to perceive.

Aristotle's world was full of events that occurred in order to realize some end or *telos*. Plants, as well as the operations of the functional parts of organisms and the developmental paths of organisms, for example, in addition to the actions of organic agents, were events and processes that had ends or goals. That such entities, events and processes had ends that were objective features of the world was evident to Aristotle from the fact that appealing to their goals, or final causes, could help to explain these natural phenomena. And, for Aristotle, there was no requirement that these natural ends had their grounds in the representational capacities of agents. No one designs the developmental path of a fetus, and Aristotle knew this. It is still the

[2] (In *Rational Animals* I distinguish sharply between agents or systems, such as plants and simple animals such as wasps, that lack intentional states because they merely act in order to achieve a goal without being responsive to reasons, and animals who do act for reasons and have beliefs and desires with intentional content, such as ourselves and our friend the lion. Roughly, agents with intentional states are those that act flexibly in unique ways in response to varying and novel conditions in their environment so as to achieve their goals and can respond to their mistaken actions as mistakes by correcting them, while the other class of entities that act in order to achieve ends cannot do this. I argue that an entity can't belong to the category of rational agents without belonging to the category of simple teleological agents, but that the reverse is not the case. In the interests of brevity I am ignoring this important distinction here.)

case, however, that according to Aristotle that path occurs as it does *because* it happens in order to bring about a freestanding organism of a certain type.

On the other hand, for Aristotle there *are* a class of natural occurrences that happen in order to achieve ends, the acts of certain organic agents, for which we can offer an explanation of why it is that the event has the goal it does that appeals to the representational capacities of the agent. As opposed to well functioning hearts, which always act in order to pump blood, well functioning lions don't always act in order to capture buffalo. Sometimes, when a lion is in a certain situation relative to a buffalo she will hunt, or act in order to capture and kill the buffalo, and sometimes she won't. So there must be something about the lion itself that is different in the two cases. But that difference is just that in the first case there is something in the lion that motivates it to act in order to catch a buffalo. As we say, the lion *desires* that the buffalo be caught and killed. In the second case, this state, this desire, is absent. So the lion being hungry in part explains what the lion does. Now, since it is essential to desire that it partially explain why an agent acts with the goal it does, nothing can have a desire that does not act for an end or *telos*. But nothing can act for an end unless it is mostly successful at getting what it *needs* in order to stay alive. But this implies that an agent will mostly have the desires it ought to have, the desires it has a reason to have, given the kind of agent it is and the requirements on that kind of agent. Mick Jagger is an Aristotelian. The lion can't always get what it wants, but it had better be the case that, if it tries sometimes, it gets what it *needs*. In the Aristotelian tradition, the ends that agents desire are either appropriate or inappropriate, relative to the natural ends of the organism, and all desires are rooted in the natural ends of the organism, in the sense that without those ends there are no desires at all. It is not that there are no ends without desires. Rather, there are no desires without natural ends.

But, even if the lion desires to capture and kill a buffalo, what she will do in detail also depends upon another factor internal to the lion, how she perceives the world to be. In general, when it is hunting, that is, is acting in order to eat, or acting because it wants to eat, a lion that perceives a buffalo in one position will move in one direction, if it perceives it in another position it will move in a different direction, and if it fails to perceive a buffalo it will act in yet another way. How can we describe these parameters, these various positions and directions?

Clearly, the crucial aspect of the perceived position of the prey that makes a difference in how the lion behaves is neither the absolute position of the buffalo nor the relative position of the buffalo in respect to the lion determined in terms of the compass points of north, east, south, and west. Rather, the lion perceives those aspects of the situation that make a potential difference to whether or not (and how) the buffalo can be successfully hunted. That is, it looks as if the lion perceives instrumentally relevant aspects of the buffalo's position, or as that other great 20th century Aristotelian, J.J. Gibson, put it, the lion perceives *affordances*. If, for example, while hunting, the lion perceives that the buffalo is upwind from her, then all things equal she will crouch in high grass if she perceives it to be present; if she perceives that the buffalo is downwind, she won't do this, but, if she also perceives that she is as yet undetected, she might start moving in a circle around the buffalo; and if the lion fails to perceive the buffalo then we just can't predict how the lion will move.

We suspect that the lion's perceptions have these various contents (e.g., that the buffalo is down or up wind) based on the fact that, given that the lion is hunting, (i.e., has the goal of catching the buffalo so as to eat it) the lion acts as if it had access to just that information. The lion does just what it *ought* to do, what it has a reason to do, if it wanted to catch the buffalo, and the buffalo were down wind. That is, the lion attempts to conceal herself, *because* doing so helps her to catch the buffalo, *because* she has a reason to hide. And, since the success of concealment is dependent on the relative positions of lion and buffalo, when as frequently happens the lion succeeds in acting so as to conceal herself, she acts in that, successful, way, because she has *perceived* the instrumentally relevant relative position of the buffalo. What the lion perceives, the content of its act of perception, serves to explain why it is that the lion acts in the fine-grained fashion that it does. She crouches because, given that the buffalo is upwind, and that she has *perceived* that the buffalo is upwind, crouching is likely to be a successful act of concealment and the lion acts as she ought to given her perception of the way things are. She doesn't crouch because, given that the buffalo is downwind, and she has *perceived* that it is downwind, crouching is unlikely to be a successful act of concealment, and the lion acts as she ought to given her perception of the way things are.

