

A CRITICAL ENGAGEMENT WITH CLARK'S ACCOUNT  
OF THE ROLE OF MOTIVATION IN A BAYESIAN  
INFORMATION PROCESSING MODEL OF THE BRAIN

ELMARIE VENTER

UNIVERSITY OF KWA-ZULU NATAL

2014

## **Declaration**

Submitted in fulfilment of the requirements for the degree of Master of Arts in  
Cognitive Science, in the Graduate Programme in College of Humanities,  
University of KwaZulu-Natal, Durban, South Africa

I, Elmarie Venter, declare that:

1. The research reported in this thesis, except where otherwise indicated, is my original research.
2. This thesis has not been submitted for any degree or examination at any other university.
3. This thesis does not contain other persons' data, pictures, graphs or other information, unless specifically acknowledged as being sourced from other persons.
4. This thesis does not contain other persons' writing, unless specifically acknowledged as being sourced from other researchers. Where other written sources have been quoted, then:
  - a. Their words have been re-written but the general information attributed to them has been referenced.
  - b. Where their exact words have been used, then their writing has been placed inside quotation marks, and referenced.
5. This thesis does not contain text, graphics or tables copied and pasted from the internet, unless specifically acknowledged, and the source being detailed in the thesis and in the references.

---

Student Name

---

Date

---

Name of Supervisor

---

Signature

## **Abstract**

The action selection problem can be approached with two goals in mind: to account for why actions are selected, i.e. what factors play a role in making the decision, and how actions are selected, i.e. what mechanisms are involved in action selection. These two goals form the main focus of this dissertation and can be referred to as the efficiency problem and the architecture problem. This dissertation examines the Environmental Complexity Hypothesis as an approach to the former goal (why actions are selected). The Environmental Complexity Hypothesis adequately explains why actions are selected but lacks consideration of the architecture and mechanisms involved in complex cognitive systems. In response to this shortcoming, the predictive processing account is offered as a solution to the second problem of action selection, the architecture problem.

The predictive processing account will be critically examined and one objection will be raised against its claim about the function of cognition. According to the predictive processing account, the function of cognition is to minimize free energy. One objection to this claim takes the form of a *reductio ad absurdum* argument and suggests that if the predictive processing account is correct about the function of cognition, then biological agents will find themselves in the state of seeking environments with no sensory stimuli. This is biologically nonsensical because environments with no sensory stimuli will result in no surprise and no prediction error which, the account claims, needs to be reduced for successful adaptation. However, reducing prediction error alone does not result in adaptive fitness and is not the sole function of the cognitive system; agents explore and participate in the environment to reproduce and to satisfy desires and preferences as well. This view is reflected in the field of neuroeconomics which provides a fairly novel solution to the problem of action selection and decision making. Neuroeconomic models of action selection can provide insight about why agents select certain actions when alternatives are available and it also gives

an account of the architecture involved in action selection. It is proposed that prediction error is coded in relation to rewards. By recommending neuroeconomics as an alternative account of action selection, nothing is lost in either the Environmental Complexity Hypothesis or the predictive processing account. Instead, neuroeconomics represents the best features of the two theories in a unified account.

## **Table of contents**

<b>Details</b>	<b>Page</b>
Introduction	1
1. Environmental Complexity Hypothesis	7
1.1. Simple detection systems	9
1.2. Robust tracking systems	14
1.3. Decoupled representation	17
2. Action Selection	20
2.1. Efficiency problem	21
2.2. Architecture problem	27
3. Predictive processing model of cognition	29
3.1. The nature of the system	30
3.1.1. Bidirectional nature	31
3.1.2. Hierarchical nature	34
3.1.3. The Bayesian brain	35
3.2. The function of the system	37
4. Free-energy formulation	38
5. Action-oriented predictive processing approach	42
6. Summary so far	47
7. The dark-room problem	49
7.1. Friston's response to the dark room problem	52
7.2. Clark's response to the dark room problem	58

8.	Neuroeconomics: an account of value and decision making	64
8.1	Classical economic theories of decision making	67
8.2	Neoclassical economic models of decision making	72
8.2.1.	Revealed preference theory (WARP)	72
8.2.2.	Generalized axiom of revealed preference (GARP)	74
8.3	Expected utility theory	77
9.	Neural representations of value	80
9.1	Evidence of neural representations of value	82
10.	How value is encoded in the brain	88
10.1	Dopamine neuron activation - an experiment	91
11.	Predictive processing and neuroeconomics	93

## References

## **Introduction**

Decision making and action selection are notable topics in cognitive science because they are central to the efficiency of animal behaviour which in turn is important to overall fitness. When an animal is threatened by a predator, it has several possible actions to choose from; it can flee, attack, or stay still in an attempt to be undetected. Whatever the animal's decision, it will directly affect its efficiency and could give the animal a competitive advantage over others. Providing an account of decision making is challenging because decision making is not only influenced by external factors of the environment, but also by internal signals and motivation. A hungry animal, for example, will risk a predator in situations in which a satiated animal will flee (Sterelny 2003: 79). In addition to this, there is no agreed upon definitions and boundaries when using terms such as 'intelligence' and 'cognition'. This is problematic because intelligence and cognition play an important role in decision making and action selection. Nonetheless, there is some common sense usage of the terms which allows progress to be made in cognitive science.

The focus of this dissertation is to provide an account of action selection that takes into consideration the two problems of action selection. The first of these problems is the efficiency problem and the second is the architecture problem; the efficiency problem asks the question, why are actions selected? A solution to this problem will take into consideration the efficiency of behavioural responses and goal states. It is proposed that efficient behavioural responses give animals a competitive and reproductive advantage over others. The complexity of the environment also plays an important role in decision making and action selection because animals are required to respond to environmental complexities by making good decisions. To explore this view, I examine the environmental complexity hypothesis in section 1 of this dissertation. The environmental

complexity hypothesis claims that the function of cognition is to control behaviour which enables the animal to cope with environmental complexity (Godfrey-Smith 2002: 235). Greater environmental complexity requires more sophisticated cognitive mechanisms. This view is reflected in Kim Sterelny's account of the evolution of cognition; Sterelny proposes that mechanisms of cognition are distinct in complexity and sophistication and that this enables efficiency in decision making and action selection. I examine Sterelny's account of the evolution of cognition in section 1.1 up to and including section 1.3. Sterelny's account of cognition takes seriously the action selection problem and offers important insight about the elements involved in action selection.

Given that the aim of this dissertation is to account for both problems of action selection, I now turn to the second problem, the architecture problem, which asks the question, how are actions selected? In other words, what architecture and mechanisms are involved in making decisions and selecting actions? The architecture problem is a problem not dealt with by Sterelny whose account includes numerous examples of mechanisms involved in simple cognitive systems but has a paucity of discussion on the architecture and mechanisms involved in more sophisticated and complex cognitive systems. I present this argument in section 2.1., and propose that one solution is to offer an account of the mechanisms and architecture of cognition that is suitable to Sterelny's account of cognition. In section 2.2., I introduce the predictive processing account of cognition as a solution to the architecture problem and as a complementary account to Sterelny's project.

The predictive processing account of cognition is offered as a general account of cognition that sees the brain largely as a prediction machine. Andy Clark, a leading cognitive scientist at the University of Edinburgh, examined the



predictive processing account in a target article published in Behavioural Brain Sciences in 2013. Clark provides an unequivocal description of the predictive processing account but it is important to keep in mind that he does not pay any attention to action selection and is mostly concerned with perception and action control. Nonetheless, his account of cognition provides information about the nature of cognition and provides a solution to the architecture problem. According to the account, cognition has the following properties: (1) cognition is bidirectional in nature, (2) cognitive processes occur at various levels in cognition which is indicative of its hierarchical nature, and (3) the brain is a Bayesian inference system. Clark also examines a claim, made by the predictive processing account, about the function of cognition. It is proposed that cognition functions to reduce prediction error which is the result of a mismatch between top-down predictions and bottom-up information. Prediction error reports the information-theoretic quantity known as free energy; it is proposed that biological organisms have a biological imperative to minimize free energy. I discuss this relationship in section 3.2. and section 4.

In section 5, I examine Clark's fairly distinctive approach to the predictive processing account called the action-oriented predictive processing account. The main aim of the action-oriented predictive processing account is to provide an account of cognition that unifies perception, action and cognition. Clark proposes that perception and action work together to fulfil the function of cognition and this is possible because perception and action follow the same computational logic. Clark's account is largely influenced by the work of physicist Karl Friston who proposes that the sole imperative of biological systems is to reduce free energy.<sup>1</sup> The brain, as such a biological system, functions to

---

<sup>1</sup> The reduction of free energy is equivalent to maximizing entropy. Free energy, in thermodynamic terms, is the difference between the energy and the entropy of a system. This is discussed further in section 4.

reduce prediction error and therefore to minimize free energy. Because the account is committed to the view that minimizing free energy is the sole imperative of biological systems, it faces a notable objection.

The objection takes shape in a *reductio ad absurdum* argument and is expressed as the dark room problem. The dark room problem says that if living organisms have the sole biological and neural imperative to minimize free energy, then they should follow a strategy that permits them to successfully minimize free energy.<sup>2</sup> One such strategy is to seek out a monotonous environment and avoid activity because by blocking out all sensory information and refraining from activity, the organism will not be exposed to surprise and will, therefore, have no prediction error to reduce. In other words, by committing to the view that biological organisms have the sole biological imperative to minimize prediction error, one is committed to the view that organisms will seek out monotonous environments that are free from surprise and activity. This strategy, however, is biologically nonsensical; how can an organism feed, reproduce and survive if it performs no actions? I describe this objection in section 7 and examine possible responses to the objection in section 7.1. and section 7.2.

Friston responds to the objection by saying that the dark room is, in fact, just a metaphor and that what he means by using the term is that organisms seek out the least surprising environment in relation to the environment to which they have adapted. He goes on to say that this strategy is not as absurd as it is made out to be because dark room agents do, in fact, exist and he provides the example of Troglaphiles (organisms that complete their life cycle in a dark cave).

---

<sup>2</sup> It may be said that internal states should be included in the minimization of free energy and that agents act when they are surprised by their own hunger, thirst or boredom. This is one aspect of Clark's response to the dark room problem, and will be discussed in section 7.1.

Friston's response is not satisfactory because what is required is a general account of cognition that tells us something about how *all* living organisms function and not a select few. Clark's response to the dark room problem is different to Friston's in that he proposes that the consequences suggested by the *reductio ad absurdum* objection do not follow.

Clark suggests that organisms expect activity and exploration and will therefore be prompted 'out of the dark room'. Clark's response to the dark room problem provides a theory of action and proposes that animals will be surprised by their own inactivity and will act on the world to reduce this surprise. An animal will, for example, be surprised by internal states and will act to reduce this surprise. It seems that Clark takes seriously the view that living organisms act only to reduce surprise because animals that live in complex and changing environments expect change and exploration and will therefore be surprised by a monotonous environment and by their own inactivity. Although Clark's response to the dark room problem is more probable and reasonable than Friston's, it does not suit what we know about the evolutionary and biological imperatives of animals. Evolutionary accounts of animal behaviour suggest that animals act because they are motivated to act by internal and external signals. Being surprised by inactivity is not enough for an animal to act; this criticism is discussed further in section 7.1. Clark offers no account of the motivations that drive living organisms to seek change, initiate motion and engage in exploration (Huebner 2012). Furthermore, Clark does not offer an account that explains how animals discriminate between better and worse strategies of action. In other words, Clark's account of cognition neglects any discussion on action selection and the motivations that drive an animal to action. Motivation and efficiency of behaviour are important elements of the efficiency problem of action selection and involves discussion on the adaptive fitness of living organisms. Clark's

account offers insight into a solution to the architecture problem of action selection but his account is not sufficient to resolve the efficiency problem because he does not pay enough attention to discrimination between better and worse strategies of action; the paucity of attention paid to action selection and motivation causes Clark's account to run into the dark room problem.

In response to the problem faced by Clark's account, I propose neuroeconomics as an alternative approach to an account of cognition. Neuroeconomics offers valuable insight about how animals discriminate between better and worse strategies of action and how value is represented in the brain. This means that a neuroeconomic account of cognition provides solutions to both the efficiency problem and the architecture problem. In section 8., I introduce neuroeconomics and examine how the discipline provides a solution to the efficiency problem. In section 8.1. up to and including section 8.3., I offer the economic theoretical framework, described by Glimcher, as a solution to the action selection problem. The suggestion made by neuroeconomists is that decision making and action selection are driven by the expectation of the outcome of a decision (reward or punishment). This suggestion is sympathetic to Sterelny's view that agents act because there is an observed relationship between action and outcome (Sterelny 2003: 83).

The neuroeconomic view, surveyed in section 9. and section 9.1., makes the argument that the value of an outcome is represented in the brain. It is suggested that value has a neural correlate, called subjective value. The objective of this discussion is to illustrate that the neuroeconomic account also takes into consideration the architecture problem because it provides an account of the architecture and mechanisms involved in action selection and decision making. In section 10, of this dissertation, I offer an example of how

neuroeconomics provides a complementary account to Clark's predictive processing account of cognition. It is complementary because it is proposed that value is encoded by a reward prediction error which is processed exactly as prediction error in Clark's account. This recommendation means that nothing is lost in Clark's account but a simple modification is made. Furthermore, by taking into consideration neuroeconomic views of value, the predictive processing account will not run into the dark room problem because biological organisms that function in the way suggested by neuroeconomics will have motivation to seek change and challenge.

### **1. The Environmental Complexity Hypothesis**

Developing tools for thinking about cognition is a challenge for many cognitive scientists because terms, such as 'intelligence' and 'cognition', do not have agreed-upon boundaries and definitions (Godfrey-Smith 2002: 233). Peter Godfrey-Smith offers one approach for thinking about intelligence and cognition; in his approach, he thinks about cognition from an evolutionary point of view and uses the terms 'intelligence' and 'cognition' in a broad and less demanding way than scholars who argue that intelligence is a uniquely human trait. Godfrey-Smith approaches cognition from an evolutionary standpoint and makes the claim that "...intelligence exists to *some* degree in a huge range of living systems." because "...the mechanisms that enable organisms to coordinate their behavior ... involve some degree of intelligence." (Godfrey-Smith 2002: 234). This is the starting point for the environmental complexity hypothesis which is aimed at providing a general account of cognition that emphasizes the adaptive value of cognition.

The environmental complexity hypothesis conceptualizes the idea that “...the function of cognition is to enable the agent to deal with environmental complexity.” (Christensen 2010: 364). This idea can be expressed in the two claims made by the environmental complexity hypothesis (Godfrey-Smith 2002: 237). The first claim is that the function of cognition is to control behaviour. The second claim is that controlling behaviour enables the animal to deal with changing and complex environments in efficient ways. These two claims are based on the supposition that all animals have sense and control mechanisms that enable them to track, and respond to, features of the environment (Sterelny 2003: 3). The capacity to interact with the environment enables animals to adapt their behaviour to match an ever changing environment. Flexible behaviour and adaptive response to a changing environment is beneficial to an animal when the benefits of appropriate responses outweigh the cost of response (Christensen 2010: 363). The complexity of the environment has an important influence on the animals that inhabit them and the greater the complexity of the environment, the more flexible and resilient an animal needs to be in its behaviour if it is to be efficient. Given this, one can expect that cognitive capacities will vary in sophistication and complexity, both between species in different environments, and within an animal. Some cognitive capacities are unsophisticated and simple, and may only elicit actions such as reflexes and behaviour tied to specific cues in the environment; other cognitive capacities are more complex, and enable planning of actions and reflection on the consequences of an action. The conceptualization of different levels of complexity in cognitive capacities is reflected in Kim Sterelny’s view. Sterelny supports the environmental complexity hypothesis developed by Peter Godfrey-Smith and offers an evolutionary account of cognition and action. In Sterelny’s expansion on the account of the evolution of cognitive systems, he proposes that even though many animals possess hybrid cognitive systems, there is a clear distinction between the different levels of complexity in cognitive systems

(Sterelny 2003: 11). In the following few sections, Sterelny's account of the evolution of cognition will be described. Surveying Sterelny's account will open up a new dimension of possibilities for discussion on action selection and decision making.

### **1.1. Simple detection systems**

Sterelny proposes that at the most basic level of cognitive complexity is a simple mechanism, called a detection system, which prompts a specific adaptive response to a particular cue in the environment (Sterelny 2003: 14). Detection systems are found in all living organisms - from bacteria which respond to chemical gradients in the environment (Sterelny 2003: 14) to humans who squint when suddenly exposed to bright light or who instinctively pull away from a hot object. Such ubiquitous response systems may be inherent or acquired by associative learning but are, nonetheless, always characteristic of behaviour that is consistently prompted by particular stimuli. The nature of detection systems suggests that for each registered feature of the environment, there is a single cue which consistently prompts a specific response or action (Sterelny 2003: 14). Despite the cognitive simplicity of detection systems, simple, cue-driven behaviour may be useful, for example, in some mating rituals or predator avoidance maneuvers; cue-driven behaviour is also cost effective because such simple responses do not require sophisticated cognitive mechanisms. An example of animals that make use of simple detection systems is the great crested grebe that performs a complex mating ritual that is initiated by simple, cue-driven responses to a mate. Sterelny (2003: 18) proposes that there is no reason to expect that each grebe is in possession of a representation of the mating ritual; it is more reasonable to suppose that each act is a cued response to the behaviour of the mate.

He illustrates this in the following example:

“The female begins by approaching a male, calling. The male responds by diving, at which point the female spreads her wings low upon the water, crouching. The male surfaces from his dive in front of her, as erect as possible, and revolves slowly until he faces the female. He then sinks slowly onto the water, still facing her, until finally the two face one another erect, shaking their heads.”

(Sterelny 2003: 18).

Despite being simple and largely cost effective, simple detection systems are subject to three weaknesses; I discuss each of these in turn. First, while some rituals and predator avoidance maneuvers may be successfully performed using cue-driven behaviour, it is important to keep in mind that the nature of the environment is not always predictable or trustworthy. To illustrate this point, Sterelny makes a distinction between transparent and translucent environments. A transparent environment is one in which the properties of the environment honestly and reliably correspond to the sensory cues that prompt behaviour (Christensen 2010: 367). An environment can become transparent to a living organism as a result of evolution; evolution selects organisms that adapt to the environment by making use of mechanisms that reliably receive and process information from the environment (Sterelny 2003: 21). This can be illustrated through an example of Adélie penguins that rely on, and respond to, the position of the sun in order to navigate through the environment. Adélie penguins migrate about 13 000 kilometers throughout the year as they “follow the sun” from their breeding colonies to winter foraging grounds (Rejcek 2010). The sun generates honest cues for navigation because the penguins have the right receptors and can respond appropriately. On an evolutionary timescale, the epistemic character of the environment can become more transparent and come to include more unambiguous and salient cues, but modification of the



environment can also take place on a shorter timescale as animals epistemically engineer their environments (Sterelny 2003: 22).

When animals engineer and epistemically modify the environment, the reliability of cues in the environment may be reduced (Sterelny 2003: 23). Unreliable cues in the environment indicate that features of the environment map onto detectable cues in complex ways causing the environment to become translucent (Christensen 2010: 367). A translucent environment is characteristically hostile in that it is made up of other living organisms that disguise themselves and deceive others, creating false negatives. For example, a bird may call an alert to the presence of a hawk when there is no hawk; this will frighten the other birds in the flock and the bird will have access to more food (Dawkins 1989: 64). Another example is mimicry which is found in many insect species to protect them from predators. These examples illustrate that biological agents present epistemic problems to animals that make use of only simple detection systems; biological agents can also hinder the efficiency of other agents when they bring about a hostile environment.<sup>3</sup> False positive information, for example, is costly to animals that are constantly alarmed by predators (Sterelny 2003: 26); it is costly for an animal to flee every time a cue that represents prey is observed because it will waste resources that could be used to feed, reproduce or flee when a real threat is posed. Imagine the birds discussed above; if the behaviour of birds involve responding to every alarm call by others in the flock, i.e. 'alarm call' is registered as 'flee', the birds will risk not having time to feed because they will be spending all their energy and resources to flee in response to the alarm call. The birds will also not have the energy required to respond to real threats. The unpredictable and untrustworthy nature of the environment is indicative of the first weakness of simple detection

---

<sup>3</sup> I use the term biological agent because mimicry and deception is not limited to animals but can also take place in plants.

systems. The weakness is that simple detection systems cannot detect when the agent is being deceived because it responds to particular cues only. Following is the second weakness of detection systems.

In addition to dealing with false negatives and false positives in hostile environments, detection systems expose animals to further weaknesses. The second weakness to be discussed involves mindless behaviour. When using simple detection systems, the animal is subject to behaving in a “sphexish” manner. “Sphexish” is a term attributed to behaviour that is mindlessly repeated or rigidly controlled by a particular cue in the environment (Sterelny 2003: 15); it is labeled as such because of the behaviour of digger wasps of the *Sphex* genus.<sup>4</sup> Digger wasps enter and inspect their nests before taking their prey into the nest; they leave the prey just outside the nest during inspection. Whilst the wasp is inside the nest, the prey may be moved, by an experimenter, to another location not far from the nest. When the wasp emerges, it will locate the prey, which has been moved, and once again leave it just outside the nest in order to complete another inspection of the nest. This procedure can be repeated several times without the wasp realizing that it is executing the same sequence repeatedly (Wooldridge 1963: 82).

