

**The Adaptive Importance of Cognitive Efficiency:  
An Alternative Theory of Why We Have Beliefs and Desires**

Armin Schulz

London School of Economics and Political Science

---

**Abstract**

Finding out why we have beliefs and desires is important for a thorough understanding of the nature of our minds (and those of other animals). It is therefore unsurprising that several accounts have been presented that are meant to answer this question. At least in the philosophical literature, the most widely accepted of these are due to Kim Sterelny and Peter Godfrey-Smith, who argue that beliefs and desires evolved due to their enabling us to be behaviourally flexible in a way that reflexes do not – which, they claim, is beneficial in epistemically complex environments. However, as I try to make clear in this paper, upon closer consideration, this kind of account turns out to be theoretically implausible. In the main, this is because it fails to give due credit to the powers of reflex-driven organisms, which can in fact be just as flexible in their behaviour as ones that are belief / desire-driven. In order to improve on this account, I therefore propose that beliefs and desires evolved, not due to their enabling us to do something completely different from what reflexive organisms can do, but rather due to their enabling us to do the same things better. Specifically, I argue that beliefs and desires evolved for making the generation of behaviour more efficient, since they can simplify the necessary cognitive labour considerably. I end by considering various implications of this account.

---

**The Adaptive Importance of Cognitive Efficiency:  
An Alternative Theory of Why We Have Beliefs and Desires**

**I. Introduction**

It is fairly obvious that much of our behaviour is driven by beliefs and desires; equally, it is fairly obvious that this is not true for many other organisms – a large number of animals (not to mention plants, plankton and bacteria) seem to act in fundamentally different ways.<sup>1</sup> This raises the question as to *why* we have beliefs and desires: given that not all living creatures have them, why do *we* have them? What makes it worthwhile for us to use them when determining how to behave, when it does not seem to be worthwhile for orchids, annelids, and starfish?

Among the various accounts that have been proposed to answer these questions, the most prominent (at least in the philosophical literature) are the closely related treatments of Godfrey-Smith (1996) and Sterelny (2003a), the common core of which can reasonably be taken as the standard account of these issues.<sup>2</sup> However, as I try to make clear in this paper, closer consideration of this account reveals it to be unconvincing: its key contention – the claim that beliefs and desires evolved for making a certain kind of flexible behaviour possible – is theoretically implausible.

Because of this, I try to move past this account; specifically, I argue that a better place to look for the reason why we have evolved beliefs and desires is to be found in the way in which our information processing abilities are *organised*. While not fully settling the debate about the

---

<sup>1</sup> This clearly assumes that we really *do* have beliefs and desires. This seems a very reasonable assumption to make (see e.g. Stich, 2009a, 2009b; Nichols & Stich, 2003); moreover, to the extent that it is still found unpersuasive, the argument presented here can be taken to provide further evidence for its truth (see also Sterelny, 2003a, p. 7). Note also that the issue here does *not* concern the (folk) concepts of ‘belief’ and ‘desire’ – rather, it concerns the actual structure of our cognitive architecture.

<sup>2</sup> Note that, with the exception of Sterelny (2003a, chap. 5), nearly all of these accounts concentrate exclusively on *beliefs*, and leave *desires* largely unmodelled. For present purposes, though, it is better not to separate these two so drastically.

biological function of beliefs and desires – which would require a detailed empirical analysis – this novel account at least clarifies several important issues that are only dimly appreciated at the moment.

The paper is structured as follows. In section II, I make clear exactly what the problem is that needs to be solved. In section III, I show that the solution offered by the standard account cannot be considered successful. In section IV, I present my own solution. I develop some implications of that solution in section V. I conclude in section VI.

## **II. Reflexes vs. Beliefs and Desires: The Question in More Detail**

In order to determine why we have beliefs and desires, it is necessary to begin by making clearer what, exactly, it means to have a cognitive architecture based on beliefs and desires. In turn, in order to do this, it is necessary to get clearer on what the major alternative to having beliefs and desires is – how are the minds of organisms structured that do *not* use these kinds of mental states to make decisions? Consider this last question first.

In the present context, the major alternative to belief / desire-based minds can be taken to be a mental architecture based on *reflexes*. Intuitively, reflexes can be characterised as fairly direct and only minimally cognitive connections between specific perceptions of the world and specific internal states on the one hand, and specific behavioural responses on the other (see e.g. Sterelny, 2003a; Dickinson, 1994; Wehner, 1997; see also below).<sup>3</sup> There are two specific reasons for why it is useful to see *them* as the key alternative to having beliefs and desires.

Firstly, doing this is empirically very plausible: most simple organisms seem to act in ways that are captured well by a reflex-based model (see e.g. Grau, 2002; Shettleworth, 2002; Wehner,

---

<sup>3</sup> In behavioural ecology, it is customary to distinguish more finely among reflexes, habits, instincts, fixed actions patterns, drives and a number of other behavioural dispositions (see e.g. Krebs & Davies, 1997; Makintosh, 1984, 1994). In the present context, though, making this finer distinction is not necessary.

1997, p. 19; Hall, 1994, p. 15). Since we have evolved from a cognitively simpler ancestor, this makes it reasonable to assume that it, too, acted mainly in a reflexive way. Secondly, reflexive actions are the most *general* alternative way of generating variable responses. That is, the framework provided by reflexive actions can be applied to most other non-belief / desire-driven ways of generating flexible actions to the world (see also Shettleworth, 1994, p. 185). For example, the mechanism a sunflower uses to determine how it leans – namely, the amount of sunlight that hits different parts of its body – can be described in terms that are analogous to the model of reflex action laid out below. This means that contrasting beliefs and desires with reflexes does not skew the discussion in untoward ways, but marks the key issue that needs to be investigated.

In more detail, a reflex-based cognitive architecture can be described as having five major elements: (1) an *input system* (i.e., a set of mechanisms for perception); (2) an *internal configuration* (i.e., a set of states that the organism's non-perceptual systems can be in – e.g. various emotions, drives, etc.); (3) an *assembler* (i.e., a mechanism that combines the organism's input states and its internal states); (4) a *decision making system* (i.e., a mechanism that connects input / internal state combinations with particular actions); and (5) an *output system* (i.e. a set of mechanisms for generating motor commands).<sup>4</sup> These systems work as follows (see also Piccinini, 2007).

When a reflex-based organism receives information from the world, its input systems are caused to be in a certain state (e.g. the organism might see a walnut). At the same time, it will be in some internal state (e.g. it might be hungry). These input and internal states are then fed into

---

<sup>4</sup> For more on input and output systems, see e.g. Fodor (1983); for more on an organism's internal configuration, see Damasio (1994). In both cases, there is much room for debate about exactly what these systems are and how their states are best characterised; for present purposes, though, this can be left open. Note also that the term 'decision making system' ought not to be overinterpreted here – in particular, it is not to be seen as referring to any kind of *representational* system.

the organism's assembler, which, in turn, does two things. Firstly, it combines the two types of states by a simple process of concatenation; secondly, it passes the resulting combination on to the decision making system. Note that there is no reason to think that this input / internal state combination cannot be quite complex in structure: it can comprise currently occurring input and internal states, past input and internal states, and any (logical) combination thereof. Finally, the decision making system considers a pre-determined list of associations of input / internal state combinations with action sequences – the organism's 'table of reflexes' – and then initiates the appropriate motor commands by passing the relevant information on to the output system.<sup>5</sup> Graphically, this model can therefore be represented as follows (where each arrow exemplifies one state of the relevant system):<sup>6</sup>

[Figure 1: A Model of a Reflex-Driven Organism]

Matters are very different in the context of a belief / desire architecture. In particular, when it comes to the latter, input and internal states give rise directly only to *beliefs and desires*, and are not themselves passed on to the decision making system.<sup>7</sup> Importantly, moreover, this process of 'giving rise to beliefs and desires' is not straightforward – in fact, it is (in different ways) both *one-to-many* and *many-to-one*.