It is important to notice that it is not that the buffalo is actually up or down wind from the lion that explains how the lion moves. If this was the explanatorily relevant factor, then we could explain the lion's behavior in simple Hempelian terms. But in fact as things stand we can't do this. The lion might not be hunting; or it might fail to perceive or misperceive the relevant features of the situation, and in all of those cases the buffalo being upwind from the lion could not be used to explain the lion's behavior. Rather, precisely because the lion doesn't always do what it has reason to do, given its goal, we need to appeal to a state of the lion in explaining its behavior. What explains how the successful agent acts is not the actual state of the world in which it finds itself. What explains how that agent acts is its acknowledgement, its perception, of the state of that world, together with its current desire that fixes its goal. Notice, on this model, the lion doesn't need to *apperceive* that it perceives the instrumentally relevant features of the position of the buffalo in order for it to be the case that it acts as it does *because* it has a reason to do so. It is the instrumental significance of the positioning of the buffalo, that is, the content of its perception, that supplies the lion with the reason that motivates its action, not the lion acknowledging, in a second order intention, that it has that reason. But this state of the agent, what it perceives, does have a systematic relationship with the actual state of the world. Since an agent acts in order to achieve goals only if it mostly succeeds in achieving them in its actual environment, any agent that acts to achieve goals will be mostly successful. As I mentioned above, the price of persistent failure is death, the suspension of the very attempt to succeed. But that the agent acts as it does is a function of its goals and its perception of the world. When it acts successfully, as it most often does, it acts as it has reason to given the actual situation in the world. But since what the agent does, whether successful or not, is always a partial function of how it perceives the environment, when the agent acts successfully, it does so because its perception of the environment 'corresponds' with the actual environment. And, since teleological agents mostly do what they ought, most of their perceptions are accurate. But since teleological agents also make mistakes, that is, act in ways that are not productive of their ends, this correspondence of reality and perception is only 'proximately and for the most part'.

The Aristotelian tradition thus takes the fact that certain entities, organic agents, act in order to achieve goals, as the primary key to understanding their behavior. From this perspective,

the activity of such agents resists the kind of explanation that appeals to subsumption under laws, statistical or non-statistical, because that activity crucially depends upon certain variable states of the agent. Sometimes the agent acts in order to achieve some end, and sometimes it doesn't. And when it does act in order to realize a goal, which goal it wants to achieve varies in ways that are characteristic of the kind of agent that it is. In acting to reach a goal, certain behaviors will facilitate success in a given environment; other possible behaviors do not have that property. There is a reason for an agent to engage in some behavior in some environment only if that behavior will help in the achievement of some goal of that agent in that environment. An agent itself has a reason of its own to act in a certain way in an environment only if that behavior would help in the achievement of some goal of that agent in the environment as the agent perceives it to be. Since, on this view, an agent only acts when it does what it does for reasons of its own, (organisms as a class are characterized by having internal principles of motion and rest), the activity of organisms can be explained by appealing to the agents' reasons. Since, on this view, agents mostly do what there is reason for them to do, agents' reasons mostly correspond to what they objectively ought to do, in order for the agent to achieve its goals. But absent a goal, there is no reason for an agent to do anything. Absent a goal, an agent can't have a reason for doing anything. And absent a goal, no agent can actually do something because it has a reason for doing it.

It is also crucial to notice that on this classical Aristotelian view what is explained by appeals to an agent's reasons are the acts of the agent, described as acts of that agent. That is, the explananda of rational explanations are the acts of the agent described in goal directed terms. The lion *hunts* because it wants to eat. The lion *conceals* herself from the buffalo, because she is hunting and she sees that the buffalo is upwind. This feature of the Aristotelian tradition distinguishes that tradition fundamentally from most standard forms of contemporary functionalism, which take the explananda of rational explanations to be behavior described in non-teleological terms. In addition to being committed to the goal directed character of the explananda of reason giving explanations, the Aristotelian tradition is also implicitly committed to a certain position regarding the character of the internal states that are appealed to in offering explanations that appeal to an agent's reasons. As agents can act for reasons only if they act in

order to achieve goals, and acting in order to achieve a goal is associated with an internal state of the agent, the agent's desire, that desire can be individuated and distinguished from other desires by identifying it as the desire to achieve the particular goal that the agent acts in order to achieve. The *content* of an agent's desire is both explanatory of and indexed to the goal of the agent's activity. Since the proximate goal of any such activity is always embedded in a broader holistic context of the life activity of the agent, however, this content has an intensional dimension. A lion can want to eat the meat, without wanting to eat the poison, even if the meat is the poison. The lion wants the item *as* meat, but not *as* poison, as Aristotle would put the point. Similarly, in this case the content of the lion's perception is that it is meat, not that it is poison, and the lion eats the meat *qua* meat, not *qua* poison. For someone in this tradition there is intentionality, and intensionality, because agents do things for reasons, and there are reasons only because actions are goal-directed. Intentionality has teleological roots.

THE KANTIAN TRADITION

"Everything in nature works in accordance with laws. Only a rational being has the capacity to act in accordance with the representation of laws, that is, in accordance with principles, or has a will. Since reason is required for the derivation of actions from laws, the will is nothing other than practical reason."[3]

"Actions just are performances for which it is appropriate to offer reasons, and offering a reason is making an assertion. So actions are not intelligible as such except in a context that includes assertional giving of reasons. Where intentional explanations are offered of the behavior of nonlinguistic creatures (those that are not understood as interpreters of others), the reasons are offered, the assertions are made, by the interpreter of the simple intentional system, who seeks to make its behavior intelligible by treating it *as if* it could act according to reasons it offers itself. That is why what is attributed in such interpretations is *derivative* intentionality."[4]

In a modern context, "final cause" is a poor translation of the term that Aristotle uses when he speaks of the goal of a process that can be appealed to in explanations of that process. That term is better translated as 'the for-the-sake-of'. What is misleading about 'final cause' is, of course, that that for the sake of which something happens is not a cause in our sense at all, and

[3] I. Kant, *Groundwork of the Metaphysics of Morals*, trans. M. Gregor, (Cambridge: Cambridge University Press, 1997), p. 24.
[4] Robert Brandom, *Making It Explicit* (Cambridge: Harvard University Press, 1994), p. 171.

could not possibly be one. It is a possible state of affairs that is to be brought about by the event or process it explains, and, as such, it is no actual event that could enter into causal relations with anything. Both in ordinary language and for an Aristotelian analysis of that language goals are invoked in order to explain an event, not to provide a cause for that event. But, in the context established by the scientific revolution of the 17th century, this presented a serious problem for the project of explaining events or processes by appealing to their goals. And this in turn created a crisis in the project of understanding how it is that agents, human or otherwise, could possibly act because of reasons.