The third weakness to be discussed states that animals, that make use of detection systems in hostile environments, are exposed to exploitation. Ants, for example, recognize and react to one another using chemical and mechanical cues which make it easy for other insects and spiders to exploit them (Sterelny 2003: 15). One example of the exploitation of ants can be found in the

---

<sup>4</sup> Douglas Hofstadter coined the term *sphexish* to describe mindless, pre-programmed behaviour. The term was made popular by Daniel Dennett in his 1984 book *Elbow Room: The varieties of free will worth wanting*.

relationship between the zodariid ground spider and ants (Korenko et al. 2013). The zodariid ground spider mimics ant behaviour in order to enter the nest; it walks on its three hind sets of legs and offers the front set of legs as antennae when coming face to face with ants. The spider gains protection from predators because ants are foul-tasting to many species that eat spiders but the zodariid spider also exploits the ants by depending on them for sustenance. The spider takes from the food the ants forage and also preys on the ants.

The three weaknesses described above indicate that in a hostile environment, action control and behaviour can benefit from more robust tracking of cues in the environment because single detection systems limit an animal's access to robust behaviour in environments that consist of predators that change the epistemic character of the environment. In other words, hostility in the environment means that complex cognitive systems and sophisticated behavioural responses will be more efficient and beneficial to an organism. Flexible and more sophisticated cognitive systems will be more efficient in translucent and hostile environments for two reasons. First, misleading information is usually found in environments that are hostile and unpredictable because prey and predators hide and disguise themselves. Second, the environment is complex and constantly in "...a number of different possible states which come and go over time..." (Godfrey-Smith 2002: 242). This means that animals that make use of detection systems in hostile environments have to deal with the cost of rigid, cue-driven responses because the animals are exposed to the presence of false negatives and false positives in the environment. It must be noted, however, that sometimes the cost of sphexish behaviour or being exploited is not high enough to prompt different behavioural responses. Plants, for example, cannot afford to have expensive cognitive mechanisms that consume large amounts of energy. However, as the

abovementioned examples show, limited responses to environmental cues may compromise an animal's behavioural efficiency. Because of this, many animals develop more complex response mechanisms because it allows them to function more efficiently in translucent and hostile environments. If cues that signal features in the environment are unreliable and have a high cost of error, animals will find it beneficial to develop and make use of more sophisticated and flexible cognitive systems and response mechanisms. In hostile environments, evolution will favour organisms that can respond adaptively to features of the environment given a variety of cues which may be unreliable and corrupted by other animals. Sterelny proposes that animals with robust tracking systems can respond to a feature of the environment given several cues (Sterelny 2003: 28). In the following section, I look at robust tracking mechanisms and how such mechanisms function to control behaviour when cues in the environment are unreliable.

## **1.2. Robust tracking systems**

In environments that are unpredictable and hostile, animals cannot solely rely on simple detection systems to control action and behaviour. This is because the cost of error when relying on a single cue to register features of the environment may be too high, or may lead to exploitation or selfish behaviour. Tracking the environment and adapting to environmental complexity enable animals to gain a reproductive advantage over others (Godfrey-Smith 2002: 242). Being able to respond to cues that signal features in the environment in a flexible way is favoured by selection when a particular cue is not sufficiently reliable (Sterelny 2003: 27). Detecting an environmental feature through more than one cue can provide greater reliability of information which is a condition for efficiency in translucent and hostile environments. It can, therefore, be said that a mechanism that can register a feature of the environment using several cues,

and that can trigger a suitable response, would be functionally more efficient. Sterelny calls such a mechanism robust tracking (Sterelny 2003: 27). An example of robust tracking is found in the behaviour of animals such as reed warblers. Reed warblers are faced with brood parasites, such as cuckoos, that abandon the task of raising their young. Cuckoos lay their eggs in the nests of warblers and then leave their young to be raised by the warbler. The reed warbler uses a variety of cues to determine whether to throw out an egg or not; these cues include egg size, temporal cues and the visible presence of cuckoos (Sterelny 2003: 28). Animals with robust tracking mechanisms have resilient and flexible behavioural responses which are more flexible and intelligent than simple detection systems (Sterelny 2003: 29).

Though robust tracking is more flexible and sophisticated than detection systems, the method is not error free: warblers, for example, may get rid of their own eggs. Robust tracking is limited in that there is a lack of response repertoire with behavioural options that are not tied to particular features of the environment. Robust tracking mechanisms are also limited because of the relationship between unreliable cues in the environment and behaviour. When cues in the environment are unreliable, they can trigger an inefficient response. This being said, hostile environments provide animals with an opportunity to track the complex environment using multiple cues which give them a reproductive advantage over animals that do not track their environment in a flexible and intelligent way (Godfrey-Smith 2002: 243). However, even though robust tracking allows the animal to track features of the environment using multiple cues, particular features of the environment are still tied to specific actions. This means that the animal has a limited repertoire of behavioural responses. When a meerkat gives a predator alarm call, "...it is arbitrary to translate its registration as 'Eagle Above!' rather than 'Run!' ..." (Sterelny 2003:

29) because there is a close connection between the feature that is detected and the action that is performed. Going back to the example of the relationship between the warbler and cuckoo eggs, the warbler responds to a variety of cues that determine whether the eggs in the nest are its own, but its response is tied to a specific action: reject. Godfrey-Smith notes that some organisms do not have the evolutionary option to “get smarter” and make use of flexible mechanisms that enable them to respond to environmental complexity (Godfrey-Smith 2002: 245). Some organisms are limited by their biological layout and overall ecology because developing complex response mechanisms come down to cost of error. This means that if the cost of complex computational mechanisms that detect cues in the environment is higher than the cost of error, it will be inefficient and impractical for more sophisticated mechanisms to develop. If sometimes throwing out their own eggs is not too costly, warblers will have no need to make use of better response mechanisms.

Flexible behaviour is advantageous and more effective in situations where the cost of limited response is significantly higher than the cost of flexible and controlled response because simple detection systems do not require expensive tracking mechanisms. In short, an animal needs to make a tradeoff between cognitive complexity and hostility in the environment so that its behaviour more efficiently fits the contingencies of its environment (Sterelny 2003: 3). For many animals, the cost of error ratio is low enough to only make use of simple detection systems but translucent and hostile environments have elements of complexity that require a mechanism more like robust tracking systems. For this reason, many animals possess hybrid action control mechanisms and respond to the environment using both simple response mechanisms and robust tracking systems (Sterelny 2003: 14). In complex and social environments, flexible and sophisticated tracking mechanisms, such as robust tracking, may not be

sufficient because it is still limited by the animal's response repertoire. This means that even though the animal can respond to the environment by detecting various cues, it has a limited response breadth. Response breadth refers to the available repertoire of behavioural responses which are registered when animals make use of the information that they gather from the environment (Sterelny 2003: 34). Animals that make use of only robust tracking and detection systems have a very narrow response breadth because their responses are tied to specific actions. Greater response breadth means a greater repertoire of available actions and thus more flexible behaviour. In complex and social environments, a greater response breadth is required because greater response breadth gives the animal options about possible actions. One system that, according to Sterelny, endows an agent with a greater response breadth is decoupled representation.<sup>5</sup> Sterelny suggests that decoupled representation is a mechanism that allows efficient tracking in hostile and translucent environments and increases the efficiency and behavioural flexibility of responses because features of the environments are not tied to specific actions (Sterelny 2003: 29).

### **1.3. Decoupled representation**

Decoupled representations are representations of features of the environment that are not tied to any specific action and are potentially relevant to various behavioural responses (Sterelny 2003: 31). Sterelny offers decoupled representation as the most complex level of cognition because representations of the environment that are registered using decoupled representation are not necessarily acquired from one cue but formed using multiple cues of features of the environment. These cues can be detected and stored to be used in a variety of situations, often in the future, and makes available a variety of possible

---

<sup>5</sup> I choose to use the term 'agent' rather than 'animal' or 'human' at this point because Sterelny is skeptical about the possibility that animals (other than humans) are capable of decoupled representation. I will make a statement about this issue in the following section.

actions to be used at any given time. The most important feature of decoupled representation is that responses are not tied to particular cues but are decoupled and can be used in various situations. Animals with decoupled representations - according to Sterelny, this definitely includes only humans and some primates - are able to track features of the environment and perform actions that are not tied to specific responses (Sterelny 2003: 31).

Sterelny supports the view that humans have the capacity to form decoupled representations but he believes that it is unclear whether non-humans have this capacity due to lack of evidence (Sterelny 2003: 50). He supports his presumption by providing two instances of experiments for decoupled representation in animals. The first involves the notion of moving through space. Sterelny proposes that it is unclear whether animals use cognitive maps (which would consist of decoupled representations) because animals may not have decoupled representations of space and the use thereof, and may navigate through space using procedural representations of spatial information (Sterelny 2003: 41). The second experiment involves the notion of tool use. Sterelny proposes that tool use may be the result of trial and error learning or it could be behaviour acquired from observing other animals as opposed to having decoupled representations of different ways in which an object can be used (Sterelny 2003: 46).

The main premises for rejecting the view that animals have decoupled representation are: 1) the experiments are limited because it is difficult to evaluate whether animals have decoupled representation when there are so many variables that could influence behaviour. The rat can navigate its way through a maze using procedural representations, by following a distinct smell of a reward at the end of the maze or by trial and error. 2) The second premise is



that behaviour that employs decoupled representations is indistinct (Sterelny 2003: 50). Current knowledge does not provide insight about which behavioural responses and processes require decoupled representation. A rat, in navigating its way through a maze, may use decoupled representation in finding its reward but could make use of automatic responses to navigate right or left. Though the question of whether animals have decoupled representation is important in the field of cognitive science, the discussion is beyond the reach of this paper. I will, however, accept the possibility that some animals (other than humans) possess the ability to form decoupled representations. This view is supported by scholars such as Griffin (2001) and Gallistel (1990).

Decoupled representation, and flexible and intelligent response mechanisms, facilitates efficient behaviour in social and complex environments more so than do rigid cue-driven response mechanisms, because the animal has a number of actions to choose from at any given time. As mentioned in the previous section, response breadth is an important element of behaviour in complex environments. Response breadth provides the animal with a repertoire of available actions; decoupled representation, which is a mechanism of behaviour in which action is decoupled from stimulus, is closely related to response breadth. As response breadth develops, options about possible actions and responses become more extensive (Sterelny 2003: 34). In other words, decoupled representation and a high response breadth imply that animals have a variety of actions to choose from. An example used by Sterelny is predator recognition and avoidance. When an animal recognizes a predator, it must assess the risk and decide upon an action (Sterelny 2003: 34-35). Sterelny supports Avital & Jablonka (2000: 119-122) in saying that "... (an) animal must decide between flight, concealment, a readiness to fight, or, if the risk is real but

not intense, continuing normal behavior but at a higher state of alertness”  
(Sterelny 2003: 35).

It is important that an animal selects the appropriate action in response to the environment to improve welfare and gain reproductive advantages over other animals. When an animal has available actions to choose from and have representations of the environment that are not tied to specific cues, response and action become less automatic (Sterelny 2003: 79). Decoupled representation enables an animal to select an appropriate action from a repertoire of available actions in a given situation and generates more efficient behaviour than response systems that detect cues that are tied to particular behaviour. At any given time, there are several actions that an animal is capable of performing but does not have the time to perform or may be too costly or impractical to perform at a given time. Sterelny’s account of the evolution of cognition offers an important account about how action control and cognition have evolved to generate efficient behaviour. Furthermore, Sterelny’s account offers important insight because it takes seriously the problem of action selection. His account leads us to ask an important question: how does an animal select the appropriate action in the situation and environment it is found? This is known as the action selection problem and has been a central topic of many papers and books.<sup>6</sup>

## **2. Action selection**

Selecting an appropriate action is very important to the efficiency and fitness of an animal (Houston, McNamara and Steer 2007: 1531) because at any given time

---

<sup>6</sup> The list is extensive. Some papers and books include Beer (1990), Brooks (1987), Firby (1987), Kristan et al. (1989), Ludlow (1980), Maes (1990, 1991), McFarland (1977), Seth (2007), Sterelny (2003), Todd (2001) and Winterhalder & Smith (2000) to mention a selected few.

an animal with decoupled representation has numerous available actions to choose from. This is central to the action selection problem which can be stated as follows: Given an agent with a repertoire of available actions, a number of goals, and knowledge about its internal states and external environment, how does an agent select the most useful or most valuable action in a given situation? (Maes 1990: 991; Seth 2007: 1545) The action selection problem constitutes two problems. First, it constitutes the efficiency problem which embodies the question, why are actions selected? Answers to this question may involve theories about motivation and desires as well as theories about action control. The second problem that forms part of the action selection problem is the architecture problem. This problem is aimed at asking the question, how are actions selected, or what mechanisms are involved in selecting an action? To answer this question, an account of the architecture, mechanisms and processes of cognition is required. I will return to the architecture problem in later sections but first I look at the efficiency problem and how the account provided by Godfrey-Smith and Sterelny helps answer the question, why are particular actions selected at particular times? This discussion is relevant to decoupled representation and response breadth in particular because simple detection systems constitute a particular response to a particular stimulus and therefore no action selection takes place.

### **2.1. The efficiency problem**

There are three features of the environment to take into account when considering action selection (Houston, McNamara and Steer 2007: 1533-1534). First, the outcome of an action provides information about future outcomes. Information is a resource that enables organisms to move through space and time. Sterelny proposes that this information, stored in the form of decoupled representations, is a “fuel for success” (Sterelny 2003: 31). The state of the

environment is a source of information which can be correlated with different variables. “When the state of a source of information is correlated with the state of another variable...the second variable carries information about the source.” (Godfrey-Smith 2002: 236). The information that is found in the relationship between the first and second variables can be the outcome of an action (reward or punishment). Reward or punishment, i.e. the outcome of an action, carries important information about the action that is selected and performed. This will be discussed in depth in section 8.1. up to and including section 10.1. of this dissertation.

Second, the environment contains other living agents which may influence the expected outcome of an action or the animal’s ability to perform an action. The selected action should be considered by taking into account the presence of other animals and other biological agents. Animals are constantly faced with dangers from predators and competition that hide, camouflage, mimic and deceive, and therefore epistemically pollute the environment (Sterelny 2003: 25). By taking other agents into account, more efficient behaviour can be prompted because better defenses against predators can be employed and more efficient detection systems will develop.

Third, the environment is in a state of constant change which makes cue-driven behaviour less than efficient. Cue-driven behaviour, if it is to be adaptive and efficient, requires that the cue be detectable and stable in the cue-world relationship (Sterelny 2003: 24). Environmental changes, and other living agents, affect specific cues and can disrupt the cue-world relationship. Recall, however, that flexible behaviour will only develop in place of simple response mechanisms if the cost of error is too high for efficient behaviour or survival. The three features of the environment that are important in action selection bring to light

two elements involved in the action selection problem that can provide insight about why particular actions are selected at particular times.

First, there are several possible actions to choose from and some actions are better than others in that they result in more efficient behaviour. Selecting actions that are better than others means the animal is more efficient in its behavioural responses than animals that do not select efficient actions. On a shorter timescale, animals with flexible behaviour can learn about better and worse responses to features in the environment; this can be achieved through associative learning. Animals with simple detection systems and robust tracking, however, may not be able to learn flexible behaviour and cannot change the responses that are tied to particular cues. In animals with simple detection systems and robust tracking, adaptation takes place and evolution selects those that are capable of survival and efficient behaviour. Efficient behaviour includes, for example, predator avoidance. A pigeon that has observed a predator is faced with several options. It can fly off the ground to avoid danger but risks being seen as vulnerable because it is travelling alone, or it can stay on the ground with its flock and rely on the protection of being in a group but this is a high risk option too (Dawkins 1989: 170). The best possible action would be to fly off the ground but make sure all the other pigeons follow so as to stay in the flock (Dawkins 1989: 170). This option involves risk because the pigeon endangers itself by making the alarm call but it is the best possible action given the other options. There is no reason to believe that pigeons have representations of the possible actions and their outcomes but evolution selects organisms that behave efficiently and have flexible behaviour when it is required. On a shorter timescale, if an animal does not select the action with the better outcome, it risks disadvantaging itself. If the pigeon chooses to stay on the ground so as to stay in the flock, it is selecting the action with little chance of

being better than sounding an alarm call. The pigeon disadvantages itself by not getting away from the threat. Another example is that of the warbler; the warbler that does not kick out the cuckoo egg is at a disadvantage because it is spending valuable resources on raising the young of another species. It is proposed that the action that should be selected is the better option given the other options.<sup>7</sup> This means that the selected action has the highest expected payoff given other available options and the context of the environment. Behavioural responses can be shaped over a variety of timescales and, on an evolutionary timescale, the agent may not know which is the better option but evolution selects agents that make better choices.

The second element of action selection is that each possible action is accompanied by consequences which are taken into account when actions are selected. Godfrey-Smith and Sterelny's evolutionary approaches to cognition and behaviour provide an account of action control that provides important insight about why actions are selected and can be offered as a solution to the efficiency problem. Sterelny proposes that many animals have a built-in hierarchy of motivation and that behaviour is driven by internal variables as well as external cues (Sterelny 2003: 81). In other words, he presumes that behaviour is a result of the motivation that arises from the combination of variables, both internal and external. Each outcome has a consequence which is translated as either a reward, or a punishment; an animal would learn to associate certain rewards and punishments with selected actions (Sterelny 2003: 82). A baboon that is faced with hostility from a superior has to make a decision regarding action. It should consider the number and location of its friends and those of his rival, the physical geography of the interaction and the value of a resource, to mention only a few variables (Sterelny 2003: 93). Most importantly,

---

<sup>7</sup> This also reflects the view of economists and neuroeconomists and will form part of the discussion on neuroeconomics in section 8 and onwards.

it should consider previous encounters and their outcomes. If the outcome of the selected action is rewarding, the action will be repeated when the opportunity arises, and if the outcome is not rewarding, the action should be avoided. In other words, is the interaction worth the reward of the resource given previous experience? The action selected by the baboon would be the one that generates the highest reward or lowest punishment given the variables and available actions. One would expect such flexible and intelligent behaviour from animals with decoupled representation, but could one expect the same from animals with simple response systems? The short answer is: probably. Experiments conducted by Tony Dickinson suggest that rats recognized the value of causation by selecting the action that in the past yielded a reward (Dickinson & Balleine 2002; Sterelny 2003: 37). This is an important discovery but it is, to date, unclear whether such appreciation of causation takes place in living organisms with cognitive systems simpler than those of rats. Though, one can expect that, through the process of evolution, animals that make bad choices will perish and eventually become extinct.

It can be said that Sterelny and Godfrey-Smith's approach to the action selection problem concerns efficiency which is an important element in selecting an action. Efficient and flexible behaviour gives animals a reproductive and competitive advantage over others. The approach offered by the environmental complexity hypothesis, and Sterelny's distinction between the different levels of complexity in cognition, provides an account of why particular actions are selected at particular times. Their evolutionary account of action control suggests that the function of cognition is to produce behaviour that enables agents to deal with complexity in the environment (Godfrey-Smith 2002: 235). Importantly though, action also depends on the internal environment of an animal (Sterelny 2003: 79). Internal cues are largely transparent and honest; an

animal can reliably know when it is thirsty or hungry. But, actions are not selected only on the basis of external environmental cues and internal cues. Sterelny proposes that animals also have motivational mechanisms such as preferences that play a role in action selection (Sterelny 2003: 83). Given this, Sterelny's account of cognition and action offers insight about the efficiency problem but does he offer a solution to the architecture problem?

The action selection problem needs to be approached with the goal to account for both dimensions of the action selection problem: the efficiency problem and the architecture problem. The two dimensions are closely related but often the focus of research is more extensive on one goal rather than the other. Sterelny attempts to unite the two views; he claims that his goal is to explore and understand the relationship between the wiring-and-connection facts and the interpretation facts (Sterelny 2003: 6). Wiring-and-connection facts refer to our internal cognitive organisation and how that internal organisation registers and tracks features of the environment; interpretation facts refer to our concepts of belief, preference and character states (Sterelny 2003: 5). The distinction between wiring-and-connection facts and interpretation facts roughly corresponds to the two dimensions of the action selection problem proposed in earlier sections. According to Sterelny, the set of wiring-and-connection facts includes "...facts about our internal organization (the wiring facts) and the facts about how that organization registers, reflects, or tracks external circumstances (the connection facts)." (Sterelny 2003). Interpretation facts can be said to be about our folk conceptions and beliefs (Sterelny 2003). These two sets of facts correspond roughly to the two dimensions of the action selection problem because the wiring-and-connection facts concerns the architecture involved in action and cognition and the interpretation facts involves beliefs and desires.



The two dimensions of the action selection problem are: the efficiency problem which is concerned with why actions are selected and the architecture problem which is concerned with the mechanisms involved in how actions are selected, i.e. the architecture involved in selecting actions. In short, the action selection problem asks the questions: what actions are selected and how are these actions selected? (Seth 2007: 1545) Sterelny's account overtly answers the former question but his solution to the latter question is more difficult to extract because his account does not offer an account of cognitive architecture.