---

<sup>5</sup> Note also that, in some cases, the internal states can by themselves initiate an action (as when the organism stops doing whatever it is doing when it is in pain), and that the same is true for the input states (as when we lower our eyelids upon registering the presence of an object near our retina). For present purposes, though, these complications can be left aside.

<sup>6</sup> For a related (and slightly more detailed) picture, see also Dickinson (1994, p. 54) and Timberlake (2002, pp. 109-110). Note, though, that Dickinson's view of these matters is fairly complex: he thinks that a reflex-based model can be seen as the *realisor* of a belief / desire-based one (see e.g. Dickinson, 1994; Dickinson & Balleine, 2000). A closer discussion of this idea is not possible here, however.

<sup>7</sup> That said, since belief / desire-driven organisms can also retain some reflexes, some input and internal states can initiate behavioural sequences directly even there. See also notes 9 and 22.

On the one hand, *many different input and internal states* can cause *the same belief or desire* to come about: for example, the belief *that it is raining* might be caused by a number of different smells, sights, and sounds; equally, the desire *that my child be happy* might be caused by a number of different emotions and perceptual representations. On the other hand, *one input state* – or *one internal state* – will normally be the cause of *many different beliefs or desires*, depending on what the other input or internal states are with which it is coupled. For example, a perception of wet streets can lead to the belief *that it is raining* – when coupled with the sound of rain drops on the window pane – but also to the belief *that it is sunny* – when coupled with the sound of garden sprinklers. These complex connections between input and internal states and beliefs and desires are the result of the facts that (a) beliefs and desires are primarily information-carrying states, and (b) the information embedded in an input and internal state need neither be unique to that state nor fixed across situations. Put differently: while input and internal states are defined by their physical realisations, beliefs and desires are defined by their informational content – and there is a many-to-many correspondence between these two.<sup>8</sup>

All of this matters, as it makes clear that, with beliefs and desires, there is no longer any straightforward connection between input / internal state combinations and specific actions. In fact, input and internal states now are relevant to an organism's behaviour only to the extent that they give rise to *beliefs and desires* – for it is these beliefs and desires that are fed into the decision making system, and which determine what the organism will do. Graphically, this

---

<sup>8</sup> Of course, there is much more that could be said about the nature of beliefs and desires; for present purposes, though, the account of the text is sufficient. Note also that this account does not entail that content internalism a la Fodor (1980) is true. In particular, there is no need to assume that the *content* of our beliefs and desires are our input and internal states.

model can therefore be expressed as follows (the arrows again represent different states of the relevant systems):<sup>9</sup>

[Figure 2: A Model of a Belief / Desire-Driven Organism]

In a nutshell, therefore, reflex-based minds differ from belief / desire-based minds, in that (1) the latter process input and internal states further (to generate the organism's beliefs and desires), and (2) the latter make decisions not by consulting combinations of input and internal states directly, but by consulting combinations of beliefs and desires. This contrast between the two mind-designs can now be used to make more precise the rather abstract question as to why we have beliefs and desires. In particular, breaking this question into two parts, the issue can now be expressed as follows:

- (i) Why did the cognitive architecture in figure 2 evolve from that in figure 1 *in our lineage*?
- (ii) Why did it *not* evolve in the case of *many other organisms*?

With these two more specific questions in mind, we are now in a better position to consider the currently most widely accepted approach to the determination of why we have beliefs and desires.

---

<sup>9</sup> Note that since a belief / desire-driven organism is likely to maintain some of its reflexes (see also notes 7 and 22), a full picture of the latter kind of architecture would also contain an assembler that feeds some combinations of input- and internal states directly to the decision making system. I have not included this in figure 2 so as to make it easier to read. Note also that a combined picture of this sort might be seen to spell out the underpinnings of dual-system architectures like those defended e.g. by Sloman (1996), Stanovich (1999), and Carruthers (2006). I thank an anonymous referee for useful discussion of this point.

### III. Behavioural Flexibility, Environmental Complexity, and Belief / Desire-Based Decision Making

At least from a philosophical point of view, the common core of the treatments of Godfrey-Smith (1996) and Sterelny (2003a) has become the standard account of the evolution of beliefs and desires in the literature.<sup>10</sup> In order to see how compelling this account is, let's begin by considering it more closely.

#### 1. *The Standard Account: Behavioural Flexibility and Environmental Complexity*

According to the standard account, the reason why we have beliefs and desires is that they evolved due to their bringing about particular form of behavioural flexibility in the face of a certain kind of environmental complexity. More precisely, the account rests on the following two theses:

(Flexibility Thesis) We evolved the ability to form beliefs and desires due to its allowing us to react flexibly to the same state of the world.

(Complexity Thesis) Reacting flexibly to the same type of state of the world is beneficial in epistemically complex environments.

Consider the two theses in turn.

---

<sup>10</sup> Note that precursors of this account can be found in the 19<sup>th</sup> century (see e.g. the work of Spencer and Dewey); however, when it comes to the more specific issue of belief / desire-based vs. reflex-based decision making, Godfrey-Smith (1996) and Sterelny (2003a) clearly do deserve special mention. Note further that, in what follows, I shall concentrate on Sterelny's (2003a) treatment, as it is a bit more detailed than that of Godfrey-Smith (1996). Finally, note that Sober (1994) presents a model that is formally similar to that in Godfrey-Smith (1996); however, he uses this model in a very different way. In particular, Sober (1994) is interested in determining *which kinds* of beliefs (variable or fixed) are adaptive under which circumstances, not with why we have beliefs *at all*; accordingly, his account will not be discussed further in what follows.



The Flexibility Thesis is meant to express the thought that an organism may find itself in a position where the presence or absence of a cue cannot be completely reliably and beneficially connected to any one action (see e.g. Sterelny, 2003a, pp. 27-29; see also Shapiro, 1999).<sup>11</sup> In particular, fine differences in how this cue came about, and what other cues are present, may make it extremely beneficial to be able to react flexibly to its presence or absence. In cases like this, the standard account claims that it is best for the organism to *decouple* the representation of the state of the world and the goal that is to be achieved from any particular behavioural response, and to determine what to do using such decoupled representations (i.e. beliefs and desires) instead (see e.g. Sterelny, 2003a, chap. 3-4).

To this idea, the standard account then adds the claim (expressed above in the Complexity Thesis) that the usefulness of decoupled representations is an increasing function of the *epistemic complexity* of the organism's environment. The idea here is that in epistemically 'open' environments, the presence or absence of various cues (e.g. shadows, moving grass) is sufficiently informative about the state of the world so as to allow the organism to connect, in a beneficial way, one behavioural response to each such feature. However, as these environments get more complex – i.e. less epistemically open – these cues become less and less strongly correlated with the state of the world, thereby making it useful to vary the behavioural response to their detection (see e.g. Sterelny, 2003a, pp. 86; 92-93).<sup>12</sup>

In short, the standard account claims that we have beliefs and desires because they allow for flexible responses to an epistemically complex environment. In these kinds of environments, it

---

<sup>11</sup> For actual examples of these sorts of cases, see Heyes (1994, p. 293) and the references mentioned there.