As I mentioned above, Aristotle thought that agents who acted in order to achieve goals because of their desires and perceptions were not the only entities whose behavior could be explained by appealing to goals. The operations of the parts of organisms, and the developmental path of those same organisms, as well as plants and simple animals, could also be explained by their 'final causes'. For Aristotle goals, as necessary explanatory principles, were objective features of the world that could exist as goals independently of being represented as such in the mind of any cognitively equipped agent. Aquinas, as a Christian influenced by Semitic ideas, could argue that, since things in nature act for ends, and nothing can act for an end unless it is designed to do so, and nothing in nature designs natural entities, such entities must be designed by a supernatural being. Aristotle could never accept this argument, because he denies the second premise; that nothing can act for an end unless it is designed to do so. In the aftermath of the scientific revolution, with its emphasis on mechanistic causality, it became difficult to accept this argument for a different reason. Once the new science had established itself it became difficult to accept the premise that things in nature act for ends. After all, as Kant puts it, "everything in nature operates in accordance with law". That is, everything in nature, including the actions of organic agents, can be explained in Hempelian fashion through an appeal to natural laws that specify the causal antecedent of the event that is to be explained. And if this is the case, then explanatory appeals to goals become otiose.

As is evident from numerous discussions in the 17th century[5] this new context created a crisis for the understanding of the concept of an action. In a traditional Aristotelian context, an act is an event that occurs *because* it helps to bring about some goal of the agent. If it occurs because it is instigated by something external to those goals, then the agent is a mere patient. The lion who hunts acts, because what it does while hunting happens because it is for the sake of the lion eating. When the lion is shot, this is not the case. But if, as seemed obvious after the 17th century, nothing in nature happens because it is in order to bring about a goal, then nothing in nature acts. Since, however, the correlated concepts of an act and an agent are central and fundamental to our ways of coping with the world and (even more fundamentally) ourselves and other human being, the fact that these concepts appeared to have no purchase in the brave new world provoked a philosophical crisis. And, while it is anything but obvious that this constellation of factors *in fact* present a real and serious obstacle to the employment of teleological explanations (I argue that they do not), or even that all 17th century thinkers thought that they do, it is uncontroversial that the rise of the new science provoked a crisis for teleological thought.

This set of problems, and the rationalist and empiricist responses to these problems, supplied the context for Kant's work in practical reason. The single most significant aspect of that context was the problem of how to understand the normative force of reasons in the absence of natural ends. Kant's solution to that problem was radically anti-teleological. Real action, that is, behavior that really happens because of reasons, is not action that is for the sake of achieving an end or goal at all. Instead, Kant took as his starting point the Rationalist insights that action is behavior that happens because of reasons, and that reasons are essentially premises in an argument that can justify claims regarding what the agent ought to do. He argued that as long as the reason for acting is to reach a goal, the justification or reason for the act must depend, ultimately, upon contingent, non-reason based, facts about the agent's psychological make-up. Because of this, even if there is an apparently inferential element to the link between an agent's desires and her behavior, such behavior falls under Hempelian laws and can be explained by the

[5] For example, see Leibniz' discussion of the concept of an act in *Discourse on Metaphysics* beginning in Section 8.

fact that the agent has the desires that, as a contingent matter of fact, she has. Because of this contingency and lack of justification, 'oughts' that follow from having a certain goal, hypothetical imperatives, are not genuinely obligatory and lack genuine normativity. Only premises for practical reasoning that can themselves be justified without appeal to the ends to be accomplished by the act could possibly provide a real reason for acting. And, since an act is a piece of behavior that happens because of the agent's reasons, in a crucial sense, no behavior that happens in order to achieve a goal is an act at all.

Kant's alternative to goal oriented practical reasoning was designed to supply a premise for that reasoning that was truly normative, necessary (that is, justified or non-contingent) and objective. To be objective, the premise needed to hold universally for all agents, rather than to be based on the subjective preferences of a particular agent. To be truly normative, the premise needed to specify what an agent ought to do, independent of any qualification. Together, these two requirements fix the form of the premise. The first premise of practical reason must have the form of a normative universal generalization that is true necessarily rather than contingently. That is, it must have the form of a law that, as opposed to a Hempelian covering law, specifies how any agent is to act, rather than specifying how any agent will act. The content of this first premise, on the other hand, is fixed by Kant's analysis of what is involved in acting for a reason at all. As the premise from which the action follows is not a law that specifies how things actually act, but only how they ought to act, one can't infer from the premise to the conclusion that any agent will act in any particular way. If an agent *actually* acts because of this reason, because of this premise, then it can't be because she causally could not have done otherwise. Rather, for the agent to act because of the reason is for her to act because she has *accepted* the law as binding on herself because she is an agent, as one from which *any* agent *ought* to reason. But to be binding of any agent in that way, the content of the law must be such that it *could* apply to any agent as such, which is not the case of any principle rooted in the contingent desires of particular agents. A universal generalization that can possibly apply to all agents, however, is a principle that could be a universal law that could actually hold of all real agents. And this provides the content of the first premise of practical reason: Always act so that the maxim (major premise) that provides the reason for your act could be a universal law of nature.