Sterelny goes into detail about the cognitive mechanisms involved in simple detection systems and robust tracking. He describes the functional properties of simple tracking and response mechanisms and provides numerous examples of such mechanisms in biological organisms (Sterelny 2003). Notably absent from Sterelny's project, however, is an account of the architecture involved in decoupled representation.<sup>8</sup> Sterelny focuses primarily on the *functional* properties of decoupled representation and states that his focus is "...not on the neural or the computational mechanisms that realize interpretative capacities" (Sterelny 2003: 9); yet his primary aim is to explain the relationship between interpretative facts and wiring-and-connection facts. In other words, even though he claims to consider both the interpretative (and intentional) elements of behaviour *and* the wiring-and-connection facts, he pays relatively little attention to the latter in human cognition. This being said, Sterelny's account of action control and selection is useful and can provide important insight into one dimension of the action selection problem. The architecture problem remains a serious topic of discussion and seeing that it is notably absent from Sterelny's account of cognition, I offer an account of cognitive architecture that fits Sterelny's account of action control and action selection. In the following

---

<sup>8</sup> One would expect such an account to be one of human cognitive architecture because Sterelny is skeptical about the possibility of decoupled representation in non-human animals.

section, I will explore an account of cognition that provides insight about the functions and mechanisms involved in cognition. The account that I examine is a general account of cognition that is aimed at providing a unified account of perception, action and cognition. The account is a suitable fit to Sterelny's project because it gives important insight about the functions and mechanisms involved in cognition; however, I am particularly interested in what it says about action control and action selection. For now, I put aside discussion on action selection and will return to it once an account of the architecture of cognition has been provided. The account that I propose as complementary to Sterelny's project is the predictive processing account of cognition.

## **2.2. The architecture problem**

In his 2013 paper published in Behavioral Brain Sciences, *Whatever next? Predictive brains, situated agents, and the future of cognitive science*, Andy Clark offers an account of cognition that one may consider an account of the cognitive mechanisms and architecture involved in cognition and action. Clark forms part of a school of thought that supports the predictive processing account of cognition and shares his view with physicists, such as Karl Friston (Friston 2002; 2003; 2010; 2011; Friston et al. 2009; Friston & Stephan 2007; Friston et al. 2012), and philosophers, such as Jakob Hohwy (Hohwy 2010; Hohwy et al. 2008). We can return to the question, how are actions selected, or what mechanisms are involved in selecting an action? Though Clark does not directly respond to this question, his account may offer a solution on the neuronal level and provides an account of cognitive architecture. It must be noted, however, that Clark does not consider action selection overtly and, in fact, mentions action selection only twice in his article.<sup>9</sup> Instead his account is one of action control

---

<sup>9</sup> Clark mentions action selection once on p.190 of his paper, *Whatever Next?*, only as a result of saying that minimizing prediction is the driving force behind action selection. The second

and perception. Clark's general account of cognition does not provide an answer to the efficiency problem because he does not take seriously enough the problem of action selection but Clark's account provides us with a possible solution to the architecture problem.<sup>10</sup> In other words, it tells us something about the mechanisms involved in cognition and action.

The predictive processing account is inspired by the ideas of Hermann von Helmholtz who depicted perception as a probabilistic, knowledge-driven process (Clark 2013: 182). The predictive processing account is characterized by the key notion that the brain uses prediction-driven processes "...to acquire and deploy hierarchical generative models of the hidden causes (sometimes called latent variables) that best explain the changing patterns of sensory input that impinge upon the agent." (Clark 2012: 760). The notion that perception is a process involving top-down inference contrasts with traditional models of perception that model the brain as a passive, stimulus-driven mechanism (Clark 2014: 23). These traditional views take the process of perception to infer from effect to cause, and model the brain as a feature detector (Hubel & Wiesel 1965; Biederman 1987; Riesenhuber & Poggio 2000). The brain as feature detector is not a model supported by those sympathetic to views such as the predictive processing account which holds that perception is the inverse of the traditional views of perception. The predictive processing account claims that instead of simply accumulating information from the world and building a model of the world, the brain tries to predict the sensory information in the world and attempts to match its predictions of sensory causes with actual sensory stimuli. This means that the brain functions by inferring the cause of sensory input and then comparing the actual input to the predicted input. The current model of

---

mention of action selection is to briefly raise a concern about the scope of the predictive processing account.

<sup>10</sup> See section 7.1 for discussion on action selection in Clark's account.

cognition that is widely accepted is that perception is a process involving both bottom-up sensory input and top-down cognitive processes (Clark: 2013 186). The activity of bottom-up sensory input and top-down cognitive processes is central to the predictive processing account of cognitive function.

Clark says: “(perception, cognition and action) conspire to move a creature through time and space in ways that fulfil an ever-changing and deeply inter-animating set of (sub-personal) expectations” (Clark 2013: 186). It seems as though Clark considers fulfilling expectations as a goal in its own right and the sole function of cognition and action. Clark may consider that reducing drives should be considered minimizing prediction error because an agent can be surprised by its own hunger or thirst and by acting on its drives, the agent is reducing surprise. This will be discussed in section 7.1. In the following few sections, I examine the predictive processing account of cognition as proposed by Clark. I will begin with discussion on the nature of the cognitive system and go on to discuss the function of cognition as proposed by the predictive processing account. I will then discuss the theoretical framework in which the account is grounded and consider how Clark’s extension of the account to action runs into difficulties.

### **3. Predictive processing model of cognition**

The predictive processing account is a general account of cognition traditionally used to explain the mechanisms and functions involved in perception, but it has recently been extended to include an account of action (Friston & Stephan 2007; Brown et al. 2011; Clark 2013). Clark’s take on the predictive processing account is novel in that he offers the “action-oriented predictive processing” account which constitutes the claim that action and perception follow the same

predictive principles. This follows from the presumption that perception and action are computationally related and work together to reduce prediction error. Prediction error is the result of a mismatch between bottom-up input and top-down predictions; this will be discussed in depth in coming paragraphs, particularly section 3.2. With regards to action selection, the account examined by Clark does not consider action selection but says that action control is driven by the minimization of prediction error (Clark 2013: 190). According to the predictive processing account of cognition, biological organisms have an imperative to minimize prediction error, and perception and action function to fulfil this imperative. Prediction error in the perceptual process is reduced by updating the hypotheses about the causes of sensory input; action reduces prediction error when agents modify or adjust the world. Friston (2010: 8) and Clark (2013: 185-186) both support the view that prediction error is minimized through selective sampling and active sculpting of the environment (Clark 2013: 186). I will return to the action-oriented predictive processing account offered by Clark in a few sections. First, it is important to understand the nature of cognition as proffered by the predictive processing account. In the following sections, I provide a descriptive account of the nature of the cognitive system as supported by the predictive processing account. This will provide important background information needed to critically examine Clark's action-oriented predictive processing account.

### **3.1. The nature of the system**

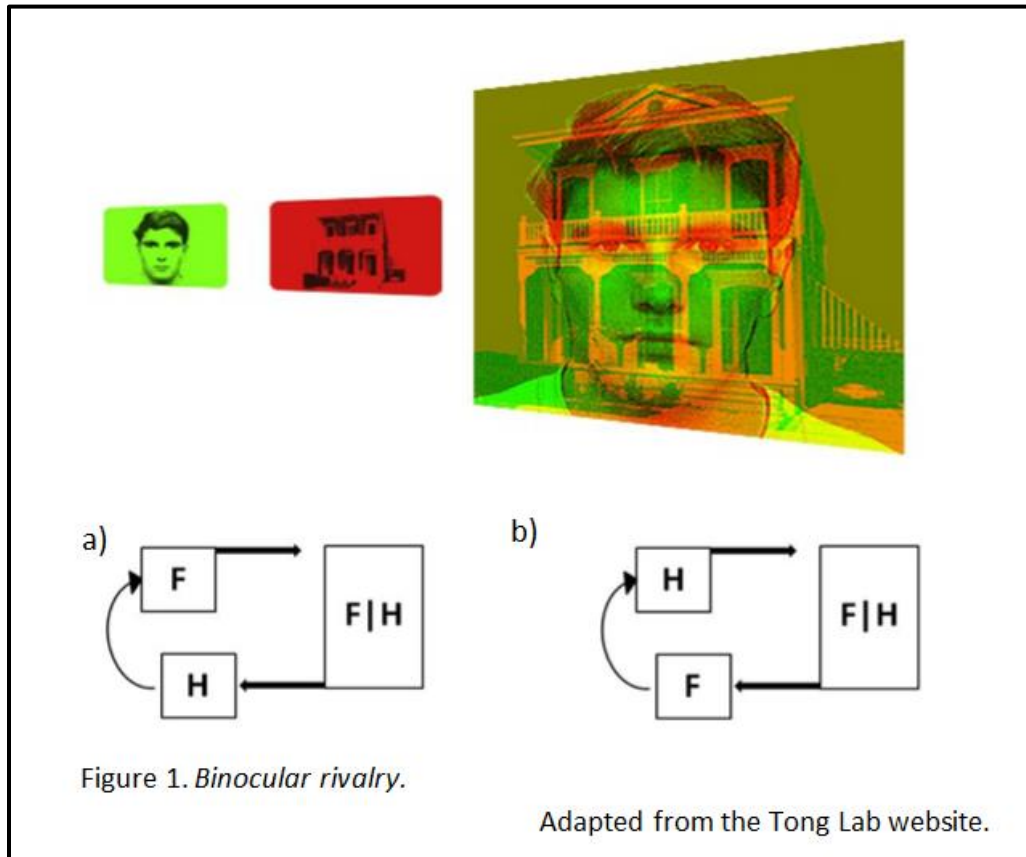
Clark's general account of cognition has three properties: 1) the system is bidirectional in nature. 2) The system is hierarchical in nature. 3) The system uses Bayesian inference in the selection of predictions/top-down information. I discuss each of these properties in turn.

#### **3.1.1. Bidirectional nature**

According to the predictive processing model of the brain, information is encoded using both top-down and bottom-up processes, and perception occurs through the encounter between bottom-up input and top-down predictions (Clark: 2013 186). In other words, cognition is bidirectional in nature (Clark 2013: 183). This view of hybrid cognition differs from many traditional theories of cognition in which top-down and bottom-up processes are considered mutually exclusive. Top-down information is predictive in nature, and functions to predict the causes of sensory input; bottom-up input constitutes stimuli and functions to update predictions. According to the predictive processing account, perception is a constant interaction between current and previous information and constitutes comparing bottom-up stimuli with top-down predictions (Clark 2013: 186). This process is known as “explaining away” and involves explaining away the incoming stimuli by matching it with expectations (Clark 2013: 187).

An example of this process can be found in an experiment that investigates binocular rivalry which is a visual experience that occurs when each eye is simultaneously presented with a different visual stimulus (Clark 2013: 184). The one eye is presented with a picture of a house and the other with a picture of a face. What is interesting about the experiment is that instead of seeing a merged picture, only one stimulus is perceived at a time (Hohwy et al. 2008: 690). The perceptual system switches between the two stimuli focusing on each picture, only for a few seconds. The top-down information will explain away only those elements of the incoming signal that match the current prediction (Clark 2013: 185) resulting in the switching of perceptual experience. This is an important experiment to help understand the relationship between top-down and bottom-up processing because the bottom-up stimuli remain constant but

the top-down information changes (Clark 2013: 185).



The information that matches the prediction is explained away and the difference between the prediction and input is transmitted as new information. This implies that perception does not necessarily include encoding the entire stimulus but encoding the difference between the sensory stimulus and a prediction. If the prediction is good, only the discrepancy between the prediction and the incoming stimulus will be perceived as input (Hohwy et al. 2008: 689). If the prediction is poor, the entire input signal is encoded as new information. This difference is known as prediction error and is reduced through updating top-down information. In Figure 1a), the prediction FACE is transmitted from higher levels to lower levels. FACE|HOUSE is the entire stimulus but FACE is explained away and only the prediction error HOUSE is transmitted backward. In Figure 1b), the prediction HOUSE is transmitted from higher levels to lower levels. FACE|HOUSE is the entire stimulus but HOUSE is explained away and only

the prediction error FACE is transmitted backward. This is a suitable account of why the subject's perception switches between the two stimuli; the prediction error that is transmitted backward updates the prediction and the other percept is experienced. It seems as though the generative model providing the top-down predictions is doing much of the perceptual work. However, the bottom-up signals provide ongoing feedback on top-down activity because perception is only guided by expectations and not enslaved by it (Clark 2013: 190). Therefore, the view put forth is not radically internalist but a hybrid view making use of both internal knowledge and external experience.

For example, consider a game ranger in a game reserve. If he is exposed to a grey object at a distance, top down hypotheses will predict what this object may be, perhaps a rhino, or an elephant, or a large grey vehicle. Until more bottom-up sensory information is received, it is just a grey object with no distinctive features. The system will deduce the probability of the object being a rhino, elephant or grey vehicle whilst taking in information, such as the context of the environment. Movement of the object may increase the probability of it being a living animal, so the prediction 'vehicle' is deleted, and the system is left with the possibilities, rhino and elephant. Bottom-up sensory information may communicate the size of the object and the prediction 'rhino' may be deleted in this process. Prediction error is minimized through the bottom up information received, that is, the movement and size of the object, and the object perceived is an elephant. Prediction error is minimized and the models of the environment are updated. An update could be formulated as follows: it is unlikely that there is an object of that size that moves like an animal, that is not an elephant.

### **3.1.2. Hierarchical nature**



This process occurs at various levels of cognition and supports the view that cognition is hierarchical in nature. A key function of the hierarchical system is to guess the next states of its neural economy - each higher level in the cognitive hierarchy predicts the informational state of the lower level (Clark 2013: 183). In other words, information in higher levels act as priors for the lower levels in the hierarchy. In the binocular rivalry case, "... a visual percept is predicted by a process of prediction operating across many different levels of a (bidirectional) processing hierarchy, each concerned with different types and scales of perceptual detail" (Clark 2013: 185). This process is explained by Hohwy:

"level 1 is the basic input level, level 2 tries to predict activity at level 1 and is also input level to level 3. Level 3 predicts activity in a top-down manner for level 2 and is paired upwards with level 4..."

(Hohwy 2010: 136)

This process constitutes backward updating of higher levels to keep prediction error to a minimum. Backward updating involves the lower levels informing the higher levels of any discrepancies between input and prediction. The nature of the levels in the hierarchy varies in that the basic input level encodes basic stimuli and higher levels encode more abstract and perspective invariant data (Hohwy 2010: 136). Each level along the cortex involves processes at varying timescales. For example, processes in the occipital lobe function, and are updated, on a shorter timescale than processes in the frontal lobe because processes in the frontal lobe involve more abstract information.

### **3.1.3. The Bayesian brain**

At each level of the cognitive hierarchy, interaction between top-down and bottom-up information takes place. Bottom-up information constitutes stimuli in the environment and top-down information constitutes predictions about the

causes of stimuli. How are these prediction generated? According to the predictive processing account, the brain performs Bayesian inference to select the models, or priors, that are transmitted downwards and that predict the causes of sensory input. This means that the prediction with the highest posterior probability is the prediction that is selected and that determines the perceptual content of the system (Hohwy et al. 2008: 688). In the Bayesian prediction model of the brain, supported by Andy Clark, stable, internal hypotheses have two important constituents: prior probability and likelihood (Clark 2013: 185). The hypothesis with the highest prior probability is the hypothesis that informs top-down predictions about incoming sensory stimuli (Hohwy 2010: 136). If the prediction corresponds well to the incoming stimuli, the likelihood of the input increases. High prior probability combined with high likelihood implies high posterior probability.

In the binocular rivalry case, first discussed in section 3.1.1., and illustrated in Figure 1, the prediction FACE and the prediction HOUSE have a roughly equal likelihood but if, for some insignificant reason, the picture of the face has a higher likelihood, then it will be selected for perceptual dominance (Hohwy et al. 2008: 690-691). The conjoint prediction FACE AND HOUSE have a higher likelihood because it can predict both the stimulus FACE and the stimulus HOUSE but it has a much lower probability because it is *a priori* very improbable to see a house-face or for the two objects to occupy the same space (Hohwy et al. 2008: 691).<sup>11</sup> The brain has learned that only one stimulus can cause sensory input at a given place and time (Hohwy et al. 2008: 691). In other words, given prior probability and likelihood, only one stimulus will be perceived at that time. The

---

<sup>11</sup> Hohwy's use of the term *a priori* is not the same as the traditional philosophical use of the term. Hohwy means that the brain has internal models of the world that function to predict the causes of sensory input. These models are *a priori* in the sense that they encounter experience after it has propagated down the hierarchical levels.

stimulus that is perceived is the one with highest likelihood given other contextual information.

Suggested by the predictive processing account is the notion that the probability density of information is encoded rather than the information (Clark 2013: 188). What this means is that when we perceive an object, the nature of the object is not represented in the brain but the relative probability of the nature of the object is encoded. This function enables the brain to deal with uncertainty, noise and ambiguity (Clark 2013: 188). For example, to encode the event TREE IN THE FOREST, the Bayesian brain will encode the conditional probability density function of this event using top-down and bottom-up processes and presents an approach where the general scene is first identified and the details are filled out (Clark 2013: 188). The general picture is presented and initially informed by internal hypotheses which are selected through Bayesian inference. In other words, perception constitutes a process that follows the following principle: forest first, then trees.

It may seem as though this is a cognitively and computationally costly and time-consuming process but this is not the case. Neurons are not suddenly turned on when new information is encountered because the brain is already active with a large set of priors which impact the processing of new information (Clark 2013: 188). This means that a large amount of context information is already in place when new information is encountered. The higher level processes influence activation in lower level processing but at each level of the neural economy, the predictions are probabilistic and take into account the uncertainty at each stage of the processing (Knill & Pouget 2004: 713). This illustrates the Bayesian nature of the cognitive system. When the system encounters new bottom-up information, the neuronal responses are expected to change significantly given

the contextual information provided by top-down predictions (Clark 2013: 189). This means that perception is a task of top-down predictions, more than it is of sensory stimulation. This view is supported by Rauss et al. (2011) who claim that "...neural signals are related less to a stimulus per se than to its congruence with internal goals and predictions, calculated on the basis of previous input to the system." (Rauss et al. 2011: 1249)

### **3.2. The function of the system**

The predictive processing account of cognition also makes an important claim about the function of cognition. According to the predictive processing account examined by Clark, perception is a process that involves top-down predictions interacting with bottom-up signals received from stimuli in the environment. This interaction results in a discrepancy between top-down predictions and bottom-up signals, called prediction error. This discrepancy/prediction error is transmitted backward in order to update predictions in higher levels of the cognitive hierarchy. The discrepancy that arises from the interaction between top-down and bottom-up processes, known as prediction error, is a result of the content that is entered into the system through sensory input but that is not predicted by higher level processes (Hohwy 2010: 135); it is the function of the system to reduce this error. The prediction that is transmitted is selected within a Bayesian framework and is the most probable prediction about the cause of the stimulus given the context. This means that the prediction with the highest posterior probability is selected and transmitted downwards; high posterior probability depends on likelihood and prior probability (Hohwy et al. 2008: 688). When prediction error is minimized, the accuracy of likelihood and prior probability of predictions increase and, as a result, posterior probability is increased. However, if the prediction is poor or absent, the entire input signal is transmitted as a prediction error and if the prediction is good, only the

discrepancy will be transmitted backwards. The claim made by the predictive processing account is that it is the function of the cognitive system to reduce this error by transmitting appropriate predictions. A process involved in prediction error minimization, is the modification and updating of higher-level hypotheses through lower-level input (Friston 2003: 1341). The discrepancies that arise as a result of the interaction between bottom-up sensory input and top-down predictions are involved in updating internal hypotheses. In addition to updating the hypotheses that feed predictions, prediction error can also be minimized through action.<sup>12</sup> This supports the claim that perception and action function to reduce prediction error and expresses the view that it is the function of biological organisms to reduce prediction error.

#### **4. The free energy principle**

The view that biological organisms function in ways that aim to reduce prediction error - by fulfilling expectations and matching predictions of sensory causes - can be expressed in terms of the free energy principle. The free energy principle states that all parts of a system that can change, will do so in order to decrease the information-theoretic quantity known as “free energy” (Friston et al. 2012: 1; Friston & Stephan 2007: 427). Free energy, as expressed in thermodynamic terms, is a measure of the difference between the energy and the entropy of a system (Friston & Stephan 2007: 419). In the domain of cognitive science as well as information theory, free energy is expressed as a measure of statistical probability distributions that is applied to the relationship between the biological system and the world (Friston & Stephan 2007: 420). It is important to note that what Friston has in mind when he speaks of a system, is any biological system; in Friston’s felicitous words: “from viruses to vegans”

---

<sup>12</sup> This refers to Andy Clark’s action-oriented predictive processing account and will be discussed in greater depth in coming sections.

(Friston et al. 2012: 3). Biological systems, such as bacteria, have an imperative to minimize free energy. *E. coli*, for example, move in a relatively straight line which is constantly interrupted by random tumbles. If the bacteria is moving towards an attractant, it is in a relatively stable state but when its direction of movement is changed by rotational diffusion, it encounters surprisal (or free energy) and will tumble sooner to find a new direction for movement (Friston 2011: 114). The bacteria do not have a representation of movement and do not choose the direction in which they move but, as a biological system, it functions to minimize free energy and responds to unexpected change in the environment (Friston 2011: 114). Friston's claim that biological systems function to minimize free energy can be extended to the neural domain because the brain is a biological system. Free energy is minimized in the brain when prediction error is reduced and it is the function of cognition to minimize free energy (Clark 2013; Friston & Stephan 2007; Friston 2010; Hohwy 2010). Free energy in the brain is minimized through updating the priors, or models, of sensory causes by transmitting prediction error. The free energy principle makes the claim that the brain minimizes free energy; from this follows that, living organisms with brains, or at the very least, cognitive capacities, have a biological and neural imperative to minimize free energy.