<sup>12</sup> Note that this usefulness reaches a maximum, and then drops off again: in epistemically *closed* environments, there is no need to use decoupled representations either. There, acting in line with the relevant objective probabilities is all that can be achieved.

pays to react in different ways to the same circumstances. In this manner, the account answers the above two questions as follows:

- (i) *We* have beliefs and desires, as it was adaptive in the environments we evolved in to be able to react flexibly to the same state of the world.
- (ii) *Many other organisms* do *not* have beliefs and desires, as their environments are too epistemically simple to make this adaptive.

However, as I try to make clearer in what follows, a closer look at these answers suggests that they are both ultimately implausible.<sup>13</sup>

## 2. *The Return of the Reflex: A Criticism of the Standard Account*

The main problem the standard account has to face is that it fails to do justice to the power of reflexive behaviours. To see this most clearly, note that this account has left out of the picture the fact that, for a reflex-driven organism, the input / internal state combinations the assembler passes on to the decision making system can be very complex in structure.

This is important, as with suitably complex reflexes, the model in figure 1 can give rise to behaviours that are *indistinguishable* from those that result from the model in figure 2. This can be seen easily by noting that an organism can react flexibly – that is, non-uniformly – to an input / internal state combination C by breaking it down into a number of other such combinations, each of which has a unique action associated with it. In this way, flexible responding to C can be realised with inflexible responding to sub-states of C (see also Grau, 2002, pp. 77; Shettleworth,

---

<sup>13</sup> Further criticisms are found in Sober (1997) and Walsh (1997). However, the issues they raise are quite different from the ones brought up here: in particular, they are more concerned with the Complexity Thesis, whereas I am more concerned with the Flexibility Thesis.

2002, pp. 126-127; Pearce, 1994, pp. 123-128; Walsh, 1997; Sober, 1997).<sup>14</sup> The following example makes this clearer.

Assume that one of the main types of predators a population of prairie dogs has to face is a snake. Furthermore, assume that, depending on the exact circumstances surrounding the presence of a snake, a prairie dog will find different behavioural responses to be adaptive: when it perceives a snake to be some distance away, it is optimal to utter its distinctive alarm call and hide in a tunnel underground; however, when it perceives a snake to be nearby it is best to stay silent and attack the snake outright (see also Slobodchikoff, 2002). How, then, is the prairie dog to decide what it is do when it is perceiving a snake?

Now, it is clear that *one* way of answering this question is to have its actions be belief and desire determined. The prairie dog can form the belief that it faces a *snake* (as opposed to one of its other predators), that the snake it faces is *nearby*, and that, in these circumstances, it is best to *attack*. It can then combine these with its *desire to stay alive*, and in this way decide to act in the appropriate manner. However, it seems equally obvious that *another* way of generating this kind of behaviour is possible – namely, via a set of (rather sophisticated) reflexes. In particular, instead of individuating the world into ‘a snake is present’ and ‘no snake is present’ (or even more broadly: ‘a predator is present’, ‘no predator is present’), the prairie dog’s reflexes can individuate it into ‘a snake is at some distance’ and ‘a snake is nearby’. Given this, the animal can act reflexively *to the particular circumstances obtaining* and still be as flexible as a belief / desire-determined organism.

More generally, what looks like flexible reactions to the same input state (the perception of a snake) can just as well be seen as inflexible reactions to ‘sub-states’ of the state in question (the

---

<sup>14</sup> Note that Sterelny (2003a, footnote 3, p. 35) seems to recognise this issue as being potentially problematic for the standard account. However, he does not consider it further.

perception of a snake that is nearby, the perception of a snake that is at some distance away).<sup>15</sup>

This shows that the standard account is inadequate: there is no reason to think that *only* by having beliefs and desires can an organism respond flexibly to a complex environment – in fact, a reflexive responder might do so just as well. In turn, this makes clear that the Flexibility Thesis is implausible: given fine enough individuations of the input and internal states of the organism, it is perfectly possible to use reflexes to generate varying responses to what, on a coarser individuation scheme, is the *same* state of the world. However, if we do not have reason to think the Flexibility Thesis is true, the standard account collapses. Before presenting an alternative theory of the evolution of beliefs and desires which can avoid these concerns, I want to briefly consider two replies on behalf of the standard account. Doing so brings out several issues of importance.

The first reply suggests that belief / desire-driven organisms are more flexible than (even complex) reflex-driven ones, as only the former can *learn*. That is, it may be admitted that a reflexive responder can become just as flexible as a belief / desire-driven one *in evolutionary time*; however, it may also be maintained that only a belief / desire-driven organism can increase its flexibility *in developmental time*. This would clearly give the latter a major advantage: its ability to adjust its behaviour in developmental time means that it will always be several steps ahead of its slower, reflex-based competition (Godfrey-Smith, personal communication; see also Cosmides & Tooby, 2000, pp. 53-54; Sterelny, 2003a; Dickinson, 1994, pp. 45-46; Dennett, 1978).

However, this reply is based on the implicit assumption that reflexes cannot be acquired during the lifetime of the individual. This implicit assumption, though, has not much to

---

<sup>15</sup> A structurally related point has been made by Block (1981) in a different context. I thank Kim Sterelny for useful remarks about this matter.

recommend it: in fact, there is good reason to think that reflexive behaviours can be acquired in developmental time, too. The entire literature on the possibilities of conditioning is testament to this: pigeons, dogs, rats, and many other animals can be led to acquire a vast array of new reflexes within very short time frames (see e.g. Mackintosh, 1983; Sherman et al., 1997, pp. 79-81). Furthermore, there is no reason to think that there are not yet other ways of acquiring reflexes as well. For example, when driving a car, a human being might act in a largely reflex-driven way; however, most of the actions she is doing in this situation are neither innate nor acquired by a straightforward process of conditioning (see also Sober & Wilson, 1998, pp. 256-260; Morris, 1994, p. 135; Shettleworth, 1994, pp. 185-204). In short: since there is no reason to think that reflexes can be altered only over evolutionary time, this reply is unconvincing.

The second possible reply concerns the amount of cognitive complexity involved in figure 1. It may be thought that in order to do justice to the behavioural abilities of even moderately flexible organisms – like bees, ants and spiders – there would have to be a truly *astronomical* amount of reflexes (i.e. the list that associates of input / internal state combinations with actions would have to be *extremely* long). Having this many reflexes, though, may seem highly implausible: surely – it may be thought – this Baroque complexity is beyond the bounds of biological credibility.<sup>16</sup>

Now, I do think this reply is on the right track – however, as it stands it is implausible. As will also be made clearer below, while there are reasons to think that natural selection will favour the model in figure 2 over that in figure 1, there is no reason to think that the *absolute number* of reflexes in the model has something to do with it (see also Pearce, 1994, pp. 126-127). How many reflexes can an organism be said to have before this is considered biologically implausible

---

<sup>16</sup> For the clearest version of this kind of reply, see Sterelny (2003a, pp. 86; 92-93) (who attributes it to Dickinson). Another version is in Carruthers (2006, pp. 73-74).

– 50? 500? 5000? What is the right order of magnitude here? It is hard to think of a plausible way to answer this question. So, while the complexity inherent in model 1 might indeed provide evolutionary grounds for favouring beliefs and desires over reflexes, these grounds need to be spelled out more carefully than is done by the above reply.

In sum, the standard account fails to give a satisfactory answer to the question of why we have beliefs and desires. Its suggestion that beliefs and desires are an adaptive response to the need for extraordinary behavioural flexibility in a complex environment turns out, upon closer consideration, to be unable to make clear what is special about beliefs and desires: reflexes would seem to do just as well for these purposes. This means that the reason why we evolved beliefs and desires must be looked for elsewhere.