The problem with the content of Kant's Categorical Imperative, that it is too abstract to be an actual premise in practical reasoning, is well known and will become relevant to my discussion soon enough. At this point, however, I want to focus on a singular feature of Kant's remarkable analysis of what it is to act for a reason. On Kant's view, if an agent acts for a reason she does not act because she falls under some Hempelian covering law. Rather, when an agent acts for a reason she acts because she acknowledges some principle, some universal generalization, as applicable or binding on herself. As Brandom puts the point: "What makes us act as we do is not the rule itself but our *acknowledgment* of rules."[6] This analysis of acting for reasons has three crucial consequences that together determine how those in the Kantian tradition conceive the nature of intentional states. First, actions that are performed because of reasons, that is, all real action, is mediated by a *representation* of a law by the agent. Second, since laws can only be represented discursively, that is, linguistically, only agents who are capable of language use are capable of acting for reasons. Third, since when an agent acts for reasons she does so because she has *inferred* that she ought to do so from a normative principle treated as a premise, only agents that are capable of explicit inference from universal premises are capable of acting for reasons, and thus acting at all.

These conclusions, together with the Aristotelian premise that the intentional content of internal states of agents is fixed by their roles in giving rational explanations of the behavior of the agent, effectively imply that, on the Kantian view, non-linguistic agents can neither act nor really have any intentional states. Strictly speaking, the lion does not want to capture the buffalo, although it is sometimes convenient for us to speak metaphorically as if she did. As far as I know, Kant himself never explicitly developed this argument. He could not, because while he inherited certain Aristotelian tendencies from Leibniz, he seems to have held to the dominant view of his age that the intentional characteristics of mental states are intrinsic features of those states, although they can be appealed to in rational explanations. As this conviction eroded in the course of the 20th century, many more contemporary Kantians have repeatedly and explicitly used this argument. For example, in one of the quotes with which I began this section Robert Brandom hints at this Kantian argument to the conclusion that non-linguistic agents have no

[6] Ibid, p.31.

intentional states. Although the entire argument is not explicitly stated in the cited passage, it is relatively easy to reconstruct it from the rest of Brandom's work. (Essentially the same argument also appears in Sellars and Davidson, among others.) Here, in condensed form, is the argument: 1. Intentional states of agents are the states of those agents that can serve essential roles in rational explanations of the actions of those agents. 2. The content of an intentional state is fixed by the specific role of that state in providing rational explanations. 3. Rational explanations appeal to the practical rationality of agents, and an agent is practically rational just in case it is motivated by its acknowledgement of norms and by the conclusions it infers from those norms (which serve as major premises, or 'principles', of its practical inferences). This is the premise that is derived from the analysis of the Kantian vision of practical reason rehearsed above. 4. An agent can acknowledge a norm or rule just in case it can make that norm explicit to itself. 5. A norm can be made explicit to an agent only if that norm is linguistically articulated. 6. No non-linguistic agents are practically rational (from 3, 4, and 5). 7. No non-linguistic agents have states that have (real, non-derivative) intentionality (from 1, 2, and 6).[7]

[7] It is important to note that this argument is quite distinct from another argument that is suggested by the first two sentences of the original quote that I cited from Brandom. According to that other, spurious, argument, the reason that non-linguistic creatures can't have practical rationality, and thus must lack (real, non-derivative) intentionality, is that actions as such are 'not intelligible' outside of a context of assertions that are offered as reasons for the actions, and therefore creatures who lack the ability to make assertions can not act for reasons. To see that this other argument is fallacious, consider the following. Let's say, for the sake of discussion, that Brandom is right in asserting that "Actions just are performances for which it is appropriate to offer reasons, and offering a reason is making an assertion." (I myself am inclined to deny the second part of this claim, but that position is irrelevant in the current context.) That is, let us assume that what it is to be an action is to be an event for which it is appropriate to offer a reason by way of explanation, and such reasons that are offered must be assertions. From these premises it really does follow that, as Brandom puts it, "actions are not intelligible as such except in a context that includes assertional giving of reasons". The reason for this is that 'action' is *defined* in terms of 'reason', and what it is to be a reason is *defined* in terms of 'assertion'. So for one who did not understand what it is to be an assertion it would be impossible to understand what it is to be an action. Action would be unintelligible as such to such a non-linguistic creature. But from this it does not follow that such a creature would be *incapable* of *acting* for reasons, and thus incapable of intentionality. All that would follow is that such a creature could not intend any creature, itself or another, as acting, and thus as intentional. That is, what follows is that such a creature would be incapable of second order intentions. But lacking a further premise, none of this implies that non-linguistic creatures lack intentionality *tout court*. On the proffered

Two singular aspects of Kant's original suggestion dominated the development of the Kantian strategy in the 19th and 20th centuries. First, in accordance with Kant's views regarding