In the neural domain, free energy is represented by prediction error which is a result of the content that is entered into the system through sensory input but that is not predicted by higher level processes. In other words, prediction error reports surprise which is a result of the mismatch between bottom-up and top-down signals (Clark 2013: 183). This implies a strong bidirectional relationship between internal models of the brain and the external environment. Prediction error updates the internal models of features and states of the environment in order to minimize free energy. The free energy principle proposes that there is

both a neural and biological imperative to minimize free energy and that the basic function of cognition is to minimize prediction error and free energy (Hohwy 2010; Friston et al. 2012; Clark 2013; Friston & Stephan 2007). Cognition evolved because it is useful, efficient, and gives some animals advantages over others but one can expect Friston to say that cognition evolved because it enables some organisms to minimize free energy more efficiently. This argument seems like a tautology but becomes less ambiguous if an evolutionary approach is taken. The free energy principle, as proposed by Friston, suggests that there is a relation between free energy and evolutionary fitness and that behaviour can be unified under one principle: avoid surprises and you will last longer (Friston et al. 2012: 2).

The free energy principle, as supported by Friston, holds the view that a biological system that minimizes its tendency to enter into surprising states is more likely to succeed on an evolutionary scale. Biological systems can act on the world to change their position in, or relation to, it (Friston & Stephan 2007: 422). Through acting on the world, biological systems "...keep their exchange within bounds and preserve their physical integrity." (Friston & Stephan 2007: 423). This implies a strong form of the free energy minimization principle and is illustrated in the following example:

"A normal snowflake will fall and encounter a phase-boundary, at which its temperature will cause it to melt. Conversely, snowflakes that maintain their altitude and regulate their temperature may survive indefinitely, with a qualitatively recognisable form. The key difference between the normal and adaptive snowflake is the ability to change their relationship with the environment and maintain thermodynamic homeostasis."

Friston & Stephan 2007: 423

(This is a slightly odd example, among other reasons because snowflakes are not alive, and don't act or behave.) Friston goes on to say that the same general processes are at play in an evolutionary setting where systems that minimize free energy will be selected over those that do not (Friston & Stephan 2007: 423). The notion behind this is that if an organism minimizes free energy on a moment to moment basis, it minimizes free energy in the long haul and increases its adaptive advantages. For example, animals that avoid surprising encounters with predators are more likely to succeed than animals that are poor at avoiding surprising encounters (Friston et al. 2012: 5). Friston's free energy account is not limited to animals and he means that *all* living organisms function to minimize free energy. Plants also react to unexpected events, such as water scarcity or drought by accumulating the hormone abscisic acid (ABA) which triggers the closure of stomata when there is a lack of water. Stomata are pores in the leaf epidermis that allow gas exchange for photosynthesis and transpiration (Mueller-Roeber & Dreyer 2007). By minimizing the free energy created when exposed to an unexpected event, time after time, plants can be more efficient and survive longer than if free energy were not minimized. Friston proposes that the mechanisms used to react to surprising situations, and to minimize free energy, are results of the process of evolution and adaptation. Implied in this principle is the view that the process by which a living organism adapts to its environment is synonymous to the process by which that living organism minimizes free energy (Friston et al. 2012: 1). In the words of Friston: "...adaptive fitness and (negative) free energy are considered by some to be the same thing" (Friston et al. 2012: 2). Adaptive fitness and negative free energy can be considered synonymous because the minimization of free energy and prediction error results in better predictions and more efficient responses to the world. The biological imperative of organisms to minimize free energy means that the organism strives towards a state of negative free energy. If an organism's responses to the world are efficient, and free energy is minimized on



a moment to moment basis, adaptive fitness is more likely (Friston et al. 2012: 5). This means that living organisms not only adapt to the environment by minimizing free energy but minimize free energy by adapting to the environment (Friston et al. 2012: 1).

Furthermore, the minimization of free energy is concomitant with an increase in the posterior probability of representational models of the world because of decreased prediction errors. Living organisms that minimize free energy do so by reducing their tendency to be exposed to situations high in surprise. Plants, for example, reduce their tendency to be exposed to unexpected situations by making use of mechanisms such as mechanical and chemical defenses. In living organisms with more complex cognition, surprise can be reduced by updating representational models of the world and changing predictions of sensory input, or by acting on the world (Friston et al. 2012: 2). Representational models of the world are updated until prediction error is minimized to an expected level of noise (Hohwy 2010: 135). Given the close relationship between prediction error and updating of representational models of the world, one expects that minimizing free energy will correspond to maximizing posterior probability. The models that match the structure of the environment have a higher posterior probability than poor models of the environment but there is constant interaction and updating of models in higher levels. Constant updating of internal models in the brain means free energy is minimized and prediction error is reduced. Updating internal models is not the only manner in which free energy is reduced in living organisms with more complex cognition; the free energy principle, as supported by Clark and Friston, also proposes that action is utilized to minimize free energy (Friston et al. 2012: 2). This view is put forth by Clark as the “action-oriented predictive processing account” and has as its main claim

that perception and action follow the same deep logic and work in tandem to minimize free energy or prediction error (Clark 2013: 186).

### **5. Action-oriented predictive processing**

In his examination of the predictive processing account, traditionally used as an account of how prediction error is minimized in perceptual processes, Clark proposes that action also functions to minimize free energy. One way in which action enables the minimization of free energy is by changing how an agent samples the environment rather than changing its expectations of the environment (Hohwy et al. 2008: 690). In the example that was used in section 3.1.1., involving the game ranger's perception of the grey object in the distance, prediction error can be minimized through action in numerous ways. The game ranger can act on the world in several ways; he can, for example, move closer to the object, shine a light on the object, or make use of binoculars depending on the need. Other forms of action may include shifting branches or other objects that may obstruct access to stimuli. This illustrates Clark's claim that an animal moves its sensors and body, and acts on the environment in ways that can be described as active sculpting of the environment to match expectations (Clark 2013: 186). Acting on the environment influences what is perceived and how it is perceived. This is not the only manner in which the relation between action and free energy can be illustrated; prediction error, or free energy, also occurs in action. Action, it can be said, exhibits a similar profile to perception but reduces prediction error by eliciting movements that change bottom-up information (Clark 2013: 186). This means that prediction error is used to adjust action to minimize the discrepancy between the consequences of action and that which is predicted (Friston 2003: 1349). I will return to this process in a moment but first I look at how the free energy principle illuminates the intimate relationship between perception and action (Friston & Stephan 2007: 419). Clark's action-

oriented predictive processing account makes two claims about perception and action. The first involves the relationship between perception and action and states that perception and action work together to move an organism through time and space (Clark 2013: 186). The second is about the nature of perception and action and states that perception and action follow the same principles; this is based on the premise that both perception and action are implemented through the same computational strategies (Clark 2013: 186).

The first claim that perception and action do not function individually but rather work together to reduce prediction error is supported by Clark as a result of the view that prediction error is minimized through the selective sampling and active sculpting of stimuli (Clark 2013: 186). Perception and action function cooperatively to minimize prediction error through active inference which involves an agent moving its sensors to actively shape the environment so that it matches expectations (Clark 2013: 186). In other words, action is used to align sensory experience with the expectations of sensory experience. Perception functions to update internal models and expectations about the causes of stimuli, and action works to fulfil these expectations. Friston illustrates:

“Perceptual learning and inference is necessary to induce prior expectations about how the sensorium unfolds. Action is engaged to resample the world to fulfil these expectations. This places perception and action in intimate relation and accounts for both with the same principle.”

(Friston et al. 2009: 12)

For example, going into an unknown area elicits predictions of what to expect in the new area, and during behaviour, free energy needs to be minimized. There is a constant interaction between acting on the environment and predictions.

Action may include responses that block out sensory input, or eye movements which directly affect how perception occurs (Friston & Stephan 2007: 428).

The second claim put forth by Clark's action-oriented predictive processing account is that perception and action are computationally similar in that perception and action follow the same computational logic. This claim is captured by Eliasmith (2007) who supports the notion because "The best ways of interpreting incoming information via perception, are deeply the same as the best ways of controlling out-going information via motor action..." (Eliasmith 2007: 380). What Eliasmith has in mind is something along the lines of a Kalman filter which functions by predicting system states and updating these predictions through incoming information. Action plays a role in this process because it enables control of action and navigation through the environment in order to update predictions and expectations. Implied in this claim is the view that both perception and action are driven by predictions and expectations (Clark 2013: 186, 200). This is because both perception and action function to fulfil the expectations of the system.

We can now return to the view, put forth earlier, that prediction error also occurs in action. Action not only functions to reduce prediction error in perception but self-suppresses so that bottom-up information match top-down predictions (Clark 2013: 186). The suggestion is that when an agent performs an action, the system uses feedback from bottom-up information to minimize free energy while completing the action. Imagine reaching for a light switch in a dark room. You reach out your arm to meet a wall but do not immediately find a light switch; bottom-up information is constantly interacting with top-down information to determine whether the action will efficiently and successfully get you to the desired goal state (to find a light and switch it on). Bottom-up

information supplies the feedback needed to move your hand in upwards and sideways motions until you find the light switch. In other words, the interaction between top-down and bottom-up information provides information about where your arm is in space. Is your arm too low to reach a light switch? Is it too far from the wall? Or are you reaching to a place where there is very unlikely to be a light switch? This example illustrates that prediction error plays an important role, not only in perception, but in action too.

Where does this leave us in terms of action selection? Clark does not discuss action selection in his article but he offers an account of action control by looking at the free energy principle's claim that the biological imperative of living organisms is to minimize free energy (Friston 2011; Friston et al. 2012). In terms of action, the claim is that an agent acts on the world and adapts accordingly in order to fulfil its expectations. Implied in this claim is the view that the selected action is the action that has the highest likelihood of fulfilling an agent's expectations or that can shape the environment in such a way that expectations are matched. For example, finding your way to the entrance of the room where you are most likely to find a light switch and switching it on. According to Friston, the function of the brain can be understood in terms of prediction-error minimization which accounts for perception, learning and attention (Friston & Stephan 2007: 418). This is the function of cognition because, according to the free energy account, living organisms have a neural and biological imperative to minimize free energy (Friston 2011; Friston et al. 2012; Clark 2013). In addition to this, the free energy account holds the view that living organisms that fail to minimize prediction error, will have inefficient responses to the world and will also lack appropriate mechanisms for action (Friston & Stephan 2007: 428 - 429).

Friston's account is one of information theoretic free energy but what might the motivation be to reduce free energy, or prediction error, in the brain? One response might be as follows: processing information is costly in the computational system and this cost increases when there is a large amount of information that requires sophisticated processing. Minimizing prediction error is in the interest of the living organism because it reduces the computational cost of processing sensory input. Computational cost is reduced when prediction error is minimized because what is transmitted is the discrepancy between the predicted and input signal. If the prediction is good, the prediction error is small and less taxing to process. Having good representations of the world and efficient responses to the environment is important to a living organism and is affected by prediction error which needs to be minimized (Clark 2013: 187). The result of unsuccessful minimization of free energy is ineffective mechanisms for action and perception (Friston & Stephan 2007: 428 - 429). If an organism cannot behave appropriately in its environment or fails to respond to threats in an efficient and cost effective manner, the organism will not survive. Survival of the individual animal depends on behaviour that allows the animal to reproduce, feed and protect itself in an efficient and cost-effective manner. According to Friston and Clark, this is successfully realized through the minimization of free energy or prediction error. Minimizing free energy, according to Friston and Clark, is "...a necessary, if not sufficient, characteristic of evolutionary successful systems." (Friston & Stephan 2007: 428)

## **6. Summary so far**

The action selection problem is a serious concern for cognitive scientists and there are many ways in which the problem can be approached. There are two important questions asked by the action selection problem: first, why are actions selected? Second, how are action selected? Peter-Godfrey Smith and Kim

Sterelny offer evolutionary accounts of cognition and action that provide insight into why particular actions are selected at particular times. They propose that actions are selected because it enables an animal to deal with cognitive complexity and enables the animal to be more efficient in its responses to the environment. Sterelny proposes that there are distinct levels of complexity in cognition, and offers numerous examples of the architecture and mechanisms involved in simple action and response mechanisms. His project, however, lacks an account of the architecture and mechanisms involved in more complex cognitive systems, such as decoupled representation. In other words, Sterelny's account does not offer a solution to the question, how are actions selected? In response to the absence of an account of cognitive architecture, I propose the predictive processing account, recently examined by Clark in a Behavioural Brain Sciences target article. The predictive processing account offers a general account of cognition and offers insight into the mechanisms involved in cognition and action. Clark does not overtly consider the processes and mechanisms involved in action selection but his account makes three claims about the nature of cognition. First, the predictive processing account proposes that cognition is bidirectional in nature which means that cognition is a constant interaction between top-down information and bottom-up stimuli. Perception, Clark claims, occurs through the encounter between bottom-up input and top-down predictions; this process of interaction between bottom-up and top-down information occurs at various levels in cognition. This is the second property of cognition put forth by the predictive processing account, cognition is hierarchical in nature. Finally, the account proposes that the processing of information is Bayesian in nature; this concerns the selection of predictions in higher levels of cognition about the causes of stimuli. There is good reason to support these claims regarding the architecture and nature of cognition given experiments such as binocular rivalry which involve an experiment where the bottom-up incoming stimuli is constant but the top-down predictions alternate. Clark's

account provides insight that can be used to develop a theory of cognition about both action selection and the mechanisms and processes involved in cognition

The predictive processing account of cognition also makes an important claim about the *function* of cognition. Friston and Clark propose that the function of cognition is to minimize free energy; in the neural domain, free energy is reported by prediction error. It is the main function of the cognitive system to reduce prediction error, which is the discrepancy that arises as a result of the interaction between bottom-up input and top-down predictions; prediction error reports free energy in the brain. Friston proposes that living organisms have a biological imperative to minimize free energy; he also claims that adaptive fitness and negative free energy are considered by some to be the same thing (Friston et al. 2012: 2). Friston goes on to say that if a biological system (a living organism) fails to minimize free energy, it will encounter interactions with the environment that may lead to its demise (Friston & Stephan 2007: 434). Negative free energy may simply be understood as the state in which an agent's expectations are met or when there is a lack of surprisal. It is therefore the biological imperative of living organisms to reduce prediction error and minimize free energy because this increases adaptive fitness. Prediction error minimization can take place in two ways. First, the prediction error can be fed backward to hypotheses in higher levels which are then updated to decrease prediction error in the future. This follows the same logic as a Kalman filter that take into account different variables when forming predictions about unknown stimuli (Friston 2002; Grush 2004; Rao & Ballard 1999). The second way in which prediction error can be minimized is through action. Living organisms act on the environment to reduce prediction error by moving through time and space in ways that fulfil its expectations. From this follows that perception and action are driven by predictions and expectations because the main function of perception



and action is to fulfil expectations. In the next few paragraphs, I discuss an objection to the claim made by the predictive processing account and, thereafter, examine the responses by Karl Friston and Andy Clark to this objection known as the dark room problem.

## **7. The dark room problem**

As a result of saying that organisms have a sole biological imperative to fulfil expectations, Friston and Clark run in to a serious concern regarding the behaviour of living organisms. The dark room problem is a *reductio ad absurdum* objection which states that accepting the claim that all living organisms have the sole biological imperative to minimize free energy leads to a peculiar, and false, consequence. This can be illustrated as follows. If the predictive processing account is correct about the function of cognition, one would expect that all living organisms with cognitive capacities will behave in ways that enable them to be in a surprise-free environment. One way of achieving this is for an organism to block out all sensory information by seeking out the least stimulating environment and staying there (Friston et al. 2012: 2). In such an environment, a living organism will be deprived of sensory stimuli and the brain will have to make no predictions which means that prediction error will be minimal, if not absent. Mumford says:

“In some sense, this is the state that the cortex is trying to achieve: perfect prediction of the world, like the oriental Nirvana, as Tai-Sing Lee suggested to me, when nothing surprises you and new stimuli cause the merest ripple in your consciousness.”

(Clark 2013, citing Mumford 1992: 247n5)

This seems absurd since it is empirical knowledge that almost no living organisms inhabit monotonous environments or seek out situations that have no sensory

stimuli.<sup>13</sup> The notion of seeking out a monotonous environment finds expression in the “dark room metaphor” which implies that agents should settle in the nearest dark room to avoid surprise and movement. The dark room is a metaphor utilized by Friston to imply that organisms should seek out environments that are least surprising to them; that is, they need to find the dark room relevant to their environment. The dark room also illustrates an objection to the view that the function of cognition is to minimize prediction error.

According to the free energy principle, as advocated by Friston, living organisms have a biological imperative to minimize free energy (Friston 2011; Friston et al. 2012). This follows from the principle embedded in the free energy formulation: avoid surprises and you will last longer (Friston et al. 2012: 2). Seeking out monotonous environments that contain the least amount of surprises enables the organism to keep prediction error to the minimum (Friston et al. 2012: 2). An organism that successfully finds a monotonous environment and avoids surprise is in a stable state and therefore in a state of negative free energy. Mumford (1992: 247) describes the ultimate stable state as a state in which the top-down predictions perfectly predict activity at lower levels and when there is no error signal to be conveyed upwards. In such a state, there would be no need to update representational models of the world and no need for an agent to act on the world so that the world matches expectations. Friston suggests that adaptive fitness is expected from organisms in an ultimate stable state. Furthermore, in the neural domain, the function of cognition in living organisms is to reduce prediction error and minimize, or avoid, surprise through action and perception (Clark 2013: 194). According to the predictive processing account, it

---

<sup>13</sup> One exception may be Troglaphiles and Troglobites (animals that complete their entire life cycle inside a cave environment). Troglaphiles can also survive in above ground habitats unlike Troglobites who are often transparent (or white) and completely blind (Chapman 1982).

would make sense for organisms to seek out and stay in a dark room because the account is committed to the claim that in order to survive, surprise needs to be avoided. As a result, the predictive processing account is also committed to the view that organisms can successfully minimize surprise by blocking out sensory information and seeking out a monotonous environment, free from surprise. Such an environment would result in the organism being in the ultimate stable state with near perfect predictions (Clark 2013: 186). This view is peculiar because it is empirically false that living organisms with at least some cognitive capacities seek out dark, monotonous corners to inhabit; the view is also nonsensical from an evolutionary and biological point of view.

Consider, for example, the arctic ground squirrel that goes into hibernation in the winter months. Friston may consider a state of hibernation an attractive or stable state as there is little to no surprise in such a state. However, after a few months of hibernation, the squirrel leaves this stable state and explores the environment for food and mates. What is required of Friston's account, then, is to provide an account of free energy minimization, where organisms have an imperative to minimize surprise, but that takes into account the obvious fact that animals do not simply seek out a nice dark corner and stay in it (Clark 2013: 193). In the following sections, I look at Friston's response to the dark room problem and his attempt to better explain what is meant by the dark room metaphor. Thereafter, I examine Clark's response to Friston's account and also critically examine Clark's response to the dark room problem. Clark and Friston's responses to the objection have important differences but they both offer solutions to the consequences of the free energy principle. Their responses exemplify two general solutions that can be offered in response to the dark room problem. First, Friston refutes the *reductio ad absurdum* argument and supports the view that the consequences of claiming that minimization of free

energy as sole biological imperative are not as absurd as it seems. Second, Clark proposes that the consequences of the account do not follow and he provides additional considerations regarding prediction and expectation of surprising activity in changing and challenging environments.

### **7.1. Friston's response to the dark room problem**

Friston responds to the dark room problem by saying that dark room agents do, in fact, exist and that we need to look no further than Troglaphiles for such examples (Friston et al. 2012: 2). In other words, Friston refutes the *reductio ad absurdum* argument by saying that the free energy minimization principle and its consequences do not reduce to absurdity because some organisms do seek out monotonous environments that are free from sensory stimuli. Hence what Friston means when he offers the dark room metaphor is not that all living organisms should seek out a dark corner to stay in. Instead, what he means is that organisms should avoid surprise which is measured in relation to an organism's model of the world (Friston et al. 2012: 1). He goes on to say that solutions to minimize free energy "...will be unique to each conspecific and its econiche" (Friston et al. 2012: 2). For example, what is surprising to Troglaphiles may not be surprising to primates and *vice versa*. In other words, primates will find a dark cave very surprising and Troglaphiles will find sunny, open spaces very surprising. Should a Troglaphile then find itself in a sunny open space, it should seek the nearest cave and stay there. Similarly, should a primate find itself in a dark cave, it should seek an open space with light and stay there. The example is simplified but gets across the important point that surprise exists in relation to the living organism's expectations of the environment (Friston et al. 2012: 2).

Surprise can be reduced by sculpting the environment to change sensory input (action), by changing predictions of the input (perception) and, so Friston claims, by changing the living organism through evolution (Friston et al. 2012: 2). According to the free energy principle, living organisms have a biological imperative to reduce surprise - both on a short term and long term basis. What this means is that living organisms can minimize free energy on a moment to moment basis and as a result minimize free energy over their lifetime (Friston et al. 2012: 5). This, according to Friston, will increase adaptive fitness because there is an intimate relationship between minimizing free energy and adaptive fitness. Minimizing free energy to increase adaptive fitness entails changing the architecture of living organisms and endowing the organisms with prior beliefs through evolution. Friston endorses the view that free energy minimization takes place over multiple timescales and proposes that adaptive fitness and negative free energy are one and the same (Friston et al. 2012: 2). Friston illustrates his view by saying that "... (organisms) can minimize free energy on a moment to moment basis (through perception and action) and implicitly minimize their average free energy over their lifetime..." (Friston et al. 2012: 5). It seems as though what is biologically most important, according to Friston, is the minimization of free energy on an evolutionary scale. Lifetime free energy is minimized through acting in ways that minimize free energy on a moment to moment basis.