#### **IV. The Adaptive Importance of Cognitive Efficiency**

The core idea of the account to be developed here is that beliefs and desires are an adaptation, not for allowing for behavioural flexibility (as expressed by the Flexibility Thesis), but for increasing the organism's *cognitive efficiency* in generating actions. To show this, I begin by making clear that increased cognitive efficiency can explain why *we* have beliefs and desires. Given this, I then show why this explanation does *not* entail that *all organisms* ought to be expected to have beliefs and desires.

##### *1. Beliefs and Desires as Adaptations for Cognitive Efficiency*

To bring out the core idea of the account to be developed here, it is best to start by returning to the key differences between the minds of belief / desire-driven and reflex-driven organisms. In this regard, remember that, for reflex-driven organisms, decision making consists in 'consulting'

a mapping between (possibly quite complex) concatenations of input / internal states and actions. For present purposes, what is most important about this process is that the decision making mechanism treats every input / internal state combination independently from every other one. In particular, it does not matter how similar the information embedded in these combinations is, or how similar the actions are that are associated with these combinations; all that the decision making system does is go through the list of input / internal state combinations it recognises and find the template that matches the one it has received from the assembler. This will tell it what action to pass on to the output system to initiate.

In contrast, in a belief / desire-based architecture, the decision making system relies only on the information embedded in the input and internal states – i.e. it only considers *grouped combinations* of input and internal states. This seemingly small difference between the two cognitive architectures has some major implications for their computational efficiency. In order to bring this out most easily, assume, for the sake of the argument, that this is the only difference between them – in other words, assume that the belief / desire decision making system also works by consulting a mapping between particular actions and particular belief / desire pairs (as will become clearer momentarily, there is good reason to think that this assumption is false; for present purposes, though, it is useful to grant it).

Given this, the belief / desire-based decision making system has much room for *simplifying* the problem it has to solve – in particular, it is in a position to determine what the organism is to do much more easily and efficiently than the reflex-based alternative can. To see this, note that reflex-based organisms are always in danger of failing to see the forest for looking at all trees: their approach to figuring out how to behave is highly reductionist – in fact, it is so reductionist that they are unable to notice *patterns* in their behaviour that it is worthwhile picking up on.

This comes out particularly clearly by noting that if a reflex-based organism has a repertoire of  $n$  different input / internal state combinations that it needs to consider in order to make a decision, an otherwise similar belief / desire-based organism might only have to consider a repertoire of  $n' < n$  different belief / desire pairs to do so.<sup>17</sup> The reason for this is that the latter organism can subsume many of the  $n$  input / internal state combinations under the same belief / desire pair. An example might make this clearer (see also Whiten, 1995).

Assume (quite realistically) that a particular prairie dog can recognise the presence of a snake in its vicinity in many different ways: it might see it move, it might see a particular kind of movement in the grass, it might hear the snake slither over the dry ground, it might smell the snake, it might hear an appropriate alarm call from a conspecific, etc.. On the reflex-based picture, each of these states would need to be distinguished and separately connected with the appropriate action ('prepare to fight' or 'hide', for example). By contrast, on the belief / desire-based picture, it is sufficient for the prairie dog to consider only the belief *that a snake is nearby*, and associate *it* with the appropriate action. Figure 3 captures this difference graphically.<sup>18</sup>

[Figure 3: Reflex-Based Decision Making vs. Belief / Desire-Based Decision Making]

The key lesson to take away from this example is that, even if decision making is based only on basic look-up tables, using beliefs and desires can considerably simplify the problem to be solved: by recognising *patterns* in their behaviour, belief / desire-based organisms can get away with considering many fewer 'prompts' (the left-hand sides of the tables in figure 3) than reflex-

---

<sup>17</sup> Of course, given the complex connections between inputs and internal states on the one hand, and beliefs and desires on the other, an organism with beliefs and desires could also (if it so 'chose') make the cognitive problem more complex. However, this is not so relevant here: the point is just that it *can* make it easier.

<sup>18</sup> For expositional convenience, I here concentrate on beliefs and input states only. However, all the arguments carry over when desires and internal states are considered as well.



based organisms have to. In turn, this is likely to make decision making much more efficient (see also Sober, 1998a; Whiten 1995).<sup>19</sup> There are three main reasons for this.

Firstly, using beliefs and desires is likely to make decision making *faster*. Instead of having to comb through a list of  $n$  possible prompts, the organism can simply consider the smaller list of  $n'$  such prompts. This will make it faster to determine the appropriate behavioural response to the state of the world: there are fewer options to consider.

Secondly, using beliefs and desires is likely to *avoid errors*. If there are fewer prompts to consider, the probability is lower that the cognitive system misconnects some of these prompts with the wrong action. This is simply because, with fewer prompts, there are fewer possible false connections between these prompts – i.e. there are fewer ways to get it wrong.

Thirdly, using beliefs and desires is likely to *save cognitive and energetic resources*. Since beliefs and desires allow the organism to accomplish the same goals with fewer internal connections, it does not have to maintain as large a memory store as a reflex-based organism needs to. In turn, this saves energy, concentration and attention. This point is strengthened by the fact that belief / desire decision making is most likely *not* based on look-up tables of the above sort anyway: since beliefs and desires have a (structured) content, they allow for the *computation* of the appropriate action. For this reason, in some contexts, there might not be a need to *store* a particular connection between some belief / desire pair and some action at all – the latter can simply be *derived* from the former.

In this way, it becomes clear that there are good reasons to think that belief / desire-decision making will be cognitively more efficient than reflex-based decision making. Interestingly, this is a consequence of the fact that appealing to beliefs and desires *simplifies* the decision making

---

<sup>19</sup> Note that I am using ‘efficiency’ in a wide sense here, as including matters of speed, reliability, and resource frugality. I thank an anonymous referee for pointing out the need to be clearer about this.

problem: it allows for the recognition of macro-level generalities that the micro-level reflex-based approach overlooks.<sup>20</sup> Before considering further what implications this has for the evolution of beliefs and desires, it is now useful to consider an objection that might come to mind at this point.

This objection notes that it seems that all the benefits that come from having beliefs and desires might also be had by simply individuating input and internal states in a more coarse-grained manner. In particular, it might seem that if an organism can make decisions more efficiently by grouping various input and internal states together, there is really no reason for it to distinguish these input and internal states from each other in the first place: it seems it could simply consider them as *one* state – and act reflexively to that state. Hence, it may seem that my proposal cannot explain the benefits of having beliefs and desires either.

However, when looked at in detail, this objection fails to be convincing. Primarily, this is because it is in fact theoretically highly implausible to consider different sense impressions (understood broadly to including introspective reports) that carry the same information to be *one* input or internal state. This is due to the fact that these different sense impressions will often originate in different modalities (e.g. sight and smell), or carry different kinds of information depending on what other impressions they are coupled with (as in the above example of the wet streets). For this reason, taking these impressions to be the same input or internal state would obviate key aspects of their structure, and sit badly with the rest of cognitive science. Hence, this objection does not, in the end, threaten my claim that beliefs and desires bring efficiency gains in decision making.

---

<sup>20</sup> This also fits with many other investigations that have found that, often, the simpler approach to a problem can be the more successful one (see e.g. Dennett, 1989; Forster & Sober, 1994; Gigerenzer & Selten, 2001; Sober, 1998b, 2001, 2008, 2009). However, unlike what is true in these other contexts, the issue here is not a methodological or epistemological point – rather it is an *engineering* one.