---

definitions, for an agent to have non-derivative intentionality all that is required is that that agent engage in some performance for which it is appropriate to offer a reason, not that that agent actually offers that reason, or is capable of offering that reason. The way in which 'action' is here defined is that an event is an action if it is *appropriate* to offer a reason in explaining it. But 'being appropriate' does not imply 'being actual'. It can be appropriate for A to do B without A, or anything else, ever doing B. When my dog drops a chewy toy at my feet and nuzzles me until I pick it up and throw it in a game of fetch, he does something for which it is *prima facie* appropriate to offer a reason, that he wants to play fetch with me and believes that cuing me in this way will result in my playing. Because it is appropriate to offer this reason, I sometimes do offer it in explaining the performance, although my dog never does. And it remains appropriate, in this sense, to offer such a reason even if no one were ever around who was capable of offering it. So, given the definition of 'action' on offer, it is correct to say that my non-linguistic dog acts, and thus is capable of intentional states, regardless of whether or not he is capable of ascribing those reasons to himself. The distinction here is just the one that Brandom himself makes in a different context, between sense dependence and reference dependence. Compare: What it is to be pleasant is to be an object that would cause pleasure in organisms like us were such organisms to encounter the object. The concept of the pleasant is thus unintelligible to a being that doesn't understand what it is to be pleasure for a being like us. So a being with no understanding of beings like us cannot understand the concept of the pleasant. But of course it doesn't follow from this that something can't be pleasant if beings like us don't exist. For in that case the things that are currently pleasant would also be pleasant, because they *would* cause us pleasure if we were to exist and if we were to encounter them. Similarly, an agent can act, be rational and have intentional states even in the case where to act is to engage in a performance for which it is appropriate to offer a reason, even in the counterfactual circumstance that there has never been any being that offers reasons, even though in that case there would be no one for whom those states were intelligible. Lacking a further premise, the most this line of argument could provide, then, is that intentionality is sense dependent on assertion. One could not infer reference dependence. An agent could act even if what it is to act is to engage in performances for which it is appropriate to offer reasons in the form of assertions and there are no agents who are capable of offering those assertions or reasons.

What is required to complete the argument, of course, is a further limitation on and specification of what is involved in it being 'appropriate' to offer a reason for a performance. And this further limitation is just what is provided by the Kantian/Brandomian analysis of practical reason. According to this analysis it is 'appropriate' to offer a reason for an event, and thus consider that event to be an action, only if the etiology of that act is mediated by the agent offering itself a reason and the explicit acknowledgment of the norm or principle embedded in that reason as binding on herself by the agent, an acknowledgment that itself is only realized in and through the acceptance or making of an assertion. So it is never appropriate to offer a reason for the events associated with a non-linguistic organism, because such organisms can never do anything *because* they have a reason for doing so.

the proper way to understand conceptual representation, as rules, the major premises for practical reasoning were understood as rules specifying behavior as appropriate and inappropriate in light of the rule. Just as the rules of chess specify that if something is a bishop, then it is only allowable and appropriate to move it on a diagonal, the major premises of practical reasoning specify that if one is an agent, it is only allowable and appropriate for one to act so that the maxim of her action could be a law of nature. In the Aristotelian context, an agent has a reason to act if that act is genuinely helpful in achieving some goal. In the Kantian context, an agent has a reason to act if she acknowledges some rule from which the action can be derived. The second aspect of the Kantian position that dominated its development in the 19th and 20th centuries was the evident abstractness of Kant's own suggestion regarding the content of the major premises of practical reasoning, pointed out most forcefully by Hegel. There are many principles that could consistently be laws of nature, and the Kantian suggestion gives one no guidance on which of those rules are genuinely obligatory. The solution to this problem that came to dominate was suggested by the rule character of the Kantian major premises of practical reasoning. According to this view, the major premises of practical reason are norms of appropriate behavior for some *type* of agent that come to have normative force within some community of agents.

For example, the social role of parent in this society has a *normative*, as well as a descriptive, dimension. The descriptive aspect of the role is quite complicated as it involves both biological and non-biological conditions, none of which are necessary or sufficient, as is perfectly evident to me since one of my four children is adopted. But whatever the descriptive conditions that qualifies an agent as a parent, if one is a parent in this society this implies that one *ought* to act in certain ways. And it is at least possible that when I act in a loving fashion towards my teenage daughter in the face of her provocation I do so for the reason that I acknowledge that I am a parent and also accept the maxim that if one is a parent then one ought to act in a loving fashion towards one's children.

The fact that the Kantian tradition came to understand all norms in terms of rules, and came to understand rules in terms of socially appropriate cultural practices, raised a whole host of problems as well as a whole range of opportunities. In the current context I want to emphasize the way in which these developments affected the Kantian understanding of intentionality. As the

quote from Brandom makes obvious, the basic Kantian model requires that if an agent acts for reasons, she must be capable of discursively articulating that reason to herself; she must be capable of understanding and using language. But, to be able to speak or to understand the speech of others one must be able to interpret the linguistic forms in such a way that one distinguishes between correct and incorrect usage. That is, language use itself is inherently normative; an agent is using language only if she does what she linguistically does for the reason that it is right, it is appropriate, to act as she linguistically does. This fact, together with the Kantian analysis of acting for a reason, and the social twist on that analysis, that the rules that one follows when one acts for reasons are to be understood as socially prescribed norms, imply that only entities that are members of social communities that have instituted discursive norms for appropriate linguistic behavior are capable of acting for reasons in any context, linguistic or non-linguistic. And, given the dominant association of the contents of desires and beliefs with an agent's reasons for acting, it follows that only agents who live in these sorts of linguistic communities can exhibit possess beliefs and desires at all.

But this result provokes an obvious question. If only entities that are members of social communities that have instituted discursive norms for appropriate linguistic behavior are capable of intentional states, what is it for it to be the case that a community has instituted a norm of appropriate linguistic behavior and for it to be the case that agents within such a community are acting because they acknowledge these linguistic norms? This is the question that I will use as a wedge to transition to the next section, in which I will try to show why it is that acting for reasons in a Kantian fashion, by following a rule, cannot be autonomous, and that only agents that act teleologically in order to achieve a *telos* can also act in virtue of acknowledging a principle.