The reduction of surprise can be learned throughout the organism's lifetime or it can be innate which means evolution has endowed an organism with an appropriate mechanism to minimize free energy. Consider a rat in an experimental environment, that is repeatedly exposed to the sound of a buzzer. It will find the sound surprising at first but, if the sound is produced in consistent intervals, the rat will eventually expect the sound and it will no longer be

surprising. If the sound is produced in a random and inconsistent manner, the rat will continue to find the sound of the buzzer surprising. The former instance is an example of learning in a controlled environment; by learning to expect the sound of the buzzer, the rat is minimizing surprise because the sound becomes expected. Another example of surprise minimization can be found in the natural environment. The first time that a rabbit exits the burrow is an experience filled with surprises - it will experience its environment, other species, and predators for the first time. Rabbits frequently exhibit hiding behaviour but one would expect that hiding behaviour is more frequent in the initial few days after the first experience with the outside world. Eventually, the rabbit will become accustomed to features in its environment and will learn that some features in the environment are threats and should prompt hiding behaviour and some features are safe to be around. If learning does not take place and the animal is constantly hiding to avoid surprise, it will compromise its efficiency which could be detrimental to the animal's survival. The animal will not be able to reproduce and would die from hunger and thirst if it does not act on the environment - acting on the environment to feed and satisfy needs might elicit surprise.

It is not obvious (though Friston would disagree) that the sole biological imperative of living organisms is to minimize surprise because all organisms with cognitive capacities do not seek out and inhabit monotonous environments and, in fact, frequently participate in activities that elicit surprise. In other words, Friston's account cannot be accepted as a general account of cognition because not all organisms seek out monotonous environments. Friston's proposition that some organisms, Troglaphiles for example, seek out monotonous environments does not suffice. The free energy minimization account is an attempt to make a prediction about the function and nature of cognition. To illustrate the criticism of Friston's account more clearly: the account is not an attempt to make a

prediction about *some* animals that inhabit dark caves; it makes the claim that *all* living organisms function to minimize free energy. But, it has been illustrated, animals do not seek out monotonous environments that are free from sensory stimuli, and participate in challenge and exploration to satisfy needs and to reproduce. This implies that the risk of not leaving descendants is too high and animals get hungry and thirsty and are required to explore the environment to satisfy needs. Animals are also required to reproduce to promote the continued existence of the species - this often calls for challenge, exploration and exposure to surprise, even if just initially. These activities are valuable to the welfare and survival of animals because evolution eliminates inefficient mechanisms of action and organisms that are unsuccessful at reproduction; evolution, instead, favours organisms that are successful at reproduction and that are efficient and successful in their behaviour as response to the environment.

The account of minimizing free energy as sole biological imperative also lacks explanatory power that enables the account to abandon discussion on value and reward. The environmental complexity hypothesis described in section 1. illustrates that some living organisms possess hybrid action control mechanisms, which consist of simple detection systems, robust tracking systems and decoupled representation. These response mechanisms enable the living organism to respond to the environment in an efficient manner (Sterelny 2003: 83). If the free energy principle is correct in that living organisms have a biological imperative to minimize free energy, then one would expect that, through evolution, response systems would have evolved to be more than efficient in predicting challenges in the environment and that there will be little to no error in predicting the outcome of an action. Yet, in some organisms, the cost of error involved in automatic responses to the environment is not high enough to prompt development of more robust responses (Sterelny 2003: 13). It

seems that such organisms do not have the sole imperative to minimize free energy or surprise but rather function as a response to the environment to promote reproduction and efficient responses to the environment. In other words, Friston proposes that surprise is unique to each living organism and its environment, and is minimized in relation to each organism and its environment through perception, action and evolution (Friston et al. 2012: 2). However, it becomes clear when considering the environmental complexity hypothesis that Friston's view lacks explanatory power as the sole biological imperative of living organisms because it does not account for organisms with hybrid response mechanisms. Animals with more robust and flexible cognitive mechanisms explore the environment and act on the world in ways that expose them to surprise because they have the imperative to survive for as long as possible and to reproduce.

Minimizing free energy is therefore not the sole biological imperative of living organisms because evolution favours efficiency and, even more, efficiency in action selection. Efficiency is a serious concern for organisms that behave in, and interact with, the environment and many organisms can behave efficiently (enough) in response to the environment without minimizing free energy entirely. Some organisms have simple cognitive capabilities with relatively automatic responses that may be good enough for survival and reproduction. An organism will only benefit from building cognitive flexibility if such cognitive flexibility will be beneficial to the organism and if it is not too costly. Recall the discussion on cost of error in section 1.2., where it was stated that more robust response systems will only develop and be utilized if the cost of error significantly outweighs the cost of adaptive response. Most animals can behave efficiently, or rather efficiently enough, even if free energy is not minimized. Consider the digger wasp discussed in section 1.1. The wasp is acting in a



'sphexish' manner, behaving mindlessly, and has not adapted its behaviour. The wasp (both at individual level and species level) has not reduced surprise or free energy and continues to encounter the same experience. If Friston's account of the biological and neural imperative of living organisms is correct, the wasp should have adapted and developed more flexible cognitive abilities. However, the cost of error is not high enough and efficiency is not hampered to the extent that more flexibility and robustness in cognition is required. Some organisms do, however, develop more intelligent and flexible cognitive capacities which shows that there is some relation between free energy minimization and adaptive fitness but minimizing free energy is not the *sole* biological imperative of living organisms, as Friston proposes. Living organisms develop cognitive flexibility to enable efficient responses to the world given the nature of cognitive flexibility as described by the environmental complexity hypothesis, put forth by Godfrey-Smith and Sterelny. Godfrey-Smith and Sterelny propose that being able to behave in ways that are sensitive to environmental situations and being sensitive to a bigger repertoire of behaviour is beneficial to organisms in complex environments in which simple detection systems are costly.

Friston's response to the dark room problem suggests that the consequences of the free energy minimization principle is not as bad as it seems and he attempts to rescue the account by offering an example of organisms that do in fact inhabit dark environments free from sensory stimuli. It seems as though Friston misses the point of the objection because he comes up with some rationalisations for some organisms but that's not what is required. What Friston's account promises is a general account of cognition that makes predictions about the function of cognition and a rationalisation that can be generalised, but the promise is not delivered. In the next section, I discuss Clark's response to the dark room problem. Clark also refutes the *reductio ad absurdum* objection to

the dark room problem; he argues against the consequences and suggests that the free energy account does not run into the dark room problem because animals should not, and will not, seek out monotonous environments free from sensory stimuli because they expect to engage in change, challenge and exploration (Clark 2013: 193).

## **7.2. Clark's response to the dark room problem**

Clark's response is different to Friston's in that Friston confirms that the consequences follow the free energy account but he plays down the significance of the consequences; Clark claims that the consequences do not follow. Clark proposes that animals do not seek out monotonous environments free from surprise because animals expect change and exploration. In other words, in response to the dark room problem, Clark acknowledges that living organisms live in different and changing environments and proposes that organisms that live in changing and challenging environments will not seek out monotonous environments because they have expectations about change and exploration and furthermore expect to act out sophisticated strategies and perform cognitively complex actions (Clark 2013: 193). The potoo bird, for example, has adapted to a complex environment such as the Amazon rainforest. The potoo's actions reflect the complexity of the environment in that it has adapted to mimic its environment in order to blend in and protect itself from predators. The potoo's feeding and reproductive behaviours are strategic and unique to its environment. These birds expect exploration in their environments and have models and expectations about reproductive strategies and about the actions required in order to prey and protect. The notion behind this is that organisms have adapted to their environments and minimize surprise in relation to their environments, i.e. organisms in complex environments have adapted to complex environments and organisms in monotonous environments have adapted to

such conditions. From this follows that the world models and priors an agent holds influence that which is surprising to an agent, and the agent's prior expectations, and that which determines surprise, are specific to different species and are unique to the individual (Friston et al. 2012: 3). Clark's argument implies that animals are models of their environment and their predictions are formed to be compatible with the environments they inhabit. Implied in this claim is that when an animal adapts to its environment, it learns and updates its priors. Clark supports the view that perception and action are primarily driven by the need to minimize free energy. From an evolutionary approach, this is limiting; adaptive fitness depends not only on minimizing free energy but also on reproduction and drive satisfaction - these activities are efficiently performed when the value of the outcomes are learned.<sup>14</sup> Reproduction and drive satisfaction requires risky and changing behaviour, and does not merely require minimizing surprise.<sup>15</sup> Associating the reward or punishment of an outcome with an action enables animals to discriminate between, and prioritize, actions. A serious concern raised against the predictive processing account is that the appeal to incentive and reward to explain behaviour is largely replaced by constructs of prediction and expectation (Clark 2013: 200). Clark acknowledges the concern by asking whether the predictive processing account of cognition omits much of "...what really matters for adaptive success: things like boredom, curiosity, play exploration, foraging and the thrill of the hunt?" (Clark 2013: 193).

---

<sup>14</sup> This forms the focus of the following section on neuroeconomics and will be discussed in depth. For now, it suffices to acknowledge that the value of an outcome is important to action selection.

<sup>15</sup> Drive reduction is not the same as minimizing surprise because an initially surprising internal signal of hunger, for example, could be explained away simply by predicting further hunger. Acting to reduce drive - in this case by foraging and eating - has obvious fitness advantages compared to expecting starvation. If sections 8 and 9 below are correct, furthermore, many real animals are reward seekers, rather than drive reducers.

Clark responds to the dark room problem by saying that "...change, motion, exploration and search are themselves predicted - and poised to enslave action and perception accordingly." (Clark 2013: 193). He supports the notion that search and exploration are key features of life (Friston et al. 2012:1) and that animals explore the environment to feed and reproduce because without exploration, the animal will perish and the species will die out. Clark's response to the dark room problem suggests that animals that have expectations about change and exploration will find monotonous environments surprising. To illustrate this informally, the animal will be surprised by its own hunger and lack of activity and be prompted to explore the environment. Clark's response is adequate, as far it goes, in that exploring the environment is advantageous because the animal can feed, reproduce and find information about the environment but notably missing from Clark's response is an account of action selection and motivation. Clark does not propose that animals are motivated to satisfy hunger or to reproduce and neglects discussion on incentive and reward; instead Clark proposes that an animal is motivated to explore and find information so as to not be surprised by its own inactivity. The lack of discussion on incentive or reward in Clark's account is not suitable to an evolutionary or biological account of behaviour because it provides no insight into motivation to act and how to discriminate between strategies.

Survival and minimization of free energy in hostile and changing environments demand motion, search and action (Friston et al. 2012: 3). From this follows that, search, change and exploration are valuable to agents that inhabit complex environments. Clark offers a general account of cognition that suggests that animals expect to explore the environment and seek information. His account in response to the dark room problem is coherent and more conceivable than Friston's response but Clark's illustration of the predictive processing account

does not match what is known about evolution and biology. Sterelny's evolutionary account of action, described in section 1.1. - 1.3., illustrates that action is the result of both external states of the environment and a motivation to change or attain a particular state in the environment (Sterelny 2003: 79). Clark's study of predictive processing offers no account of the motivation that living organisms have to seek change, initiate motion and engage in exploration (Huebner 2012). Without such an account, the neural and biological strategy stays the same - reduce surprise (Friston et al. 2012: 5). Clark's response (that change and exploration is expected) is therefore not satisfactory because it lacks an account of a mechanism for selecting actions that, not only promotes adaptive fitness, but that encourages reproduction and discrimination between actions. Such a mechanism of motivation is required especially in an environment that delivers signals which are noisy and somewhat unreliable as well as functionally ambiguous (Sterelny 2003: 81).

Clark's general account of cognition lacks an account that explains how animals discriminate between strategies of action and how certain actions can be better or worse. It also lacks discussion on the motivations that prompt behaviour and does not offer an account of how the values of actions are represented. This is not to say that it is always necessary for the value of an outcome to be represented; motivation can be based on various internal drives and sensations (Sterelny 2003: 79). This means, for example, that water will have a higher value to a thirsty animal than to a satiated animal. Animals with robust and flexible cognitive mechanisms, such as decoupled representation, are able to update their goals according to the value of a reward (Huebner 2012). Animals with simple detection systems may act only on sensations and internal drives (Sterelny 2003: 79). The predictive processing account is largely correct about the mechanisms involved in perception and action but the account requires a

theory of value because learning, motivation and decision making are all processes that require valuation (Huebner 2012) and are important elements in the story of cognition. This means that although selecting the action that minimizes prediction error may fulfil an agent's expectations, adaptive fitness is dependent on more than just minimization of prediction error. A living organism increases the probability of adaptive fitness through generating efficient responses to the environment; this requires taking into account the value of the outcome of an action and the cost of acting to attain the outcome (Huebner 2012). Sterelny's evolutionary account of action provides insight into how cognition has evolved to enable animals to select actions that are efficient by taking into account the value of the outcome of an action and the cost of acting to attain the outcome. Though this insight is valuable, Sterelny does not offer an account of the mechanisms and architecture involved in the process of complex cognition.

To fill this gap in Sterelny's theory of cognition, I examined Clark's general account of cognition. Clark offers an account of cognition that provides insight about the mechanisms and functions of cognition, perception and action. The account of cognition faces a serious concern illustrated by the *reductio ad absurdum* objection, the dark room problem, which was discussed in section 7 of this dissertation. Clark responds to the dark room problem by asserting that the consequences of accepting the free energy minimization principle do not follow and that the account does not reduce to absurdity because animals expect change and exploration and therefore do not seek out monotonous environments. The response given by Clark does not fit what is known about evolution and biology because it is suggested by evolutionary models of cognition that animals seek reward and Clark's account neglects any discussion on reward and motivation. Reward-seeking is a fundamental property of

behaviour and have been recognized as such by many models of cognition (Barron, Sjøvik and Cornish 2010); notably, Sherrington (1906), Tinbergen (1951), and Lorenz (1965), and more recently Arias-Carrión, Ó., & Pöppel, E. (2007), Dickinson, A., & Balleine, B. (2002), Glimcher, P. W., & Sparks, D. L. (1992) and Glimcher, P. W., Dorris, M. C., & Bayer, H. M. (2005). Importantly, reward values are not merely attached to the outcomes of actions but are predictions that are updated by making use of experience (Spurrett 2012). By introducing the notion of value to Clark's predictive processing account, I do not refute his account but offer an improvement of the account because if Clark's account had taken into consideration the problem of action selection, and the role of value in action selection, the account would not run into the dark room problem. In the following sections, I set out an argument for taking seriously the concept of value in cognition and action. I begin the argument by describing the history of decision making and choice as described by Paul Glimcher in his 2011 book, *Foundations of neuroeconomic analysis*. Glimcher's illustration of the history of economic decision making with regards to value provides an appropriate theoretical framework to work from. Following discussion on the importance of value in decision making, I discuss an experiment conducted by Platt and Glimcher (1999) which illustrates that value is subjective and not found "in the world", but is encoded and registered in the brain. I conclude with discussion on reward prediction error which demonstrates that value can be integrated into the predictive processing account without refuting the account because the accounts are computationally the same.

## **8. Neuroeconomics: an account of value and decision making**

How do animals decide which actions are better or worse? Is there an objective model of value that can be used to measure the value of actions, or is the measure of value subjective? For an account that explains how decisions are

made, given available options, and how certain actions can be better or worse, I look to neuroeconomics. Neuroeconomics is aimed at providing models of decision making by relating economics, psychology and neuroscience (Glimcher 2011: 35). The goal of neuroeconomics (according to some) is to ground economic theory in neuroscience; this can be achieved through studying the economic theoretical framework and conducting experiments in neuroscience. Finding neuroscientific evidence that can describe choice behaviour will also enable us to make predictions about behaviour and decision making. What I hope to accomplish in the following sections is to show that a theoretical construct, such as value, can be related to a theory of brain function, and also provide us with a theory of action selection that will improve Clark's predictive processing account of cognition. Neuroeconomics provides important insight into behaviour and gives us an account of cognition that is computationally the same as Clark's predictive processing account. I, therefore, do not refute Clark's account of cognition but propose that the neuroeconomic account of cognition is complementary to Clark's account. Central to the model of decision making proposed by neuroeconomists is the concept of value which is a concept notably neglected in Clark's predictive processing account. The predictive processing account neglects discussion on action selection and value and, therefore, yields the dark room problem.

Neuroeconomics has the potential to provide a different perspective on the dark room problem because the neuroeconomic account of cognition does not neglect action selection. In fact, action selection and decision making are foundational concepts in studies of neuroeconomics which means that the account does not run into the dark room problem. Because neuroeconomics takes into account action selection it yields a constructive account that does not refute Clark's account but, if accepted as a complementary account of cognition,



prevents Clark's account from running into the *reductio ad absurdum* objection. The neuroeconomic account of cognition does not run into the dark room problem because action selection is an important part of the account; by seriously considering action selection, the neuroeconomic account implies that animals have motivation to act other than minimizing free energy or merely seeking activity to minimize free energy. Furthermore, neuroeconomics provides an opportunity to make Clark's account more powerful because research and ideas in neuroeconomics deeply complement the ideas that Clark offers about sensory and motor processing (Spurrett 2012). Value, it is suggested, is processed and registered as predictions and are reported by reward prediction error which means that the concept of value, as proposed by neuroeconomics, is computationally the same as the sensory and motor processes which are reported by prediction error in Clark's account.

Neuroeconomic models of decision making and behaviour are developed within the theoretical framework of economics. In economic theories about human behaviour, it is proposed that people are utility maximizers, which means that we shape our behaviour to maximize our utility (Edwards 1954: 382). Following classical economic models of decision making, utility is calculated by multiplying the value of an outcome with the probability of an outcome (Sanfey et al. 2006: 109). Given the existing economic theoretical framework, a theory of cognition that explains decision making can be developed. The neuroeconomic account of decision making proposed by Glimcher offers a suitable answer to the problem that Clark briefly considers in the closing paragraphs of his target article. Clark suggests that the predictive processing account lacks the explanatory power to abandon "...more traditional appeals to value, reward and cost..." (Clark 2013: 200). Even though Clark raises this concern, he does not spend much time discussing the notion of value, reward and cost. The lack of attention paid to

value causes Clark's account to run into problems, such as the dark room problem, because the account does not explain why organisms are motivated to seek change, initiate motion or engage in exploration (Huebner 2012). The neuroeconomic approach complements Clark's general account of cognition because it offers a theory of cognition that incorporates a theory of value and avoids the problems Clark's account runs into. Most importantly, the neuroeconomic approach offers an account that can explain how certain strategies of action can be better or worse by taking into account the cost of acting and the value of the payoff. In the following sections, the decision making models proposed by neuroeconomists will be examined as a theory of brain function that avoids running into the dark room problem because it provides a model for action selection. In order to understand neuroeconomic models of decision making and behaviour, I will begin by providing an account of economics which will describe the variables that play a role in decision making and explain why these variables are important. Starting this discussion with a theoretical account of value will establish the links required to relate economic models of decision making and neuroscience. My exposition of neuroeconomics is largely influenced by Glimcher and the account of the economic theoretical framework is essentially an account of economics as proposed and described by Glimcher.

### **8.1. Classical economic theories of decision making**

Glimcher discusses what he believes to be the most important ideas in decision making theory that can contribute to developing the field of neuroeconomics; starting with classical economic theory, he offers a general view of the economic theoretical framework. Classical economic theory begins with the *a priori* hypothesis that humans behave to maximize expected utility (Gauthier 1986; von Neumann & Morgenstern 1944; Bentham & Bowring 1843). Though the

notion of utility is usually attributed to Bentham (Bentham & Bowring 1843; Edwards 1954; Loewenstein 2000), Glimcher offers Pascal's model of decision making as a starting point for describing the history of economic models of decision making.<sup>16</sup> Pascal offered a normative model of decision making which proposed that a choice could be right or wrong. Choice behaviour, according to Pascal's theory, was guided by two variables: the magnitude of gain or loss and the probability of realizing the gain or loss (Pascal [1670] 2003; Glimcher 2011: 38).<sup>17</sup> Pascal proposed that the correct choice could be made by multiplying the probability of the event with the expected outcome of the event; one should then select the option that has the outcome with the largest expected value (Glimcher 2011: 38). An example illustrated in Glimcher 2011 is as follows: imagine a man has R45, should he pay R45 for a lottery ticket with a 50% chance of winning R100, or should he keep the R45? Pascal recommended that the man should select the action that yields the highest expected value (Glimcher 2011: 38) which is calculated by multiplying the probability of winning with the amount to be won. These calculations show that the expected value of buying the ticket is 50 whereas the expected value of not buying the ticket is 45, so the man should buy the lottery ticket.

Pascal's algorithm, however, was violated in many real life situations because many people seemed to select the option associated with the sure amount rather than selecting the option that was associated with the risk of losing (Glimcher 2011: 41). This evidence led Bernoulli to challenge Pascal's theory because it was unsuccessful at predicting choices under conditions of high risk (Glimcher, Dorris & Bayer 2005: 219). Bernoulli proposed a somewhat different model to resolve the problem encountered by Pascal's expected value theory; Bernoulli added two variables to Pascal's model. The first variable he added was

---

<sup>16</sup> It should be noted that Pascal does not make use of the term utility but his ideas lead to important developments in economic theory, including the use of the term utility.

<sup>17</sup> These variables are important in the discussion of economics and choice behaviour and will be returned to at a later stage.

the overall wealth of the chooser. Bernoulli proposed that gains were represented by a roughly logarithmic function of value that included the chooser's current wealth (Bernoulli 1954: 28; Glimcher, Dorris & Bayer 2005: 219). Adding a variable that could be calculated and measured strengthened the decision making model. The second variable added by Bernoulli was the concept of utility which replaced the concept of value used in Pascal's theory of choice; utility could not be measured because it was "...hidden inside the chooser..." (Glimcher 2011: 42). The value of an object or outcome is not objective and should be based on the subjective utility it yields (Bernoulli 1954: 24). Bernoulli's model of decision making differed from Pascal's theory of expected value in an important way; rather than multiplying the value of an outcome with the outcome probability as Pascal had done, Bernoulli suggested that we take the logarithm of the value to yield the hidden variable, utility, which can then be used to calculate the better option (Bernoulli 1954: 28). Illustrated in the 1954 translation of his work, is the following rule:

"Any gain must be added to the fortune previously possessed, then this sum must be raised to the power given by the number of possible ways in which the gain may be obtained; these terms should then be multiplied together. Then of this product a root must be extracted the degree of which is given by the number of all possible cases, and finally the value of the initial possessions must be subtracted therefrom; what then remains indicates the value of the risky proposition in question."