These efficiency gains are important here, as they provide a reason for why natural selection might have preferred the evolution of a cognitive architecture based on beliefs and desires to one based on reflexes.<sup>21</sup> The reason, quite simply, is that natural selection will favour the more efficient cognitive design over the less efficient one. Not much argument is needed to establish this point: cognitive efficiency seems quite obviously to be something that provides fitness benefits to an organism (see also Cosmides & Tooby, 1992; see also Darwin, 1859, pp. 108-109, 130). However, if more concreteness is desired, it can simply be noted that the three factors cited above for why belief / desire-based minds are more computationally efficient than reflex-based ones can also be used as reasons for why the former are *more adaptive* than the latter. Specifically, it seems quite obvious that faster, more reliable, and more frugal decision making will, *ceteris paribus*, be favoured by natural selection – all of these factors make organisms better able to deal with *any kind* of environment it might find itself in. Hence, there is good reason to think that beliefs and desires will be favoured by natural selection over reflexes.<sup>22</sup>

Note that, unlike what is said in the standard account, this reason for why beliefs and desires are more adaptive than reflexes is not located in *external* matters; instead, it is located in *internal* matters.<sup>23</sup> That is, natural selection is here shown to favour belief / desire-based minds over reflex-based ones not because the former allow the organism to do things that the latter do not; rather it is that the former allow the organism to do same things *better*. In other words, the

---

<sup>21</sup> This does not mean that this must be the *only* reason why natural selection has favoured beliefs and desires over reflexes. The point is just that cognitive efficiency – unlike flexible responding – is *one* (key) reason for why it might have done so; other such reasons may exist.

<sup>22</sup> Note that this statement should not be read as implying that belief / desire-driven organisms get rid of their reflexes altogether (see notes 7 and 9). In fact, it is quite likely that an organism will want to retain some of its reflexes even after it switched to using beliefs and desires. This is due to the fact an organism's table of reflexes need not have a uniform potential to be simplified: in particular, some of an organism's actions might only be connected to very few prompts, so that there will not be major efficiency gains from making these actions belief / desire-driven.

<sup>23</sup> Alternatively, it could be said that whereas the standard account focuses on the fact that one input state can be associated with many different beliefs (and similarly for internal states and desires), my account focuses on the fact that many different input states can be associated with only one belief (and similarly for internal states and desires).

argument here is that there is selection for *efficiency in mind design*, not for an *ability to do novel things*.

In this way, therefore, we have developed a new answer to question (i) of section II. However, so far, this answer would seem to suggest that *all* organisms ought to have beliefs and desires. How, then, can the above account be used to answer question (ii)? The next sub-section addresses this issue.

## 2. *The Limits of Adaptationism*

The reason why the above account does not entail that all organisms ought to be expected to have beliefs and desires rests on the fact that what is adaptive and what evolves are not necessarily the same thing: not everything that is adaptive needs to evolve, and not everything that evolves needs to be or to have been adaptive (see e.g. Sober, 1993; Orzack & Sober, 1994; Godfrey-Smith, 2001). This distinction between adaptiveness and evolution matters here, as it provides the basis for the explanation of why it will often be more reasonable to infer that a belief / desire architecture did *not* evolve, even though it was more adaptive than a reflex-based one.

To see this in more detail, note firstly that evolutionarily *assembling* a belief / desire architecture is very likely to be costly (see also Sterelny 2003b, p. 293). In particular, it is likely to require radical alterations to the ancestral organism's reflex-based architecture – such as a complete resetting of the neural network that underlies that organism's mind, or a major physiological change to the basic structure of its brain (such as an increase in its size, with all the concomitant changes that this brings) – which cannot be done without major interruptions to the workings of the organism's decision making machinery. In turn, this means that an organism that

tries to accomplish the switch from reflexes to beliefs and desires is very likely to have to cross a *fitness valley*: its fitness will have to decrease before it can increase from its current position (for more about fitness valleys, see also Sober, 2008, chap. 3).<sup>24</sup> Graphically:

[Figure 4: A Non-Monotonic Fitness Function between Reflexes and Beliefs and Desires]

The existence of a fitness valley between reflex-based minds and belief-desire based ones is important, as it entails that there will frequently be cases where we would predict belief / desire-based minds *not* to evolve – despite their being adaptive. In particular, given the fitness function represented in figure 4, the evolution of beliefs and desires depends on either the occurrence of a macromutation that takes the organism all the way from point A to point D in the diagram, or the occurrence of a succession of micromutations that takes it gradually from point A over points B and C to point D. However, the first option is unlikely to come about, as it goes against the nature of genetic mutations (which are based on small changes in the relevant genes); and the second of these options is unlikely to come about, as it is being selected against at every stage (the fitnesses of the mutants will be lower than that of the ancestral organism all the way until D). Therefore, in many contexts, one would expect beliefs and desires *not* to have evolved.

Given this, one might wonder if, conversely, there are situations in which one should expect that these unlikely occurrences *did* come about – in other words, are there circumstances that make it easier for an organism to evolve beliefs and desires than others? Now, it is not clear that

---

<sup>24</sup>A non-psychological example might make this clearer (see also Sober, 2008, chap. 3). Assume that the mammalian camera eye is the best available eye design. Should we therefore expect that all organisms have these kinds of eye? Not necessarily – since evolution is based on ‘gradient climbing’ (i.e. the incremental improvement of a trait from a given starting point), whether an organism evolves a camera eye depends on what kind of eye design it *started out* with. Now, in some circumstances, it may be the case that switching from a *non*-camera design to a camera design requires changes that, intermittently, lead to *worse* eyesight. In these circumstances, the relevant organism might *not* evolve a camera eye – and that is so even if the latter eye design is more adaptive than the one it currently has, and even if there are no constraints on the workings of natural selection.

a general answer to this question can be given: it is very plausible that the evolution of beliefs and desires has taken different paths in different types of organisms, and that it is not possible to state in general what it would take to make this evolution more likely. That said, some rough guidelines about this evolution can be stated nonetheless.

Given figure 4, the likelihood of the evolution of beliefs and desires is greater, (a) the flatter the fitness function is (i.e. the smaller the fitness difference is between point A and the trough of the curve), and (b) the more narrow the fitness function is (i.e. the smaller the horizontal difference is between point A and the trough of the curve). Concretely, this means that beliefs and desires are more likely to evolve either when the necessary changes to mind of the reflex-based ancestral organism do not lead to great losses in fitness, or when they do lead to losses in fitness, but only briefly. There seem to be many ways in which these circumstances can be satisfied. For example, it might be that the relevant organisms have few predators, ample resources, and are relatively numerous; in these cases, small losses in fitness might be tolerated for quite some time.

Of course, at this point, all of this is just a conjecture; for a fully fledged account of the distribution of beliefs and desires across the animal kingdom, a detailed empirical study of the relevant fitness functions and the importance of various evolutionary constraints would be necessary. What matters here, though, is just that, firstly, we now at least know what we are looking for – we have a much more precise hypothesis on the table that we can test. Secondly, it is now clearer that saying that minds based on beliefs and desires are *universally adaptive* need not commit one to saying that these kinds of minds are *widespread* – and that without belittling the importance of natural selection for this case.

In sum, therefore, my account suggests that beliefs and desires – where they have evolved – are an adaptation for *cognitive efficiency*. In this way, the two questions raised in section II can be answered as follows:

- (i) *We* (or our ancestors) have evolved beliefs and desires, because they brought benefits in cognitive efficiency when compared to sticking to complex reflexes.
- (ii) *Other* animals have *not* evolved beliefs and desires, since for them, the costs of switching to this radically different cognitive architecture were too great to make this evolution feasible.