ON WHY THERE ARE NO REASONS WITHOUT GOALS

The acceptable answers to the question of what is involved in an agent acting linguistically because of her acceptance of a norm instituted within her community are limited by a deep problem first articulated by the late Wittgenstein. It just can't be the case that, *in general*, an agent stands under a norm just in case that agent acknowledges the representation of the rule expressing that norm, and it just can't be the case that an agent belongs to a community of agents

that all stand under some norm just in case she acknowledges the representation of the rule expressing that norm. As Wittgenstein pointed out, no principle or rule interprets itself. There are always multiple ways in which any rule for action can be applied in particular circumstances, consistent with past cases. If it is the case that an agent stands under a norm and acts for a reason only if she acknowledges an explicit reason or rule that she offers to herself, then if she is to act rationally when she applies some principle to a given case in one way rather than in some other, deviant fashion, then that agent must have some explicit reason, or principle, for choosing that application of the rule rather than the rejected interpretation. But that meta-principle, whatever it is, also needs to be interpreted. And, given the view of rational agency and liability for normative assessment on the table, for the agent to act rationally in applying that principle or to stand under that rule, she would need to accept some further principle that she offers to herself from which she could infer the correctness of one application of that first meta-principle rather than others. And so the regress goes.

The implications of these considerations would seem to be obvious. It just can't be the case that it is a necessary condition on an agent acting for a reason, that that agent offers herself and acknowledges that reason. At least some acts are performed for reasons that are not explicitly represented or acknowledged by the agent of those acts. But if this is the case, then there must be some way for an agent to act because of reasons that does not involve any acknowledgment of those reasons. The Kantian analysis of practical rationality can't be an adequate analysis of what it is to act for a reason generally.

It is obvious in the light of these considerations that if the Kantian, rationalist analysis of what it is to act for a reason is treated as a general analysis of acting for a reason that is capable of defining what it is for an agent to be subject to a norm or provide a criterion for membership in a normative community, then the analysis is fatally flawed. There must be a way for an agent to act for a reason and be subject to a norm without that agent explicitly acknowledging and representing that norm, even if much of the time acting for a reason can be understood in Kantian terms. If being subject to a norm, being part of a normative group, and acting for a reason is ever to occur in the way that Kant suggests, in and through acknowledgment of principles, there must

be some other way in which an agent can be subject to a norm, be part of a normative group, and act for a reason, a way that does not involve such a recognition of rule.

Given these considerations, one might think that an agent can act for a reason and be subject to a norm only if she belongs to a group all of whose members respond in the same fashion in the same circumstances, perhaps as a result of a process of training instituted within the group. After all, even if it can't be the case that, for example, people around here are responding to the acknowledgment of some explicit principle when they say 'blue' in response to the question, 'What color is that shirt?', it does seem to be a noteworthy fact that most of us have been trained to respond in this way, and this fact and those like it do seem to be necessary for us to apply linguistic rules at all. So perhaps the kind of non-reflective responsiveness to norms in our actions that is necessary if we are to act for explicit reasons can be understood in terms of simple uniformity of action within a group. a uniformity that perhaps arises as the result of a process of training in which, at most, the teacher, but not the learner, represents the rule of the regularity to herself.

In its crude form the suggestion that simple uniformity of action within a group according to some principle is a necessary condition on acting for a reason is also an obvious non-starter. Mere regularity of behavior can't be a criterion for implicitly standing under a normative rule. Someone can be a parent and thus subject to the norms of parenting in a society, without acting as a parent ought to act. Someone can act, as a parent ought to act in some society, without acting that way for the reason that she is a parent. In general the normative structure of acting for a reason requires that it be the case that an agent ought to act in some fashion or can have a reason to act in some fashion, and still fail to do so. But if what it were to have a reason is partially defined in terms of acting according to rule, then this would be impossible. Similarly, it isn't even the case that the rule in question must be followed most of the time by an agent for that agent to stand under the norm implied by that rule. There are some notoriously bad parents out there, and to evaluate them as bad in that respect is to judge them as subject to the norms of parenting. Some ideals, in some communities, are rarely if ever realized by anyone in that community, and for all that those ideals still function as norms in the community. The problem is that, while some uniformity of behavior within a group seems to be a necessary condition for

acting for a reason, or because of a norm, mere uniformity of behavior within a group is never sufficient by itself to constitute even implicitly being responsive to a reason or being subject to a norm.

On the other hand, the second aspect of the above suggestion, that implicitly following a rule without explicit representation of that rule is best understood as being the result of a process of training producing a rough uniformity of behavior within a group, is much more promising. In fact, a group of latter-day Kantians, the Sellarsians, has made this suggestion basic to their understanding of acting for a reason. The move has two stages. First, if a teacher already has the ability to linguistically represent a rule to herself, then she can make it a principle of her action to train non-linguistic learners so that they come to follow just that rule. The result of that training will then be that the non-linguistic learner comes to (mostly) act according to the rule, *because it is that rule*, even though she cannot represent the rule to herself. And this, the Sellarsians hold, is the way to understand how it is possible to act because of a reason even when the agent cannot acknowledge that reason explicitly. But this solution leaves the residual problem of understanding how it is possible that the practice of acting according to rules could ever arise in the first place. Since there was certainly a time at which no one could represent rules to herself, the notion of a training regime is inadequate to account for the origin of acting for a reason. The second aspect of the Sellarsian suggestion, then, is to generalize the notion of training to include assessment of individual behavior within a group by other members of the group. Even if non-linguistic agents can't train each other, they can assess each other and this assessment can play a crucial role in the development of rule following acting for a reason. On this view, for an agent to implicitly stand under a norm, without that agent having been trained to do so by a linguistic agent who represents that norm, is for that agent to be assessed by other agents in a group as abiding, or failing to abide, by that norm, and for that agent to act as she does because of that assessment. While as far as I know this suggestion was first advanced by Sellars himself in "Some Reflections on Language Games"[8], by far the most well articulated view along these lines has been developed by Brandom. "…there is another move available for understanding what it is for norms to be implicit in practices. This is to look not just at what is done – the performances