(Bernoulli 1954: 28)

Bernoulli's model of decision making is illustrated in a famous example taken from his 1738 paper published in St Petersburg. The example has been translated and paraphrased many times; Glimcher's interpretation is, for the most part, as follows (Glimcher 2011: 41-44): a beggar picks up a lottery ticket

with a 50% chance of winning a large sum of money. Suppose that sum is R20 000. If one multiplies the value of the outcome with the probability of the outcome, as suggested by Pascal, the beggar is in possession of an object with an expected value of R10 000 ( $R20\ 000 \times 0,5\%$ ). Imagine a wealthy man offers the beggar R7 000 for the ticket. The beggar has to decide whether he will select the option of R7 000 with 100% probability or the option that has a 50% probability of yielding R20 000. Pascal suggests that the beggar should not sell the ticket. In other words, the beggar should not select the option with 100% probability of gaining R7000 over the 50% probability of gaining R20 000. However, given the current wealth status of the beggar, it seems that the rational option for the beggar is to sell the ticket and take the definite R7 000; this is the opposite of what Pascal would recommend.

The beggar is faced with a problem: if he sells the ticket he is 100% sure that he will gain R7000 and if he keeps the ticket he risks a 50% probability of not gaining anything. Bernoulli proposed that two additional variables, the current wealth of the chooser and utility (a logarithmic function of value), be incorporated into a model of decision making (Bernoulli 1954: 28). Following Bernoulli's model, one comes to the following conclusion. The utility of the lottery ticket can be calculated as the logarithm of its value, R20 000, and is 4.3 utils. In Bernoulli's model this is multiplied by the probability 50%, 0.50, and equals 2.15 utils. The expected gain for the beggar, if he keeps the lottery ticket, is 2.15 utils. The calculated outcome if the beggar chooses to sell the ticket is the logarithm of the value of the sale, R7 000, so 3.8 utils. We multiply this by the 100% probability, so the expected utility is 3.8 utils. According to Bernoulli's model, it is rational for the beggar to sell the ticket to the rich man because the expected utility is higher if he chooses to sell the ticket (Bernoulli 1954: 24). This calculation can be also be used to calculate the utility, and expected utility, of the rich man given

the beggar's choice. Suppose the rich man has a starting wealth of R1 000 000, this is equivalent to 6 utils. If the lottery ticket wins, he adds R13 000 (R20 000 minus the R7 000 that he paid) to his wealth, and if it loses, he subtracts R7 000 from his wealth. The utility of these two possibilities can be calculated as follows:

$$R1\ 013\ 000\ (6.0056\ \text{utils}) \times 50\% = 3.0028$$

$$R993\ 000\ (5.9969\ \text{utils}) \times 50\% = 2.9984$$

The expected utility of the ticket is simply the sum of these two calculations, 6.0012 utils. Bernoulli's model suggests that buying the ticket increases the expected utility of the rich man and selling the ticket increases the expected utility of the beggar. It is, therefore, in the best interest of the beggar to sell the ticket because "...utility resulting from any small increase in wealth will be inversely proportionate to the quantity of goods previously possessed." (Bernoulli 1738: 25).

Bernoulli's model of decision making added several key elements to the contemporary model of decision making (Glimcher 2011: 45). First, probability is represented objectively. The probability that a coin will land tails up is 50%; this is independent of any subjective feelings or beliefs.<sup>18</sup> Second, people represent utility which includes the total current wealth of the subject and a logarithmic calculation.<sup>19</sup> Third, the product of these representations is calculated. Fourth, the option with the highest product of the represented values is selected. High expected value meant that utility was maximized which was something that was important in the choice behaviour of people. These four elements provide

---

<sup>18</sup>Subjective probability is based on personal judgments and beliefs and is unreliable in calculating measures such as utility. An example of subjective probability is estimating the probability that a sports team will win based on the opinions of the fans. A probability can still be calculated as a percentage but there is no observed measure that can be reliable enough to be used in calculating utility.

<sup>19</sup> Bernoulli and other economic scholars of the time were not overly committed to the logarithmic function as the only possible way to calculate utility (Glimcher 2011: 45n3).

important data to construct a good model of decision making and prompted the development of increasingly complex models of decision making in the two centuries following Bernoulli. The models of choice that followed Bernoulli's included more hidden and complex variables to predict behaviour and choice (Glimcher 2011: 45-46). Most of the models that were developed explored the predictions about choice behaviour that could be made by models of utility functions; the classification of utility functions differed too, ranging from natural logarithms to power laws and ratio scales. These increasingly complicated theories of decision making made many predictions about choice and were systematically related to ideas such as happiness (Glimcher 2011: 48).<sup>20</sup> The models that immediately followed Bernoulli's model of decision making had several weaknesses because they could be (and in fact were) replaced many times over and were also grounded in the concept of utility which, in the context of post-Bernoulli models of decision making, is a subjective quantity and can apparently not be measured. According to Glimcher, what was required, following Bernoulli's model of decision making was a model of decision making in which utility can be observed and measured. This task was tackled by twentieth century economists such as Pareto, Samuelson and Houthakker who were motivated by a new dimension of study; they knew that utility is subjective and that it was an important element in models of decision making. In the following section, I discuss neoclassical economic models of decision making as an important stepping stone to expected utility theory.

## **8.2. Neoclassical economic models of decision making**

The neoclassical economic movement was prompted by economists who questioned the ad hoc nature of classical economic theory. Neoclassical

---

<sup>20</sup> Notable ideas that related utility to happiness include Jeremy Bentham (1789), John Stuart Mill (1861), and Henry Sidgwick (1907).

economists questioned the classical models of decision making that followed Bernoulli because the models consisted of hidden and immeasurable variables. In the following two sections, I look at two neoclassical models of decision making, the weak axiom of revealed preference (WARP) and the generalized axiom of revealed preference (GARP). These models are discussed by Glimcher because they form part of the logical primitives of neuroeconomic thought.

### **8.2.1. Revealed preference theory (WARP)**

Bernoulli's model of decision making, and the economic models that immediately followed his, were faced with a problem: utility is subjective and apparently cannot be measured. In his 1938 paper, *A note on the pure theory of consumer's behaviour*, Samuelson was troubled by the classical, Bernoulli-style models of decision making because the models did not offer any predictions about utility given observed choice in real market situations (Mas-Colell 1982: 73). In the classical models, predictions can be made about choice given what is known about utility because if utility is a power function of magnitude, then in any situation a prediction can be made about exactly what choice will be observed (Glimcher 2011: 52). The problem that particularly bothered Samuelson was that this "...relationship is not invertible." (Glimcher 2011: 53). In other words, predictions can be made about choice given what is known about utility, but, observing choice does not allow predictions to be made about utility. In response to this problem, Samuelson brought to light the notion of revealed preference - a notion he proposed could provide insight about how utility can be measured.<sup>21</sup> What Samuelson was targeting was a theory of decision making that used observed choice to predict what can be known about utility rather than using utility to predict choice (Glimcher 2011: 52).

---

<sup>21</sup> The initial terminology used by Samuelson is 'selected over' but this terminology was replaced by 'preference'.



Samuelson's theory about revealed preference was grounded in the observation that some objects, or states, are 'selected over' others (Varian 2006: 100). Samuelson proposed that selecting x over y means not selecting y over x (Samuelson 1938: 65). This seems like an incredibly simple (and almost playful) statement to make but the statement makes important predictions about choice and utility (Glimcher 2011: 53). Samuelson's statement about revealed preference is illustrated more clearly when one considers that what Samuelson intends is to explain the notion of preference; **selecting** x over y means not **preferring** y to x. A useful distinction to keep in mind, at this stage, is between choice and preference. One interpretation of Samuelson's theory makes the inference that a subject prefers x over y if she selects x over y, but this is incorrect. Choice cannot be determined from preference because preferences are only constructs that summarize predictions about choice (Bernheim & Rangel 2008: 158). In other words, preferences cannot be observed, but choices can be observed. All that can be observed in choice is that the subject is choosing x over y; that the subject prefers x over y is merely an inference. It is inferred that the subject prefers x over y because she cannot prefer y over x if she selects x. This is based on the notion that when a subject is presented with a given set of variables, she will always choose the same set of items (Samuelson 1938: 63). The logic can be stated as follows: "If a subject is observed to **choose** A over B, then we can assume that he cannot **prefer** B to A" (Glimcher 2011: 54). This is the simple but powerful statement made by the Weak Axiom of Revealed Preference (WARP) which offers a theory of decision making that enables us to make fairly good predictions about choice and preference.

### **8.2.2. Generalized Axiom of Revealed Preference**

An important turn in neoclassical economic theories of decision making involved developing a theory that could make even better predictions about choice. In response to Samuelson's 1938 paper, economists developed the WARP approach. One notable modification to Samuelson's ideas is the theory offered by Houthakker. Houthakker proposed that preferences ought to be independent, continuous, reflexive, transitive, complete, and monotonic (Houthakker 1950).<sup>22</sup> The generalized axiom of revealed preference (GARP) proposes that observing the choice A over B enables inferences to be made about the choices that will be selected when a new set of alternatives is offered (Glimcher 2011: 58). This can be illustrated by an example taken from Glimcher (2011). Imagine a situation in which Peter is presented with \$6 and the following price situation: \$1 for apples and \$1 for oranges. Peter has to pick a fruit bundle of his choice for \$6. He picks fruit bundle A after which the price of oranges is halved and the price of apples is doubled. Given the changed price situation, Peter can now select a different fruit bundle (shown by the dotted line). Following the GARP model of choice, predictions can be made about possible new selections that Peter can make, given previously observed choices - selection of fruit bundle A (indicated by the solid line). GARP enables predictions about Peter's choice if an alternative is offered (indicated by the dashed line). According to GARP, he cannot choose anything to the left of point A because his previous selection revealed that he did not prefer that combination of fruit to A (Glimcher 2011: 58).

---

<sup>22</sup> Independence, continuity, transitivity and completeness are axioms of contemporary economic choice theory and are representative of rational choice. A utility function is monotonic when it follows the principle "more is better" (Glimcher 2011: 58n9).

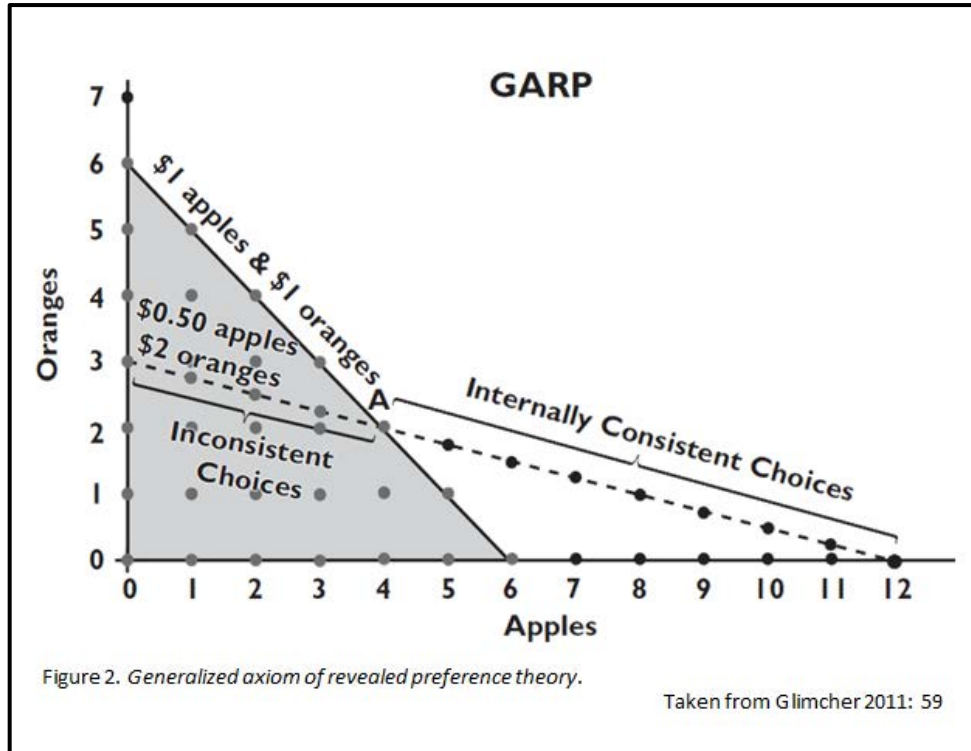


Figure 2. Generalized axiom of revealed preference theory.

Taken from Glimcher 2011: 59

GARP allows us to predict future choices based on previous choices and tells us something important about the utility function by adding the notion “more is better than less” (Glimcher 2011: 55). The notion that more is better describes what is meant by a monotonic utility function; the monotonic utility function increases as the quantity increases (Glimcher 2011: 58n9). The generalized axiom of revealed preference theory proposes that any chooser, who is consistent in her choices, behaves as though she is trying to maximize a weakly monotonic utility function (Glimcher 2011: 58) and that such a chooser with consistent preference is a rational chooser.

A rational chooser is, therefore, expected to have preferences as follows:

apples > oranges

oranges > pears

apples > pears

If Peter had inconsistent preference and, for example, preferred pears to apples, we can construct a situation as follows: imagine Peter has a pear and \$6 and is asked to participate in a series of trades.<sup>23</sup> First, he trades a pear and 1 cent for an orange. Then, he trades an orange and 1 cent for an apple because he prefers apples to oranges. We can now offer Peter his original pear for an apple and 1 cent. He will accept the offer because he has inconsistent preferences and he will be left with his original pear and be 3 cents poorer. Such trades can be made until Peter is left with no money and his original pear. The strength of GARP as a theory of decision making is in the fact that it is a theory about the representation of choice that can be falsified because if we observe a subject with inconsistent preferences, then we can conclude that the subject's model of choice does not rest on a monotonic utility function and the subjects preferences are not independent, continuous, reflexive, transitive and complete (Glimcher 2011: 62). Given this, GARP allows us to ask the question: "Is there at least one member of the class of all possible monotonic functions that can account for the observed data?" (Glimcher 2011: 62). If such a subject can be found, the theory can be tested and observable data can enable predictions to be made about choice.

### **8.3. Expected Utility Theory**

If a subject obeys the axioms of GARP in her choices, we can make predictions about that subject's future choices. GARP is successful in providing information about how one option is better than another (for a particular chooser) but GARP is limited in that it cannot tell us how much better one option is over another. For example, GARP can give an account of Peter's preference of oranges over apples but GARP does not give an account of how much better it is, for example, to have 100 oranges as opposed to 5 oranges. Is it 20 times better or 100 times

---

<sup>23</sup> Example adapted from Glimcher 2011

better? This notion can be extended to choices involving probabilities. For example, Peter may prefer a 75% chance of receiving an orange over a 50% chance of receiving an apple, but does he prefer a 35% chance of receiving an orange over a 33% chance of receiving an orange? What this illustrates is that theories such as GARP do not provide the tools to treat outcomes with similar probabilities as having related utilities (Glimcher 2011: 65). Describing choice under uncertainty was not a particular concern to classical economists but choice under uncertainty became an important concern for neoclassical economists such as von Neumann and Morgenstern who developed Expected Utility Theory. Expected Utility Theory describes utility more explicitly by studying choice under uncertainty (Glimcher 2011: 65). It was developed on the same core principles as GARP but with additional axioms. The additional axioms were added to provide a better theory of value that can make predictions about choice behaviour. Von Neumann and Morgenstern's proposal was similar to Bernoulli's proposal in that utility was an important element in choice. Von Neumann and Morgenstern's approach was different in that they paid more attention to studying uncertain or probabilistic events (Glimcher 2011: 65).

The first axiom that Von Neumann and Morgenstern added to neoclassical theories of decision making is continuity (Von Neumann & Morgenstern 1944: 73). This axiom can be illustrated as follows: if a subject prefers apples to oranges to pears, then the subject prefers a 100% chance of winning an apple to a 100% chance of winning an orange. The axiom states that even if a small chance of winning a pear is added to the certainty of receiving the orange, the subject's preferences should not change (Glimcher 2011: 66). In other words, the subject should prefer 100% chance of winning an apple over 100% chance of winning an orange plus 25% chance of winning a pear. Adding a probability of obtaining another object to the less preferred option should not change the

subject's choice. The continuity axiom will play an important role in constructing the argument about neural representations of utility.

The second axiom that is important to expected utility theory is that of independence (Glimcher 2011: 67). The axiom is based on the view that objects can be made up of separate things and states that if one adds a common prize to each option, the agent's preferences should not change. For example, if an agent prefers an apple over an orange, he must also prefer an apple plus an extra piece of apple over an orange plus an extra piece of apple (Glimcher 2011: 67). Also, if a subject prefers 50% chance of winning an apple over 25% chance of winning an orange, then he should also prefer a 50% chance of winning an apple plus 10% chance of winning a pear over 25% chance of winning an orange plus 10% chance of winning a pear. The claim made by von Neumann and Morgenstern is that if a chooser is consistent in her preferences and obeys GARP and the axioms of continuity and independence, it is equivalent to say that she has some monotonic utility function (Glimcher 2011: 67). A subject with a monotonic utility function behaves to maximize the monotonic utility function because the utility function is based on the idea that more is better (Glimcher 2011: 58n9). Such a chooser is considered a rational agent because in economic terms "...rationality is consistency and nothing more..." (Glimcher 2011: 61). A rational chooser, then, behaves as if she makes decisions by multiplying probability by utility.

Glimcher proposes that there are several limitations found in the neoclassical approach to decision making and that it leaves many behaviours poorly described (Glimcher 2011: 99). Allais, for example, developed an experiment that tested, and proved false, the independence axiom offered by Von Neumann and Morgenstern (Allais 1953). Examining the experiment conducted by Allais is not

particularly important to the scope of this dissertation but the findings of the experiment offer an historical insight into the development of neuroeconomics, as we know it today. Allais found that when offered two lotteries with different probabilities and magnitude of reward, the option that a subject chose depended on the subject's degree of risk aversion (Glimcher 2011: 101). When offered another set of lotteries that yielded exactly the same gambles, but were phrased differently, the subjects violated the independence axiom. There are more examples of behaviour that is poorly described by expected utility theory, including the Ellsberg paradox, the endowment effect and risk aversion. These behaviours will not be discussed here and it will be accepted as given that these behaviours illustrate the inadequacies of neoclassical theories of behaviour.<sup>24</sup> Glimcher proposes that one way in which we can overcome these inadequacies and develop better theories of decision making is by taking an inter-theoretic approach that acknowledges the strength of neuroscience and the models it can develop to describe choice behaviour while staying true to the neoclassical theoretical framework because the neoclassical theories and their sets of axioms capture choice behaviour well (Glimcher 2011: 100). This approach includes developing a theory of cognition that explains decision making processes at a neuronal level. Glimcher proposes that the concept of utility has a neural correlate: subjective value; he says that "Subjective values...are neurobiological objects..." (Glimcher 2011: 136) that are "...linearly related to utilities." (Glimcher 2011: 137). What Glimcher aims to do by making this claim is to develop a theory that allows examination of concepts at different levels in science. In other words, decision making examined at the behavioural level by economic theories can be tested at the neurological level by neuroscience.

---

<sup>24</sup> Glimcher 2011 discusses each of these behaviours in depth.

To qualify the claim that economic theories of decision making can be reduced to study in neuroscience, I look at experiments conducted by Platt and Glimcher which show that utility is represented in the brain by a neural correlate called subjective value. The discussion in the following paragraphs will examine the variables that play a role in how decisions are made when a subject is faced with several possible actions and outcomes. This view provides a cognitive account of how subjects discriminate between strategies of action and it explains how certain actions can be considered better or worse. The neuroeconomic approach to decision making offers a complementary improvement to Clark's account because it takes seriously the notion of value in action selection which is notably absent in Clark's account. I will return to the mechanisms and processes involved in representing value in later sections.

## **9. Neural representation of value**

Many traditional models used to describe action and behaviour are rooted in the classical view that behaviour is a reflex or response to the environment.<sup>25</sup> This view has been supported, and studied, by many scholars, particularly behaviourists, such as Skinner and Pavlov. The starting point of developing the theory of decision making, proposed in this dissertation, is the notion that behaviour is not merely a reflex or response to the environment but that behaviour is a result of taking into consideration the expected outcome of the selected action or choice and then selecting the action with the highest expected reward. Experiments by Glimcher & Sparks (1992), Schall & Hanes (1993), Basso & Wurtz (1997) and Platt & Glimcher (1999) are some of the many experiments that show that the processes that connect sensory and motor systems involve processes other than classical reflex mechanisms. What Glimcher offers is a

---

<sup>25</sup> Certain cognitive mechanisms do function in this way. A look at the Environmental Complexity Hypothesis, discussed in section 1. of this dissertation, make a distinction between such mechanisms and more sophisticated and flexible mechanisms of cognition.