I bring out several further implications of these answers in the next section.

## **V. Some Further Implications**

Two major implications of the idea that beliefs and desires are adaptations for improving cognitive efficiency are worth mentioning here. The first concerns the importance of representations for higher cognition, and the second concerns the study of animal minds. Consider them in turn.

### *1. The Non-Representational Mind*

The first implication of the account I have presented concerns the importance (or lack thereof) of mental representation for the study of cognition. Recently, a number of researchers have questioned the widespread assumption in cognitive science that representational (i.e. belief / desire-based) processing is somehow the key to higher thinking (see e.g. Clark, 1997; Brooks,

1991; van Gelder, 1996, Rowlands, 1999). According to these researchers, it is much more plausible to view the mind as being largely non-representational: much of the information it needs to function is in the world anyway, and does not need to be represented internally (see e.g. Clark, 1997; Brooks, 1991; Gibson, 1979). That is, these researchers think that, in many cases, making use of representations when dealing with the world is a needlessly costly and wasteful luxury – the world is rich enough ‘to be its own model’ and does not need to be ‘copied’ in an internal representation (see e.g. Brooks, 1991).

For present purposes, what matters most about these anti-representationalist arguments is that my account partly supports and partly diverges from them. It supports them in emphasising that representationalism really may not always be the most plausible assumption about how the mind of some organism is structured – in many cases, it truly is more reasonable to think that this organism processes information in a non-representational way. In particular, when the costs that come from building a representational architecture out of a reflex-based one were great, it might well be more reasonable to suppose that this architecture did not evolve.

However, my account also diverges quite considerably from the anti-representationalist picture. In particular, it stresses the fact that, in many cases, representationalism still *is* very likely to be a key feature of the mind. Importantly moreover, it does so by noting that – contrary to what Clark et al. might claim – there are circumstances where the use of representations makes cognitive processing *easier*, not harder. That is, the account shows that, in many cases, representationalism is not a luxury that could be done without, but rather something that is positively useful for efficient cognitive processing (see also Prinz & Barsalou, 2000; Sober, 2009). In this way, my account can help clarify the discussion surrounding the importance of representations in an organism’s cognitive life.



2. *Animal Minds*

The second important implication of the account I have presented concerns the question of which other animals have beliefs and desires. In order to see what the above account suggests concerning this matter, it is best to begin by considering the current state of the debate. Here, the positions vary quite widely: for example, Carruthers (2006, pp. 70-83; see also Wilcox & Jackson, 2002; Gould, 2002; Slobodchikoff, 2002; Tomasello & Zuberbuhler, 2002; Gallistel, 1994) suggests that most animals do have beliefs and desires. He bases this view on the fact that many animals can complete complex behavioural feats – in particular, they can act in ways that require highly sophisticated uses of variable environmental information. He suggests that this warrants our attributing beliefs and desires to them (see e.g. Carruthers, 2006, pp. 70-83). In contrast, Sterelny (2003a, chap. 4) suggests that it is not at all obvious that many animals have beliefs and desires. He comes to this conclusion via his commitment to the Flexibility and Complexity Theses: in particular, he thinks that for many animals (maybe even including some of the higher primates), the kind of environmental complexity that makes beliefs and desires adaptive might not have been present. Accordingly, he doubts that there are compelling evolutionary reasons for thinking that many non-human animals have beliefs and desires.

The account defended here can clarify the terms of this dispute. In particular, it shows that both Carruthers (2006) and Sterelny (2003a) may have gone wrong in their reasoning. Carruthers's (2006) reliance on the fact that ants, spiders and many other animals can behave in very flexible ways misses the point that flexible behaviour can be had without beliefs and desires (see also Wehner, 1997; Shettleworth, 2002; Rendell & Owren, 2002, pp. 310-311). By contrast, Sterelny's (2003a) suggestion that the kinds of environments many animals live in might not be sufficiently complex for them to need to evolve beliefs and desires misses the point that

environmental complexity is not necessary for beliefs and desires to be adaptive. In other words, the above account makes clear that neither Carruthers (2006) nor Sterelny (2003a) base their case on solid arguments. It shows that, instead of their appeals to behavioural or environmental complexity, what is important is how the *information-processing abilities* of an organism are organised.

How, then, is one to answer the question as to whether ants, spiders, birds, and macaques have beliefs and desires? The above account remains mute on this issue. However, neither should it be expected to pronounce on it: what it aims to do is merely to make clearer what we ought to consider *in order to* determine whether other animals have beliefs and desires. It does *not* aim to actually *give* an answer to this question: doing so is an empirical matter that cannot be determined from the armchair.

## **VI. Conclusion**

In this paper, I hope to have achieved three aims. Firstly, I hope to have made clearer what, exactly, the question about why we have beliefs and desires is about: it concerns the reasons for the evolution of the cognitive architecture depicted in figure 2 from that depicted in figure 1. Secondly, I hope to have thrown doubt on the common view that the reason why we have beliefs and desires is based on a combination of environmental complexity and behavioural flexibility: this fails to do justice to the powers of reflex-driven organisms in generating flexible responses to a variable environment. Finally, I hope to have provided a compelling alternative theory of why beliefs and desires might have evolved: namely, one in which they are adaptations for efficient information processing. Importantly, apart from their inherent interest, these

conclusions are noteworthy also for their broader implications for the nature of our minds and those of other animals.

## Bibliography

- Block, N.: 1981, Psychologism and Behaviorism, *Philosophical Review* 90: 5-43.
- Brooks, R.: 1991, Intelligence without Representation, *Artificial Intelligence* 47: 139–159.
- Carruthers, P.: 2006, *The Architecture of the Mind*, Cambridge University Press, Cambridge.
- Cosmides, L., and Tooby, J.: 1992, The Psychological Foundations of Culture. In L. Cosmides, J. Tooby, and J. Barkow (eds.), *The Adapted Mind*, Oxford University Press, Oxford, pp. 19-136.
- Clark, A.: 1997, *Being There*, MIT Press, Cambridge, MA.
- Damasio, A.: 1994, *Descartes' Error*, Harper, New York.
- Darwin, C.: 1859, *The Origin of Species*, Facsimile of the 1<sup>st</sup> Edition, Harvard University Press, Cambridge, MA.
- Dennett, D.: 1978, Skinner Skinned. In *Brainstorms*, MIT Press, Cambridge, MA, pp. 53-70.
- Dennett, D.: 1989, *The Intentional Stance*, MIT Press, Cambridge, MA.
- Dickinson, A.: 1994, Instrumental Conditioning. In N. Mackintosh (ed.), *Animal Learning and Conditioning*, Academic Press, San Diego, pp. 45-79.
- Dickinson, A. & Balleine, B.: 2000, Causal Cognition and Goal-Directed Action. In C. Heyes and L. Huber (eds.), *The Evolution of Cognition*, MIT Press, Cambridge, MA, pp. 185-204.
- Fodor, J.: 1980, Methodological Solipsism Considered as a Research Strategy in Cognitive Psychology, *Behavioral and Brain Sciences* 3: 63-72.
- Fodor, J.: 1983, *The Modularity of Mind*, MIT Press, Cambridge, MA.