[8] Philosophy of Science, Vol. 21, No. 3 (Jul., 1954), pp. 204-228

that might or might not accord with a norm (appropriate or inappropriate) – but also at *assessments* of propriety. These are attitudes of taking or treating performances *as* correct or incorrect."[9] "…one way to demystify norms is to understand them as *instituted* by the practical attitudes of those who acknowledge them in their practice. Apart from such practical acknowledgment – taking or treating performances as correct or incorrect by responding to them as such in practice – performances have natural properties, but not normative proprieties; they cannot be understood as correct or incorrect without reference to their assessment or acknowledgment as such by those in whose practice the norms are implicit."[10]

The suggestion is straightforward enough. The basic idea is that, even if the rules that on a Kantian model provide agents with a reason for acting are not explicitly represented in either a teacher or a learner, those rules can still be instituted, as Brandom says, by a process of social enforcement within a group, a process that implicitly enforces the rule that can be made explicit once the group develops an explicitly discursive practice. An action is implicitly incorrect, if it is *corrected,* by the members of the group, even if no agent in the group represents the principle under which the action is incorrect. An action is correct, if it is *enforced,* by the members of the group, even if no agent in the group represents the principle under which it is correct. The pattern of instances of correction and enforcement establish the boundaries of correct and incorrect behavior, and insofar as this pattern of correction is applied to an individual agent that agent is subject to the norm established by that pattern. If the agent in fact acts in conformity to the pattern established by the uniformities of enforcement, and does so because he has been subject to that process of enforcement, then one of the conditions for her implicitly acting because of the rule established by that pattern of enforcement is satisfied. If in addition the agent *enforces* this same rule in her assessments of others, then the agent is a full member of the normative community. The rough uniformity of behavior within the community that is necessary for acting in virtue of a reason is a result of this process of assessment, and this process of implicit assessment is also capable of instituting a rough uniformity of *verbal* behavior within the community. When the community in this way becomes capable of using the if-then they also

[9] R. Brandom, *Making It Explicit* (Cambridge, Harvard University Press, 1994), p. 63.
[10] Ibid.

become capable of representing the rule that was implicitly used in the group training process and that has been internalized by the group members. Then, for the first time these agents can act because that rule is their reason for acting, and they come to have intentional states. Intentionality enters the world with the 'if-then' linguistic form.

Now it seems to me that intuitively something like this view must be right. Social proprieties have a great deal to do with the behavior that the members of the society treat as correct and incorrect. In general, an act is socially acceptable if it is accepted by one's peers; it is unacceptable if it is corrected by one's peers. Beyond this, whatever else is involved in acting for an explicit reason in Kantian fashion, it is at least necessary that any agent that acts for an explicit reason must be able to represent that reason to herself linguistically. And whatever else is involved in linguistic representation, (and I would argue that a fair amount of instrumental rationality must be involved, but that is a topic for another time), linguistic competence does require adhering to social proprieties. So social assessment of correctness and incorrectness is a necessary condition on Kantian rationality, just as Brandom asserts. But as we have seen assessment comes in two forms. When a teacher or assessor is capable of representing the principle that she is using to assess the behavior of a learner, the act of assessment can be assimilated to the Kantian paradigm of acting for a reason. But this assimilation is impossible in the cases of assessment in which such representation by the assessor is impossible. And, given that the capacity to linguistically represent principles of action must have arisen in an animal community that initially lacked this capacity, that this capacity to represent linguistically necessarily involves the application of norms, and the Brandomian insistence that all normativity enters the world through a process of social assessment, there must be a form of social assessment of correctness and incorrectness of behavior that does not involve appeal to a representation of a rule. But what form can this implicit assessment of the correctness and incorrectness of behavior take?

This is an important question, as whatever is necessary for practical assessment will also be necessary for Kantian rationality. The obvious answer to the question refers to punishment and reward. Since it is the process of assessment that institutes the contours of the rule, that process must in general have enough teeth to affect the performance of the agent. But what is it

to reward or punish? Once again the answer seems obvious. One rewards and punishes when one acts in ways that in fact positively or negatively reinforces the behavior. "What counts as reward or punishment might be construed naturalistically, for instance as any response that positively or negatively reinforces the behavior responded to."[11] Significantly, Brandom himself does not accept this obvious answer to the question of the character of assessment, however. Nor does he reject it. And it is good for his program that he does neither. The reason that it is good for Brandom's program that he does neither is that this obvious answer to the question of what is involved in implicit assessment of behavior just *can't* be right. But the Kantian tradition in general, and Brandom in particular, has no other answer to the question available.

The problems with the 'naturalistic' notion of reinforcement, which in this context is the *behaviorist* notion of reinforcement, are the problems with behaviorism in general as a theory of the psychology of complex organisms. As a scientific theory of psychology, behaviorism has generally been replaced by cognitivism because it has proved impossible to empirically explain the behavior of complex organisms solely by appealing to schedules of reinforcement, without alluding to the intentional states, the beliefs and desires, of the organic agents. But, of course, since the Kantian denies that non-linguistic agents have any intentional states, they can't appeal to beliefs and desires either in a genetic account of how normativity and intentionality enter the world or in an analysis of the necessary conditions on acting for a reason or having an intentional state. If an agent has intentional states only if she can act for a reason, and she can act for a reason only if she belongs to a community in which norms are established by a process of community assessment, and reward and punishment are necessary for a process of community assessment, then reward and punishment had best not be defined in terms of the desires and beliefs of either the patients or the agents of reward and punishment. On the other hand, treated as part of an analysis of implicit normative assessment, the behaviorist notion of reinforcement is causally too strong, and because of that, normatively too weak. The problem here is analogous to the dread disjunction problem that plagues causal theories of intentional content. Using Brandom's terms, let us accept that we define punishment as any response that in fact negatively reinforces some behavior, and reward as any response that positively reinforces some behavior.