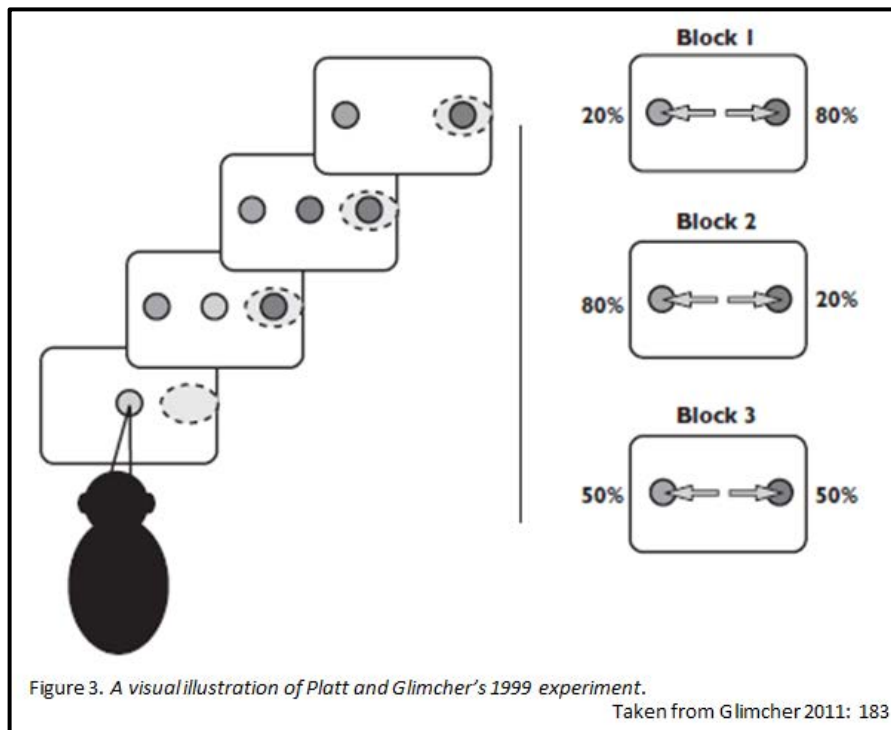
solution to the action selection problem that also provides an account of the neural architecture and mechanisms involved in action selection and decision making. The approach proposed by Platt & Glimcher (1999) has two classes of input: current sensory data and stored representations (Platt & Glimcher 1999: 233). From this follows that, decision making, according Glimcher's approach, involves a combination of current sensory data and the agent's best estimate of the outcome of an action. This view differs from Clark's in that the estimates of the outcome of an action are combined with a loss function that specifies the value of all possible losses or gains (Platt & Glimcher 1999: 233); these are represented in the brain and influence decision making.

It can, therefore, be said that the neuroeconomic account of cognition, as proposed by Glimcher, offers an account of why animals are motivated to act and how animals discriminate between better and worse strategies of action. Glimcher supports the view that animals act to maximize a utility function and goes on to say that utility has a neural correlate, which he calls subjective utility (Glimcher 2011: 136). Subjective utility is represented in the brain and an important element in action selection. The view that the value of an outcome is represented in the brain has been supported and tested by many scholars; of particular interest are the experiments done by Platt & Glimcher (1999), Kable & Glimcher (2007), Schultz (1998), and McClure et al. (2004). The scope of this paper allows discussion of only two of these experiments, Platt & Glimcher (1999) and Schultz (1998). In the following paragraphs, I will discuss the experiment developed by Platt and Glimcher and make use of their findings to show that subjective value is represented in the brain and plays an important role in decision making and action selection.

### **9.1. Evidence of neural representation of value**

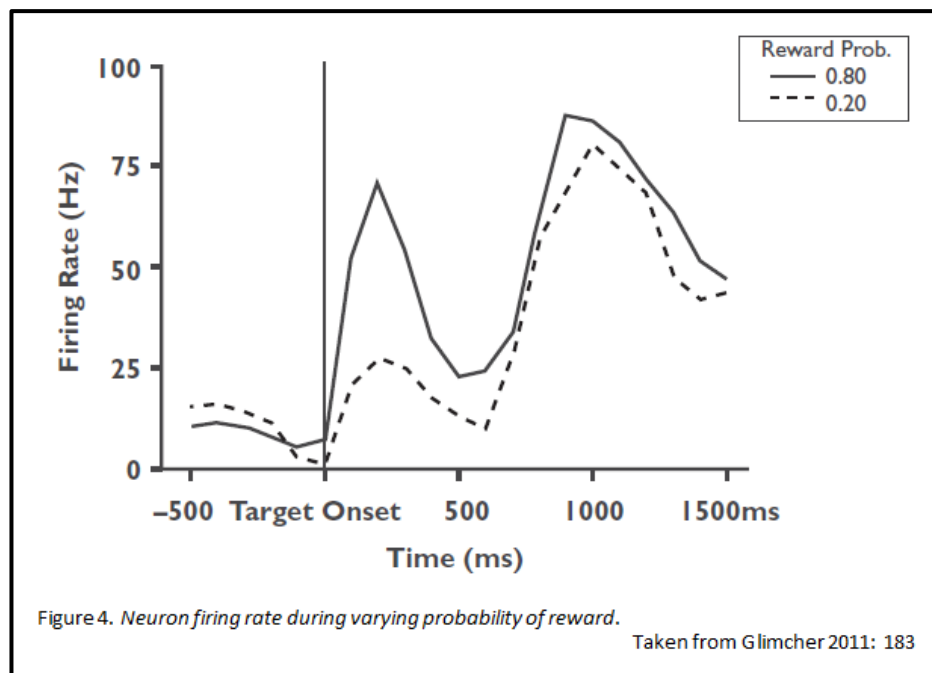
The set of experiments conducted by Platt and Glimcher (1999) were developed to examine whether decisions are made on the basis of variables, such as the probability of reward and magnitude of reward, and to examine whether these variables were related to behavioural responses (Platt & Glimcher 1999: 233). To examine the idea that neurons in certain brain areas - in the following experiment, area LIP - encode expected subjective value, Platt and Glimcher developed an experiment illustrating that the firing rate of neurons in area LIP was positively correlated with the probability and/or magnitude of a reward (Glimcher 2011: 182). In the context of this dissertation, the experiment is interpreted as developed to study the relationship between action selection/decision making and utility (rather, its neural correlate, subjective value). The study conducted by Platt and Glimcher consisted of two experiments developed with the aim to examine whether there is a link between neuronal activity in the area LIP and an animal's choices in an oculomotor task (Platt & Glimcher 1999: 234). Platt and Glimcher conducted experiments in the visual system because the topographical arrangement of neurons in the visual system is easier to translate even though it is predicted that information from our visual *and* other sensory systems is arranged topographically in the brain (Glimcher 2011: 164). What is meant by topographical arrangement is that neurons are arranged in such a way that neurons are activated in roughly the same spatial orientation than stimuli in the environment. Due to the nature of the visual system, it is easier to translate the topographical information than it is in other sensory systems. In the experiment, trained rhesus monkeys were trained to perform a cued saccade task in which the monkeys fixated on a yellow dot, while two other targets, one green dot and one red dot, were presented on the screen. After a random interval delay, the yellow dot changed colour to either red or green. Following this change of colour, the fixated dot would disappear and the animal will be instructed to select one of two possible actions in order to receive

a reward (Platt & Glimcher 1999: 234). The magnitude of reward and the probability of reward were varied in blocks of trials but the sensory and motor related variables remained constant. Figure 3 is a visual representation of the experiment.



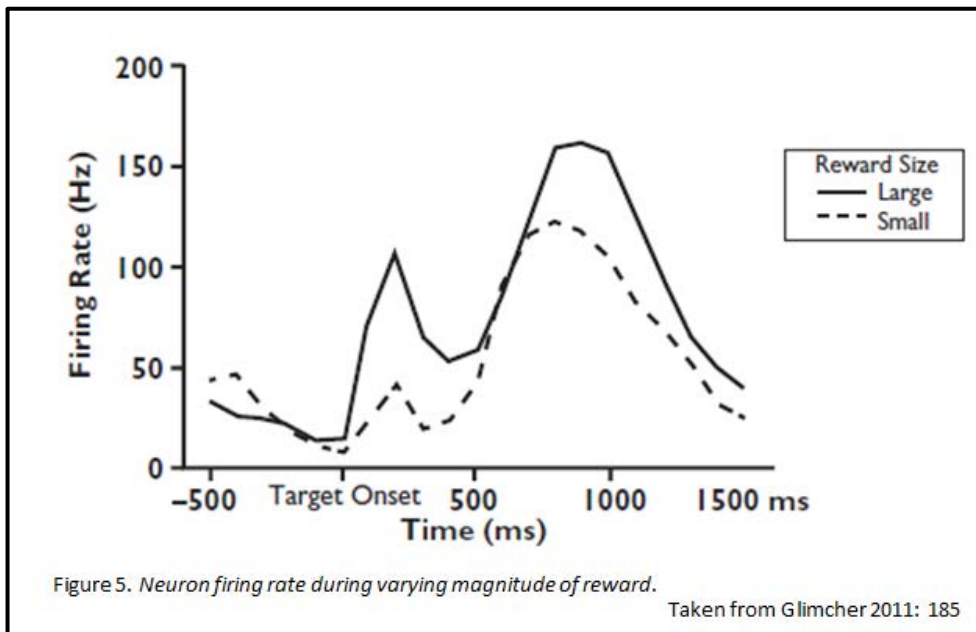
In the first set of the experiment, the monkeys were presented with a block of 100 trials in which the probability of reward associated with the green dot was 80% and the reward probability associated with the red dot was 20%. These probabilities were reversed in a second block of trials. What Platt and Glimcher found was that when there was a higher probability that selecting the red dot would be rewarded, neurons encoding movement associated with selecting the red dot responded strongly (Glimcher 2011: 183). Conversely, when the probability was low, the associated neurons responded weakly. This is illustrated in Figure 4 below; here, one can see that the firing rate increased significantly after target onset when the probability of reward was higher. To

explain, the solid line indicates that the reward probability is higher because there is a significant increase in the firing rate of neurons. The dotted line indicates a lower firing rate which correlates with a lower probability of reward. The aim of experiment 1 was to study the correlates of reward probability by examining the firing rate of neurons in area LIP (Glimcher 2011: 184). Platt and Glimcher found that the probability that a particular response will be rewarded systematically increase the activation of neurons in posterior parietal cortex (Platt & Glimcher 1999: 237).



In the second stage of the experiment, Platt and Glimcher explored whether neurons carried information about the magnitude of the reward (Glimcher 2011: 184). In different blocks of trials, they varied the amount of reward that the animal would receive when selecting a target. In the first block of trials, the monkey would have received 0,2 milliliters of fruit juice in trials where he correctly looked at the green dot and 0,4 milliliters of fruit juice in trials where he correctly looked at the red dot (Glimcher 2011: 185). These values were reversed in a second block of trials. What Platt and Glimcher found was that certain neurons fired more strongly during trials where the reward was higher

and the firing rate was weaker when the expected reward was lower. As illustrated in the graph below, the firing rate was significantly higher when the expected reward was large (indicated by the solid line). The firing rate is inversely inclined in the opposite case; the firing rate was lower when the expected reward was small (indicated by the dotted line). In other words, Platt and Glimcher found that when an animal is instructed to choose between actions, the gain expected from each possible action exerts a correlated influence on both the choice behaviour of the animal and the activation of posterior parietal neurons (Platt & Glimcher 1999).



One limitation of these experiments is that it examined a fairly limited range of choices. Platt and Glimcher acknowledged this limitation and developed the “free choice task” in which monkeys were allowed to choose freely with no reinforced targets. In the free choice task, red and green dots appeared in blocks where the expected value of each movement varied. Platt and Glimcher inferred that the probability that a monkey would look at a certain dot was

roughly equal to the expected value of the dot and that the neurons continued to track the expected utility of the two presented dots (Glimcher 2011: 186). The experiments conducted by Platt and Glimcher show that activity in certain brain regions serve as a topographic map of expected subject value of each action in the repertoire of available actions (Glimcher 2011: 187) and that “Instantaneous firing rates under all of these conditions would encode expected subjective value.” (Glimcher 2011: 188). This means that there was more activity in an area that predicted the occurrence of reward and less activity in the surrounding areas.

An important observation made by Platt and Glimcher was that neurons that represent subjective values fired before the related actions were triggered (Glimcher 2011: 186). This observation illustrates the importance of subjective value in decision making and action selection mechanisms. Following the neuroscientific evidence, it can be said that selecting an action is reduced to the process of identifying the movement with the highest expected subjective value. This view is in agreement with the economic theoretical framework discussed in section 8 and supports the idea that actions are selected to maximize a monotonic utility function. Most importantly, these experiments also provide the neural evidence required for an account of cognition that proposes that the brain represents subjective value (the neural correlate of utility), probability of the outcome and the product of subjective value and probability. The account also offers a description of action selection and proposes that there is a neuro-computational mechanism that selects the choice that has the highest expected subjective value (Glimcher 2011: 188). In essence, the monkeys in Platt and Glimcher’s experiments are solving an action selection problem and are faced

with two options from which they can select only one action.<sup>26</sup> The monkeys can choose to look at the target on the right, the target on the left or they can choose not to look at either; in other words, they have a variety of possible actions to choose from. Glimcher makes the cautionary note that, at this stage, the evidence can be observed in at least one brain area but that the brain is interconnected and therefore one can expect that brain areas such as the FEF, superior colliculus and area LIP work with other brain areas to encode subjective value and to select the action with the highest expected subjective value. It is expected that studies in other brain areas will harvest similar results, or at the very least supporting results. Given the current evidence, we can develop a theory of cognition that explains how value is represented in the brain. Experiments by Mirenowicz & Schultz (1994) (to be discussed in the next section) indicate that the account of cognition proposed by neuroeconomics is largely the same as Clark's general predictive processing account with one important difference; neuroeconomics provides an account of how value is represented and avoids running into the dark room problem by offering an account of why animals initiate motion and engage in exploration. More importantly, neuroeconomics offers an account of how animals discriminate between better and worse strategies of action. The outcome of the Platt and Glimcher experiments show that the magnitude of expected reward and the probability of reward are represented in the brain and influence decisions. In the following section, I expand on the idea that value is represented in the brain and look at how subjective value is represented in the brain.

---

<sup>26</sup> In Platt and Glimcher's experiment, the monkeys were only offered two options but other experiments such as Basso & Wurtz's (1997) task, monkeys were presented with eight targets and were able to select only one action.

## **10. How value is encoded in the brain**

In earlier sections, the importance of the value of an outcome was illustrated through a description of economic theories and through the neuroscientific experiments conducted by Platt and Glimcher. The economic theoretical framework had run into some problems because it inadequately explained certain behaviours illustrated by the Allais paradox and the Ellsberg paradox. Glimcher's solution to the problem that certain behaviour is poorly described by neoclassical economic theories was to provide an inter-theoretic reduction. The main objective of Glimcher's approach is to provide a theory of decision making that takes into account the economic theoretical framework and to relate this framework to a field that can provide empirical evidence; this field is neuroscience. Experiments done by Platt and Glimcher bring to light three important characteristics of value. First, value is encoded and registered in the brain. Glimcher proposes that the utility of an object or state can be described as the "...neural activation that encodes the desirability of an outcome in a continuous monotonic fashion." (Glimcher 2011: 134). Second, though utility and action are separable, the two concepts are causally related (Glimcher 2011: 195). The third follows from the second; the value of a reward plays a role in decision making. In neuroeconomic theory of decision making, value plays an important role in choice and action selection because the choices made by animals are affected by the associated value of the outcomes of the choices or actions. The value of an action is measured in relation to the expected gain or loss and the probability that the gain or loss will be realized.

An important idea to retain from the discussion on neuroeconomics is that events, objects and system states are all labeled with motivational values (Schultz 1998: 2). Events and system states obtain motivational significance by how they affect welfare, survival and reproduction and can be appetitive or



aversive; in other words, events or system states can be rewarding or punishing (Schultz 1998: 2). Given the knowledge that value is represented in the brain and that value plays a role in decision making, the question that remains is, how are these values represented in the brain? This question is notably absent from Clark's account of cognition and he acknowledges that "...there are neither goals nor reward signals..." in the predictive processing account (Clark 2013: 200). One approach is to keep a register of the motivational values of the available options and to update the registers when feedback is received from performing an action (Shea 2012: 5). The amount of attention paid to value and reward is a notable difference between the account supported by Clark and work by neuroeconomists, and is illustrated in the 1998 paper by Wolfram Schultz who provides an answer to the question, how are values represented in the brain? I will return to discussion on Clark in later sections and will first be focusing on the views on value that are offered by Schultz. The value of an outcome (i.e. reward or punishment) plays an important role in the behaviour of animals and influences the animal's life conditions. Schultz is aware of the importance of value in explaining behaviour and conducted a careful study of the representation of value to determine how value is represented. The experiments conducted by Schultz were grounded in the idea that the dopamine system played a role in encoding reward. For decades dopamine was thought to be merely associated with noradrenaline and adrenaline but it was discovered that dopamine was a neurotransmitter and that the dopamine system was involved in encoding reward; this discovery prompted the development and exploration of important ideas concerning neural valuation circuitry and the role of dopamine in reward processing and opened up a new dimension of study for those interested in dopamine; particularly important was studying the role of dopamine in learning and valuation. This study was undertaken by many scholars, including Hyman, Malenka and Nestler (2006), Arias-Carrión & Pöppel (2007), Waelti, et al. (2001) and Shea (2012) to name a few. Shea makes the

claim that reward prediction error drives learning (Shea 2012: 5). Reward prediction error is identical to prediction error (as described by Clark) in that it functions in exactly the same manner. Reward prediction error is calculated in each step of the process and is the result of the mismatch, or difference, between the predicted payoff and the actual payoff received (Shea 2012: 5). This will guide the focus of discussion in the following paragraphs.

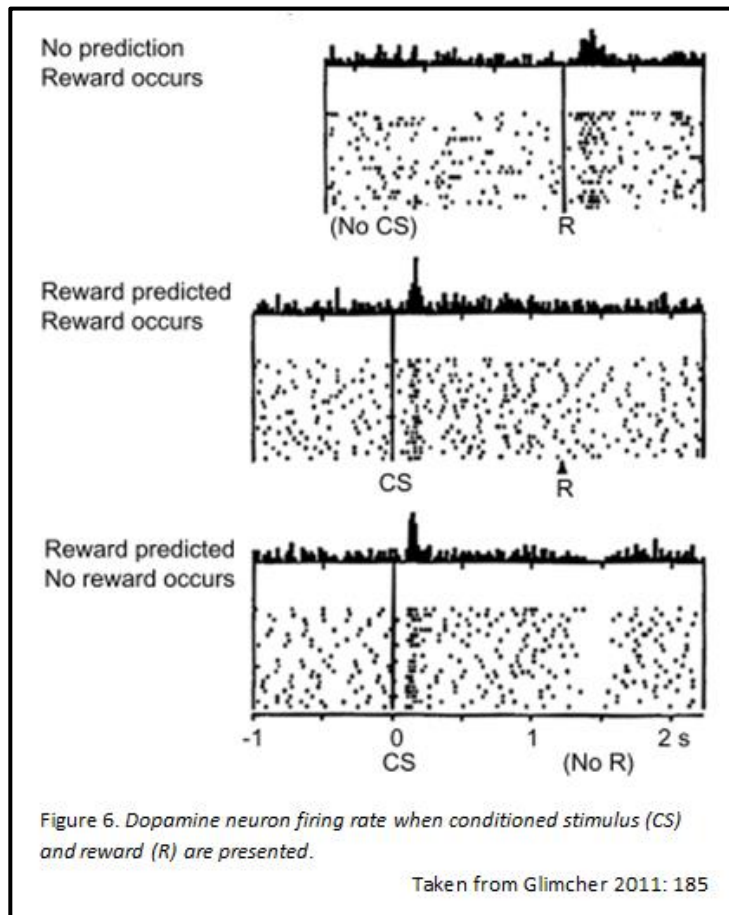
For decades, the idea that dopamine plays a role in processing rewards was supported and it was believed that the firing rate of dopamine neurons carried a utility-like signal (Glimcher 2011: 300). This view was problematic because it was unclear how exactly dopamine, and the neural circuitry it was involved in, functioned but what scientists did know was that humans and animals were willing to work to artificially activate dopamine neurons but they also came to know that natural rewards, such as food and sex, did not always activate dopamine neurons (Glimcher 2011: 300). In other words, it was believed that dopamine played a role in motivated behaviour but the specific relation became uncertain as more studies were conducted. The view that dopamine is the neural equivalent of reward lost support in the late 20<sup>th</sup> century and a new school of thought was founded. The new school of thought proposed the “wanting” hypothesis of dopamine of which the central statement was that dopamine encoded the experience of wanting a reward (Glimcher 2011: 301). This implied that dopamine neurons fired before a reward was obtained; this became the framework within which Schultz developed his experiments.

### **10.1. Dopamine neuron activation - an experiment**

In response to the question about the function of dopamine, Jacques Mirenowicz and Wolfram Schultz developed an experiment to study the activity of single

dopamine neurons. The experiment involved monkeys participating in a Pavlovian conditioning task (Glimcher 2011: 301). The purpose of the experiment was to explore the relation between dopamine neurons, reward and reward prediction. Dopaminergic neurons in certain brain areas, such as the ventral tegmental area, have a firing profile that corresponds to reward prediction error (Shea 2012: 6). In Mirenowicz and Schultz's experiment, thirsty monkeys were placed in front of a spout. At random intervals, a tone was produced; the tone functioned as a conditioned stimulus (CS). After a short delay, a drop of juice was dripped from the spout; this was the reward (R). The objective of the experiment was to study the relationship between spikes in firing rates of dopamine neurons during the presence of the conditioned stimulus and during the delivery of reward. Dopamine neurons in monkeys have a baseline firing rate of three to five spikes per second. Initially, the firing rate of dopamine neurons remained constant when the tone (CS) was produced and increased when the juice (R) was dripped from the spout (Glimcher 2011: 301).