- Forster, M. and Sober, E.: 1994, How to Tell when Simpler, More Unified, or Less Ad Hoc Theories Will Provide More Accurate Predictions. *British Journal for the Philosophy of Science* 45: 1-35.
- Gallistel, C.: 1994, Space and Time. In N. Mackintosh (ed.), *Animal Learning and Conditioning*, Academic Press, San Diego, pp. 221-253.
- Gibson, J. J.: 1979, *The Ecological Approach to Visual Perception*, Houghton-Mifflin, Boston.
- Gigerenzer, G. and Selten, R. (eds.): 2001, *Bounded Rationality: The Adaptive Toolbox*, MIT Press, Cambridge, MA.
- Godfrey-Smith, P.: 1996, *Complexity and the Function of Mind in Nature*, Cambridge University Press, Cambridge.
- Godfrey-Smith, P.: 2001, Three Kinds of Adaptationism. In S. Orzack & E. Sober (eds.), *Adaptationism and Optimality*, Cambridge University Press, Cambridge, pp. 335-357.
- Gould, J.: 2002, Can Honey Bees Create Cognitive Maps?. In M. Bekhoff, C. Allen, and G. Burghardt (eds.), *The Cognitive Animal*, MIT Press, Cambridge, MA, pp. 41-46.
- Grau, J.: 2002, Learning and Memory without a Brain. In M. Bekhoff, C. Allen, and G. Burghardt (eds.), *The Cognitive Animal*, MIT Press, Cambridge, MA, pp. 77-88.
- Hall, G.: 1994, Pavlovian Conditioning: Laws of Association. In N. Mackintosh (ed.), *Animal Learning and Conditioning*, San Diego: Academic Press, pp. 15-43.
- Heyes, C.: 1994, Social Cognition in Primates. In N. Mackintosh (ed.), *Animal Learning and Conditioning*, Academic Press, San Diego, pp. 281-305.
- Krebs, J. and Davies, N.: 1997, *Behavioural Ecology: An Evolutionary Perspective*, Blackwell, Oxford.
- Mackintosh, N. J.: 1983, *Conditioning and Associative Learning*, Clarendon Press, Oxford.

- Mackintosh, N. J.: 1994, *Animal Learning and Conditioning*, Academic Press, San Diego.
- Morris, R.: 1994, *The Neural Basis of Learning with Particular Reference to the Role of Synaptic Plasticity: Where Are We a Century After Cajals Speculations?*. In N. Mackintosh (ed.), *Animal Learning and Conditioning*, Academic Press, San Diego, pp. 135-183.
- Nichols, S. and Stich, S.: 2003, *Mindreading*, Oxford University Press, Oxford.
- Orzack, S. and Sober, E.: 1994, *Optimality Models and the Test of Adaptationism*, *American Naturalist* 143: 361-380.
- Pearce, J.: 1994, *Discrimination and Categorization*. In N. Mackintosh (ed.), *Animal Learning and Conditioning*, Academic Press, San Diego, pp. 109-134.
- Piccinini, G.: 2007, *Computing Mechanisms*, *Philosophy of Science* 74: 501-526.
- Prinz, J. and Barsalou, L. W.: 2000, *Steering a Course for Embodied Representation*. In E. Dietrich & A. Markman (eds.), *Cognitive Dynamics: Conceptual Change in Humans and Machines*, MIT Press, Cambridge, MA, pp. 51-77.
- Rendell, D. and Owren, M.: 2002, *Animal Vocal Communications: Say What?*. In M. Bekhoff, C. Allen, and G. Burghardt (eds.), *The Cognitive Animal*, MIT Press, Cambridge, MA, pp. 307-314,
- Rowlands, M.: 1999, *The Body in Mind: Understanding Cognitive Processes*, Cambridge University Press, Cambridge.
- Shapiro, L.: 1999, *Presence of Mind*. In V. Hardcastle (ed.), *Biology Meets Psychology: Constraints, Connections, Conjectures*, MIT Press, Cambridge, MA, pp. 83-98.
- Sherman, P., Reeve, H., Pfennig, D.: 1997, *Recognition Systems*. In J. Krebs & N. Davies (eds.), *Behavioural Ecology: An Evolutionary Perspective*, Blackwell, Oxford, pp. 69-96.

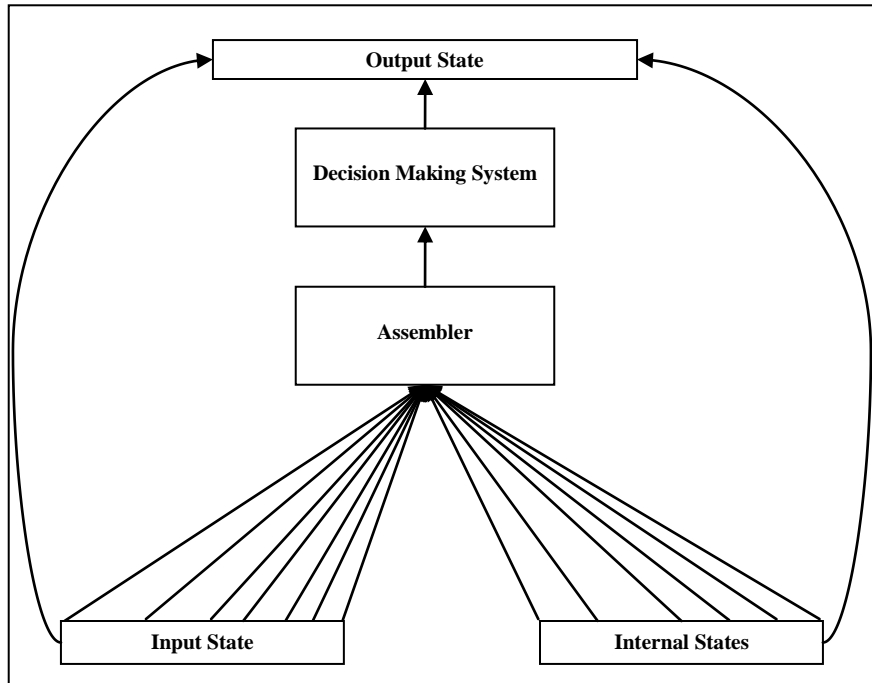
- Shettleworth, S.: 1994, *Biological Approaches to the Study of Learning*, In N. Mackintosh (ed.), *Animal Learning and Conditioning*, Academic Press, San Diego, pp. 185-219.
- Shettleworth, S.: 2002, *Spatial Behavior, Food Storing, and the Modular Mind*. In M. Bekhoff, C. Allen, and G. Burghardt (eds.), *The Cognitive Animal*, MIT Press, Cambridge, MA, pp. 123-128.
- Slobodchikoff, C. N.: 2002, *Cognition and Communication in Prairie Dogs*. In M. Bekhoff, C. Allen, and G. Burghardt (eds.), *The Cognitive Animal*, MIT Press, Cambridge, MA, pp. 257-264.
- Slovan, S.: 1996, *The Empirical Case for Two Systems of Reasoning*, *Psychological Bulletin* 119: 3-22.
- Sober, E.: 1993, *Philosophy of Biology*, Westview Press, Boulder.
- Sober, E.: 1994, *The Adaptive Advantage of Learning and A Priori Prejudice*. In *From A Biological Point of View*, Cambridge University Press, Cambridge, pp. 50-70.
- Sober, E.: 1997, *Is the Mind an Adaptation for Coping with Environmental Complexity?*, *Biology and Philosophy* 12: 539-550.
- Sober, E.: 1998a, *Black Box Inference*, *British Journal for the Philosophy of Science* 49: 469-498.
- Sober, E.: 1998b, *Morgan's Canon*. In D. Cummins & C. Allen (eds.), *The Evolution of Mind*, Oxford University Press, Oxford, pp. 224-242.
- Sober, E.: 2001, *The Principle of Conservatism in Cognitive Ethology*, In D. Walsh (ed.), *Naturalism, Evolution, and Mind*, Cambridge University Press, Cambridge, pp. 225-238.
- Sober, E.: 2008, *Evidence and Evolution*, Cambridge University Press, Cambridge.