[11] Ibid.

By hypothesis, what the agent actually does after a process of reinforcement has been a result of that process. So, in principle, whatever the agent does after such a process of assessment is in accordance with how the members of society have punished and rewarded her past behavior. And this is true of the actual behavior of every member of the community. It follows from this, of course, that whatever any member of the community does is in accordance with the rule established within the community, even when that behavior is further punished by the other members of the community. (Of course, as a piece of behavior that happens as a result of social reinforcement, that punishment is also in accordance with rule.) If the only resource available to define reward and punishment is behavioral reinforcement, the normative character of rules breaks down.

It is perhaps for this reason that Brandom fails to explicitly endorse the proposal that implicit assessment is to be understood 'naturalistically', that is, behaviorally. But he also declines to reject the suggestion. The only other option he mentions is that reward and punishment are to be understood normatively, "for instance in terms of the granting of special rights or the assignment of special obligations".[12] As it stands, without support from the behavioral model, this proposal has the problem that it depends upon an unredeemed normative bank draft. All norms are to be understood in terms of implicit assessment. Implicit assessment in practice is impossible without reward and punishment. There is no reward and punishment in practice without – rights and obligations, where there is no attempt to articulate what it is for there to be a process of assessment that institutes rights and obligations.

This problem with articulating what is involved in assessment that is implicit in behavior but does not depend upon the representation of a principle, because one has no resources for understanding what it is to reward and punish apart from a reference to the intentional states of an assessor, leaves the Kantian tradition with no resources for understanding how its preferred model of acting for a reason is so much as possible. The problem, of course, arises from the fact that rewarding and punishing are, doubly and irredeemably, teleological notions. It is of course possible for intentional agents or environmental factors to influence the future behavior of an entity even if that entity lacks intentionality. But that influence counts as reward and punishment,

[12] Ibid.

and the entity influenced counts as implicitly learning, only in the context provided by that entity having goals. Some agent is rewarded only if she gets what she wants, (or, in simpler organisms, what it needs), even if what she wants isn't good for her, and even if getting what she wants doesn't in this instance reinforce her previous behavior. Some agent is punished only if she gets something she wants to avoid. Because of this structure, some agent can assess some other agent negatively by *attempting* to punish that other agent, and failing to do so, because the assessor has false beliefs regarding what the actor wants. Some critic can assess someone else positively by attempting to reward, and failing to do so, because of false beliefs. So it is possible that an agent can be subject to a process of assessment carried out through a process of reward and punishment and it still be the case that the learner fails to act as she ought, given the norms instituted by community assessment. And this is just what needs to be the case, if such institutions establish genuine norms.

Now, given these facts, it is certainly possible that there is a community of teleological agents, such as humans, that attempt to affect each other's behavior by attempting to reward and punish. Contingently, it is also possible that such attempts turn out to be pretty successful in punishing and rewarding. And, given the actual facts of human psychology, most such punishment and reward will in fact establish behavioral patterns in accordance with the norms implicit in these acts of assessment. So in these cases the attempt to act so as to reward and punish in general brings about the rough uniformity of behavior necessary for explicit action according to the representation of a rule. But the fact remains that since an agent is only capable of action according to the representation of rule if the agent also acts in such a way that actions in accord with the rules represented are implicit in her practice, and there is only such normativity implicit in the agent's practice if she is embedded in a community in which norms are implicit in mutual acts of assessment, and there is only such assessment if the agents in the society act *in order to* reward and punish, it follows that only agents capable of acting in teleological fashion in order to achieve ends are capable of acting in a Kantian fashion because they acknowledge a representation of a rule. And the reverse is not the case. So acting for a reason in Kantian fashion is not an autonomous ability. Only organic agents that act in order to achieve ends, and because of that, act because in the circumstances those goals give them reason

to act, can act because of the representation of law. Aristotelian action in order to reach goals is at the root of Kantian action according to principle. Where there is not the first, there is not the second. And, since agents that act for reasons are those agents that have intentional states, Aristotelian non-linguistic agents, as well as Kantian discursive agents, have intentional states. When my dog reaches for his chewy toy, and when I reach for my cup, we both display the behavioral marks of intentionality.

CONCLUDING UNSCIENTIFIC POSTSCRIPT

If it is genuinely the case that an agent acting in order to achieve a goal is an asymmetrically necessary condition on that agent acting for an explicitly represented reason, as I maintain, then one can't analyze acting in order to achieve a goal in terms of acting in light of linguistically represented reasons. This leaves us with two pressing problems. First, what is it for an agent to act in a goal directed fashion, if it is not for the agent to act in a way that is motivated by its representations? Second, related to this, given the causal order of the world that has been progressively revealed since the 17th century, how is the explanation of acts by appeal to goals compatible with causal accounts of these very same events?

In my recent book, *Rational Animals: The Teleological Roots of Intentionality* I attempt to answer just these questions. After having struggled to come up with something informative but not misleading to say about my answers to these questions in the minus five minutes I have left to me, I'm afraid that the best I can do is: 'Read the book."