After a few trials, the frequency of dopamine neuron activation increased at the time the tone was presented (CS) and returned to baseline when the juice (R) was delivered. This indicates that the firing rates of dopamine neurons decrease in response to the reward and increase in response to the conditioned stimulus. If the juice (R) was delivered without the experimenter producing a tone (CS), the dopamine neurons responded to the reward which indicated that the reward had not lost its ability to activate dopamine neurons. In the final experiment, the tone (CS) was produced but the juice (R) was not delivered. The finding was that dopamine neuron activation dropped below baseline at the time of expected reward in the trials when the tone was produced but the juice was not delivered (Glimcher 2011: 302). These findings are illustrated in the graphs below.



The experiments conducted by Mirenowicz and Schultz illustrate that reward predictability plays an important role in dopamine neuron activation. If the reward is predicted by the presence of a conditioned stimulus, dopamine neurons have a higher firing rate at the time of conditioned stimulus presentation than at the time of reward presentation. If the reward is not predicted, dopamine activation occurs at the time the reward is presented. Dopamine neurons report “...rewards according to the difference between the occurrence and the prediction of reward...” (Schultz 1998: 7). Predicted rewards can be innate or learned. We learn to associate certain rewards with stimuli when the system reduces reward prediction error. Reward prediction error occurs when the reward is not fully predicted by the conditioned stimulus (Schultz 1998: 11). The process can be described as follows: the system

processes an event, predictions are generated, and the difference between the prediction and actual event is computed (Schultz 1998: 12). The system makes use of prediction errors to optimize performance and predictions. This illustrates that the process of updating reward predictions is the same process as updating prediction error in Clark's predictive processing account. There is strong evidence that reward prediction error plays a role in learning and guides decision making in many species (Shea 2012: 17). From this follows the inference that evolution has set up living systems to maximize the delivery of reward to the agent (Shea 2012: 20).

### **11. Predictive processing and neuroeconomics**

The predictive processing account of cognition is offered as a general account of cognition that provides a solution to the architecture problem and makes the claim that the function of the brain is to reduce prediction error (Clark 2013: 192). Prediction error is the result of a mismatch between predicted and actual information and plays an important role in learning; in the neural domain, prediction error is the measure that reports the free energy in a biological system, the brain. The account offered by Clark supports the view that living organisms make use of perception and action to minimize free energy and to reduce prediction error (Clark 2013: 200). Implied in the predictive processing account of cognition is the view that the sole biological imperative of living organisms is to minimize free energy and reduce prediction error; this is achieved when animals actively sculpt and selective sample the environment (Clark 2013: 194). By making the strong claim that minimizing prediction error is the sole neural and biological imperative of living organisms and that living organisms act only as a response to surprise in the environment, the predictive processing account runs into a problem. The problem takes shape as a *reductio ad absurdum* objection, called the dark room problem, and can be described as

follows. Supporters of the predictive processing account of cognition are committed to the view that the sole biological imperative of all living organisms is to minimize prediction error. One strategy that enables living organisms to minimize prediction error successfully is to seek out a monotonous environment, such as a dark room, and stay there. It follows that, by being committed to the view that reducing prediction error is the sole imperative of living organisms, supporters of the predictive processing account also commit to the strategy proposed (seek out a monotonous environment to avoid surprise). This strategy is biologically nonsensical because animals need to feed, reproduce and actively participate in the environment. Such activity is unlikely to take place if animals only inhabit monotonous environments free from stimuli and exploration.

Clark acknowledges that the consequence that follows from the function of cognition is absurd and biologically nonsensical but says that we are not to worry because the *reductio ad absurdum* consequence does not follow. Clark suggests that the consequence does not follow because animals have expectations about challenge and exploration and, therefore, expect to act on the world and engage in challenge and exploration. Animals that have adapted to changing and challenging environments will find monotonous environments surprising and will seek activity. If a lizard, for example, found itself in a dark or sensory monotonous environment, it will be surprised by the lack of sensory activity and stimuli and will be prompted to seek information. This applies not only to perception but to action too. For example, if a wild cat minimizes free energy in the manner suggested by the free energy principle, and seeks out a dark cave - it will have no prediction errors because it will have nothing to predict - but it will eventually be surprised by its own inactivity. In other words, Clark seems to think that animals do not seek out monotonous environments because they have expectations about exploration but Clark is still committed to the view that the

function of cognition is to reduce prediction error. In other words, Clark's response to the dark room problem implies that animals participate in activity only because they expect to do so and he proposes that an organism moves through time and space in ways that minimize prediction error which is the sole function of cognition, perception and action. This, however, does not support what is known about the evolutionary and biological imperatives of animals.

It is known that animals act on the world because they have motivations to attain certain goal states. Animals get hungry and thirsty, and act on the world to satisfy those needs; animals also have an evolutionary imperative to reproduce and need to employ various strategies to accomplish this successfully. Clark's account lacks any discussion on the evolutionary motivations that initiate action and the strategies employed to attain certain states in the world. In section 2 of this dissertation, I discussed action and offered the efficiency problem as one dimension of the action selection problem. The environmental complexity hypothesis was explored as a solution to the efficiency problem and the account brought to light two elements that are important in behaviour and action selection. First, at any given time, an animal has several possible actions to choose from and some actions are better than others. Second, each possible action is accompanied by consequences (good or bad) and these consequences are learned and used in future behaviour. The predictive processing account offered by Clark (both before and after his response to the dark problem) does not offer any account that considers the two elements that are important in action selection. Clark does not acknowledge in his account that animals must recognize strategies as better or worse and that they must be able to update their goals given the value of an outcome (Huebner 2012). According to the predictive processing account, encoded expectations guide perception and action; in other words, behaviour is a result of the living organism's imperative

to minimize prediction error. What this means is that reward signals and cost functions are replaced by expectations (Clark 2013: 200). Clark fails to provide an account of how we navigate the world, given what we know about the outcomes of a selection, and how we discriminate between better and worse possible actions, given that there is value attached to the outcome of an action.

Clark is not concerned with value and motivation and therefore misses an opportunity to strengthen the predictive processing account of cognition. Various resources, such as rewards and the value of outcomes, are utilized by animals to maintain welfare, survival and reproduction but these resources are not explored by Clark. Value relates to action selection because reward prediction error carries the following information: "...perform the most recent action, or for negative values: do not perform the most recent action" (Shea 2012: 23). The predictive processing account is reduced to the dark room problem because an account of value and motivation is not provided. If an account of value and motivation was included in the predictive processing account, the account would have also been committed to a view in which animals had biological motivation to act (other than minimizing free energy). Furthermore, by providing an account of cognition that includes encoding the value of the outcome of actions means that the living organism has the ability to discriminate between better and worse possible actions. What is required, then, is an account of cognition that tells us something about value and how it is represented in the brain. In his 2010, *Nature Reviews Neuroscience* paper, Friston acknowledges that value is an important element in theories about the brain. He largely supports the notion put forth by expected utility theory that we select the option with the highest value.<sup>27</sup> Importantly, Friston's view differs from traditional economic theories because he proposes that "...value is

---

<sup>27</sup> Friston means that biological systems increase value using mechanisms such as temporal difference learning.



inversely proportional to surprise.” (Friston 2010: 8). In other words, Friston’s proposal is that in a biological system, value is increased when surprise is decreased. Friston’s account of value goes full circle as it leads back to free energy minimization as sole biological imperative and, therefore, leads us back to the dark room problem.

In light of this, I propose that a different theory of value, and how value is represented, be explored and look to neuroeconomics. Neuroeconomics provides important insight that can be useful in developing an account of cognition that pays attention to value. Neuroeconomics is useful as a theory of value because it combines the economic theoretical framework, which makes fairly good predictions about behaviour, with neuroscience. Grounding economic theories in a field such as neuroscience means that choice behaviour can be observed and predictions about behaviour can be made. The neuroeconomic account of cognition offers solutions to both the efficiency problem and the architecture problem because the account provides information about the value and outcome of actions and how that influences decision making and action selection; the account also offer important insight about the mechanism involved in action selection and decision making, and proposes that value is represented in the brain. Experiments conducted by Platt and Glimcher in 1999 also illustrate that the magnitude and probability of rewards directly correlate to the firing rate of neurons in the brain. This indicates that value is encoded in the brain and plays a role in action selection. Neuroeconomic models of cognition, therefore, differ from Clark’s account in an important way: the value of an outcome is central to neuroeconomic models of cognition but strangely neglected in Clark’s account. By providing an account of how value is represented and the role it plays in behaviour, neuroeconomic models of cognition avoid problems that Clark’s account runs into.

Calculating the value of action can influence behaviour in ways that encourage efficient behaviour by influencing perception and action. Predictions about the expected value of an outcome allow an organism to evaluate situations before they occur and enable the organism to select and prepare appropriate behavioural reactions and increase the likelihood of approaching or avoiding situations with motivational value (Schultz 1998: 2). Furthermore, living organisms act on the world to maintain homeostatic balance and to reproduce (Schultz 1998: 1). This requires registering the outcome of actions and for an animal to be able to discriminate between better and worse strategies of action. How is value represented in the brain? The experiment by Mirenowicz and Schultz discussed in earlier sections suggest that value is encoded by a reward prediction error which is processed by temporal difference learning, a method of prediction also employed by the predictive processing account. The activation of dopamine neurons signal expected reward, and report rewards relative to their prediction rather than reporting rewards unconditionally (Schultz 1998: 7). This similarity between neuroeconomic models of value and Clark's account of cognition strengthens the argument that the two views are compatible. By incorporating reward prediction error into the predictive processing account, Clark would avoid running into the dark room problem because an account of motivated behaviour and discrimination between tasks would be provided. Neuroeconomics a suitable account of cognition because it provides solutions to both the efficiency and architecture problem of action selection.

## References

- Allais, M. 1953. Le comportement de l'homme rationnel devant le risque: critique des postulats et axiomes de l'école américaine. *Econometrica: Journal of the Econometric Society*, **21**(4): 503-546.
- Arias-Carrión, Ó. & Pöppel, E. 2007. Dopamine, learning, and reward-seeking behavior. *Acta neurobiologiae experimentalis*, **67**(4): 481-488.
- Avital, E. & Jablonka, E. 2000. *Animal traditions: Behavioural inheritance in evolution*. United Kingdom: Cambridge University Press.
- Barron, A. B. Sjøvik, E. & Cornish, J. L. 2010. The roles of dopamine and related compounds in reward-seeking behavior across animal phyla. *Frontiers in behavioral neuroscience*, **4**(163).
- Basso, M. A., & Wurtz, R. H. 1997. Modulation of neuronal activity by target uncertainty. *Nature*, **389**(6646), 66-69.
- Bentham, J. & Bowring, J. 1843. *The works of Jeremy Bentham* (Vol. 7). Edinburgh: William Tait.
- Bernheim, D. & Rangel, A. 2008. Choice-Theoretic Foundations for Behavioral Welfare Economics. In Caplin, A & Schotter, A. (eds.). *The Foundations of Positive and Normative Economics: A Handbook*. Oxford: Oxford University Press.
- Bernoulli, D. 1954. Exposition of a New Theory on the Measurement of Risk. *Econometrica*, **22**(1): 23-36.
- Biederman, I. 1987. Recognition-by-components: a theory of human image understanding. *Psychological review*, **94**(2), 115.
- Brown, H., Friston, K. & Bestmann, S. 2011. Active inference, attention, and motor preparation. *Frontiers in psychology*, **2**(218).

- Chapman, P. 1982. The Origins of Troglodites. *Proceedings of the University of Bristol Spelaeological Society* **16**(2): 133–141.
- Christensen, W. 2010. The Decoupled Representation Theory of the Evolution of Cognition—A Critical Assessment. *The British Journal for the Philosophy of Science*, **61**(2): 361-405.
- Clark, A. 2012. Dreaming the Whole Cat: Generative Models, Predictive Processing, and the Enactivist Conception of Perceptual Experience. *Mind*, **121**(483): 753-771.
- Clark, A. 2013. Whatever next? Predictive brains, situated agents, and the future of cognitive science. *Behavioral and Brain Sciences*, **36**(3), 181-204.
- Clark, A. 2014. Perceiving as Predicting. In Mohan, M., Biggs, S., & Stokes, D. (eds.). *Perception and Its Modalities*. New York: Oxford University Press.
- Dawkins, R. 1989. *The selfish gene*. New York: Oxford University Press.
- Dickinson, A. & Balleine, B. 2002. The role of learning in the operation of motivational systems. In C. R. Gallistel. (ed). *Steven's handbook of experimental psychology: Learning, motivation and emotion*. Volume 3. 3rd edition. New York: John Wiley & Sons.
- Edwards, W. 1954. The theory of decision making. *Psychological bulletin*, **51**(4): 380-417.
- Eliasmith, C. 2007. How to build a brain: From function to implementation. *Synthese*, **159**(3): 373-388.
- Friston, K. 2002. Functional integration and inference in the brain. *Progress in neurobiology*, **68**(2): 113-143.
- Friston, K. 2003. Learning and inference in the brain. *Neural Networks*, **16**(9): 1325-1352.

- Friston, K. 2010. The free-energy principle: a unified brain theory? *Nature Reviews Neuroscience*, **11**(2): 127-138.
- Friston, K. 2011. Embodied inference: Or I think therefore I am, if I am what I think. In Tschacher, W. & Bergomi, C. (eds.). *The implications of embodiment: Cognition and communication*. Exeter: Imprint Academic.
- Friston, K. J., Daunizeau, J. & Kiebel, S. J. 2009. Reinforcement learning or active inference?. *PloS one*, **4**(7): e6421.
- Friston, K. & Stephan, K. E. 2007. Free-energy and the brain. *Synthese*, **159**(3): 417-458.
- Friston, K., Thornton, C. & Clark, A. 2012. Free-energy minimization and the dark-room problem. *Frontiers in psychology*, **3**(130).
- Gallistel, C. R. 1990. *The organization of learning*. Cambridge: MIT Press.
- Gauthier, D. P. 1986. *Morals by agreement*. New York: Oxford University Press.
- Glimcher, P. W. 2011. *Foundations of neuroeconomic analysis*. New York: Oxford University Press.
- Glimcher, P. W., Dorris, M. C. & Bayer, H. M. 2005. Physiological utility theory and the neuroeconomics of choice. *Games and economic behavior*, **52**(2): 213-256.
- Glimcher, P. W. & Sparks, D. L. 1992. Movement selection in advance of action in the superior colliculus. *Nature*, **355**(6360): 542-545.
- Godfrey-Smith, P. 2002. Environmental complexity and the evolution of cognition. In Sternberg, R. & Kaufman, J. (eds.). *The evolution of intelligence*. Mahwah: Laurence Erlbaum Associates.

- Griffin, D. R. (ed.). 2001. *Animal minds: Beyond cognition to consciousness*. Chicago: The University of Chicago Press.
- Grush, R. 2004. The emulation theory of representation: motor control, imagery, and perception. *Behavioral and brain sciences*, **27**(3): 377-396.
- Hohwy, J. 2010. The hypothesis testing brain: Some philosophical applications. In Christensen W., Schier, E. & Sutton, J. (eds.). *ASCS09: Proceedings of the 9<sup>th</sup> Conference of the Australasian Society for Cognitive Science*. Sydney: Macquarie Centre for Cognitive Science.
- Hohwy, J., Roepstorff, A. & Friston, K. 2008. Predictive coding explains binocular rivalry: an epistemological review. *Cognition*, **108**(3): 687-701.
- Houston, A. I., McNamara, J. M. & Steer, M. D. 2007. Do we expect natural selection to produce rational behaviour? *Philosophical Transactions of the Royal Society B: Biological Sciences*, **362**(1485): 1531-1543.
- Houthakker, H. S. 1950. Revealed preference and the utility function. *Economica*, **17**(66): 159-174.
- Hubel, D. H. & Wiesel, T. N. 1965. Receptive fields and functional architecture in two nonstriate visual areas (18 and 19) of the cat. *Journal of neurophysiology*, **28**(2): 229-289.
- Huebner, B. 2012. Surprisal and valuation in the predictive brain. *Frontiers in psychology*, **3**(415).
- Hyman, S. E., Malenka, R. C. & Nestler, E. J. 2006. Neural mechanisms of addiction: the role of reward-related learning and memory. *Annual Review of Neuroscience*, **29**: 565-598.
- Kable, J. W. & Glimcher, P. W. 2007. The neural correlates of subjective value during intertemporal choice. *Nature neuroscience*, **10**(12): 1625-1633.

- Korenko, S., Schmidt, S., Schwarz, M., Gibson, G. & Pekar, S. 2013. Hymenopteran parasitoids of the ant-eating spider *Zodarion styliferum* (Simon) (Araneae, Zodariidae). *ZooKeys*, **262**: 1-15
- Knill, D. C. & Pouget, A. 2004. The Bayesian brain: the role of uncertainty in neural coding and computation. *TRENDS in Neurosciences*, **27**(12): 712-719.
- Loewenstein, G. 2000. Emotions in economic theory and economic behavior. *American economic review*, **90**(2): 426-432.
- Lorenz, E. N. 1982. Atmospheric predictability experiments with a large numerical model. *Tellus*, **34**(6): 505-513.
- Maes, P. 1990. Situated agents can have goals. *Robotics and autonomous systems*, **6**(1): 49-70.
- Mas-Colell, A. 1982. Revealed preference after Samuelson. In Feiwel, G. R. (ed.) *Samuelson and neoclassical economics*. Netherlands: Springer.
- McClure, S. M., Laibson, D. I., Loewenstein, G. & Cohen, J. D. 2004. Separate neural systems value immediate and delayed monetary rewards. *Science*, **306**(5695): 503-507.
- Mirenowicz, J. & Schultz, W. 1994. Importance of unpredictability for reward responses in primate dopamine neurons. *Journal of neurophysiology*, **72**(2): 1024-1027.
- Mueller-Roeber, B., & Dreyer, I. 2007. Ion homeostasis: plants feel better with proper control. *EMBO reports*, **8**(8): 735-736.
- Mumford, D. 1992. On the computational architecture of the neocortex. *Biological cybernetics*, **66**(3): 241-251.
- Pascal, B. 2003. *Pensees*. Translated by Trotter, W. F. New York: E.P. Dutton & Co.

- Platt, M. L. & Glimcher, P. W. 1999. Neural correlates of decision variables in parietal cortex. *Nature*, **400**(6741): 233-238.
- Rao, R. P. & Ballard, D. H. 1999. Predictive coding in the visual cortex: a functional interpretation of some extra-classical receptive-field effects. *Nature neuroscience*, **2**(1): 79-87.
- Rauss, K., Schwartz, S. & Pourtois, G. 2011. Top-down effects on early visual processing in humans: a predictive coding framework. *Neuroscience & Biobehavioral Reviews*, **35**(5): 1237-1253.
- Rejcek, P. 2010. *Researchers follow Adélie penguin winter migration for the first time*. [online] 13 August 2010. Available from: [www.antarcticsun.usap.gov](http://www.antarcticsun.usap.gov) <http://antarcticsun.usap.gov/science/contenthandler.cfm?id=2230> [Accessed: 21 October 2014]
- Riesenhuber, M. & Poggio, T. 2000. Models of object recognition. *Nature neuroscience*, **3**: 1199-1204.
- Samuelson, P. A. 1938. A note on the pure theory of consumer's behaviour. *Economica*, **51**(17): 61-71.
- Sanfey, A. G., Loewenstein, G., McClure, S. M. & Cohen, J. D. 2006. Neuroeconomics: cross-currents in research on decision-making. *Trends in cognitive sciences*, **10**(3): 108-116.
- Schall, J. D. & Hanes, D. P. 1993. Neural basis of saccade target selection in frontal eye field during visual search. *Nature*, **366**(6454): 467-469.
- Schultz, W. 1998. Predictive reward signal of dopamine neurons. *Journal of neurophysiology*, **80**(1): 1-27.



- Seth, A. K. 2007. The ecology of action selection: Insights from artificial life. *Philosophical Transactions of the Royal Society B: Biological Sciences*, **362**(1485): 1545-1558.
- Shea, N. 2012. Reward Prediction Error Signals are Meta-Representational. *Noûs*, **48**(2): 314-341.
- Sherrington, C. S. 1906. *The Integrative Action of the Nervous System*. Connecticut: Yale University Press.
- Spurrett, D. 2012. What is to be Done? Why Reward is Difficult to Do Without. *Frontiers in psychology*, **3**: 412.
- Sterelny, K. 2003. *Thought in a hostile world: The evolution of human cognition*. Oxford: Blackwell Publishing Ltd.
- Tinbergen, N. 1951. *The study of instinct*. Michigan: Clarendon Press
- Tong Lab. *Rivalry and awareness*. [online] Available from: [http://www.psy.vanderbilt.edu/tonglab/web/research/rivalry\\_awareness.html](http://www.psy.vanderbilt.edu/tonglab/web/research/rivalry_awareness.html) [Accessed: 21 September 2014].
- Varian, H. R. 2006. Revealed preference. In Szenberg, M., Ramrattan, L. B. & Gottesman, A. A. *Samuelsonian economics and the twenty-first century*. Oxford: Oxford University Press.
- Von Neumann, J. & Morgenstern, O. 1944. *Theory of Games and Economic Behavior*. Princeton: Princeton University Press.
- Waelti, P., Dickinson, A. & Schultz, W. 2001. Dopamine responses comply with basic assumptions of formal learning theory. *Nature*, **412**(6842): 43-48.
- Wooldrige, D. E. & Knight, D. E. 1963. *The machinery of the brain*. New York: McGraw-Hill.