- Sober, E.: 2009, Parsimony and Models of Animal Minds. In R. Lurz (ed.), *The Philosophy of Animal Minds*, Cambridge University Press, Cambridge.
- Sober, E., and Wilson, D. S.: 1998, *Unto Others: The Evolution and Psychology of Unselfish Behavior*, Harvard University Press, Cambridge, MA.
- Stanovich, K.: 1999, *Who Is Rational?*, Lawrence Erlbaum, Hillsdale.
- Sterelny, K.: 2003a, *Thought in a Hostile World*, Blackwell Publishing, Oxford.
- Sterelny, K.: 2003b, Darwinian Concepts in the Philosophy of Mind. In J. Hodge and G. Radick (eds.), *Cambridge Companion to Darwin*. Cambridge University Press, Cambridge, pp. 288-309.
- Stich, S.: 2009a, Reply to Egan. In D. Murphy & M. Bishop (eds.), *Stich and His Critics*, Blackwell Publishing, Oxford.
- Stich, S. 2009b, Reply to Godfrey-Smith. In D. Murphy & M. Bishop (eds.), *Stich and His Critics*, Blackwell Publishing, Oxford.
- Timberlake, W.: 2002, Constructing Animal Cognition. In M. Bekhoff, C. Allen, and G. Burghardt (eds.), *The Cognitive Animal*, MIT Press, Cambridge, MA, pp. 105-114.
- Tomasello, M., and Zuberbuhler, K.: 2002, Primate Vocal and Gestural Communication. In M. Bekhoff, C. Allen, and G. Burghardt (eds.), *The Cognitive Animal*, MIT Press, Cambridge, MA, pp. 293-300.
- van Gelder, T.: 1996, Dynamics and Cognition. In J. Haugeland (ed.), *Mind Design II*, MIT Press, Cambridge, MA, pp. 421-450.
- van den Bos, R., Houx, B., and Spruijt, B.: 2002, Cognition and Emotion in Concert in Human and Nonhuman Animals. In M. Bekhoff, C. Allen, and G. Burghardt (eds.), *The Cognitive Animal*, MIT Press, Cambridge, MA, pp. 97-104.

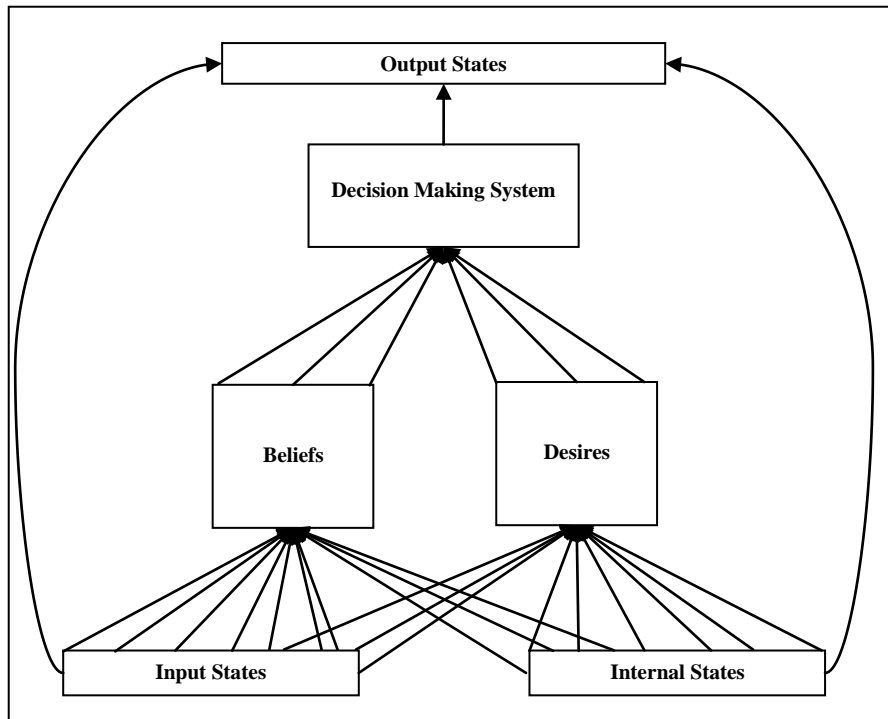


- Walsh, D.: 1997, Review of P. Godfrey-Smith – Complexity and the Function of Mind in Nature, *British Journal for the Philosophy of Science* 48: 613-617.
- Wehner, R.: 1997, Sensory Systems and Behaviour. In J. Krebs & N. Davies (eds.), *Behavioural Ecology: An Evolutionary Perspective*, Blackwell, Oxford, pp. 19-41.
- Whiten, A.: 1995, When does Smart Behavior-Reading Become Mind-Reading?, In P. Carruthers and P. Smith (eds.), *Theories of Theories of Mind*, Cambridge University Press, Cambridge, pp. 277-92.
- Wilcox, S., and Jackson, R.: 2002, Jumping Spider Tricksters: Deceit, Predation, and Cognition. In M. Bekhoff, C. Allen, and G. Burghardt (eds.), *The Cognitive Animal*, MIT Press, Cambridge, MA, pp. 27-34.

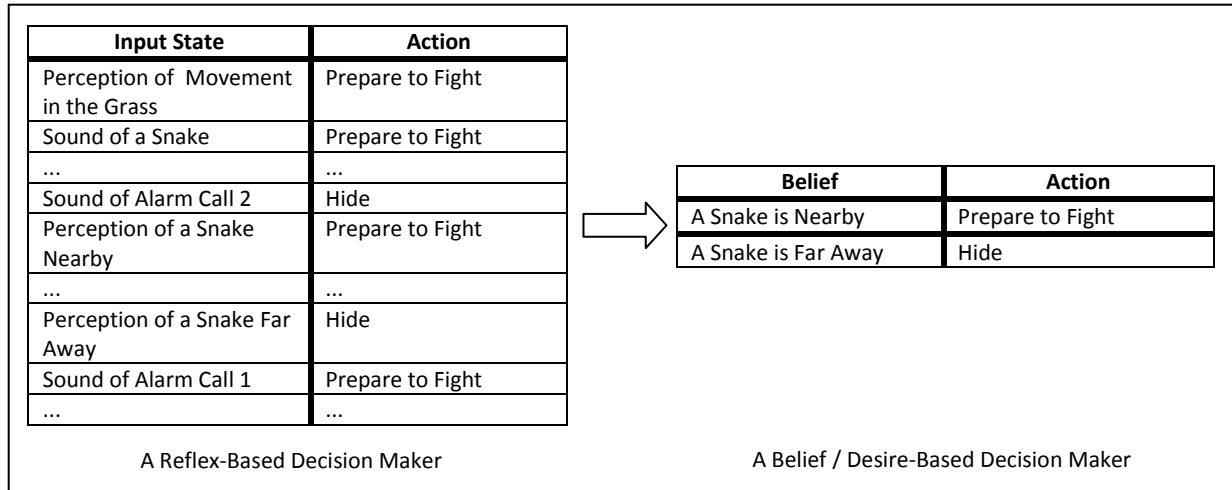
### Diagrams and Figures




[Figure 1]



[Figure 2]



[Figure 3]



[Figure 4]