

THE FUTURE EVOLUTION OF CONSCIOUSNESS

John STEWART¹

The Evolution, Complexity and Cognition Group, Free University of Brussels.

ABSTRACT. What potential exists for improvements in the functioning of consciousness? The paper addresses this issue using global workspace theory. According to this model, the prime function of consciousness is to develop novel adaptive responses. Consciousness does this by putting together new combinations of knowledge, skills and other disparate resources that are recruited from throughout the brain. The paper's search for potential improvements in the functioning of consciousness draws on studies of the shift during human development from the use of implicit knowledge to the use of explicit (declarative) knowledge. These studies show that the ability of consciousness to adapt a particular domain improves significantly as the transition to the use of declarative knowledge occurs in that domain. However, this potential for consciousness to enhance adaptability has not yet been realised to any extent in relation to consciousness itself. The paper assesses the potential for adaptability to be improved by the conscious adaptation of key processes that constitute consciousness. A number of sources (including the practices of religious and contemplative traditions) are drawn on to investigate how this potential might be realised.

1. INTRODUCTION

An improved capacity to develop novel adaptive responses has often been given as the reason why evolution favoured the emergence of consciousness. Consciousness is increasingly seen as a process that confers evolutionary advantage by enhancing the ability of an organism to discover new and better behavioural adaptations (Baars 1988; Dennett 1991; Metzinger 2003 and DeHaene and Naccache 2001).

Global Workplace (GW) theory attributes this capacity of consciousness to its ability to assemble novel combinations of knowledge, skills and other resources for the development of new adaptive responses (Baars 1983, 1988 and 1997). This enables consciousness to, for example, recruit the resources needed to construct composite mental representations of alternative responses and their consequences, enabling the most adaptive response to be identified.

This paper explores the extent to which the adaptability conferred by consciousness can be enhanced in humans. In particular, it seeks to identify the potential for changes in the functioning of consciousness to improve its ability to discover better behavioural adaptations.

¹ Contact addresses: JES999@tpg.com.au ; and PO Box 8076 Rivett, ACT, 2611, Australia

Addressing this question has been greatly assisted by the recent development of information processing theories of the functioning of consciousness. An understanding of consciousness from an information processing perspective enables us to assess its potential for further improvement. We can judge how well consciousness performs its functions, and whether changes in the processes that constitute consciousness could overcome any limitations and enhance its ability to adapt behaviour effectively.

This paper uses the information processing framework embodied in GW theory to assess the potential for improvement in consciousness. Of the competing information processing accounts of consciousness, it currently attracts the widest support (Baars 2002; DeHaene and Naccache 2001; Kanwisher 2001; and Dennett 2001).

The paper focuses on those improvements that can emerge through the processes of cultural evolution, rather than through genetic change. It therefore considers only changes that can be learnt and can be transmitted culturally.

We begin in Section 2 by outlining and developing the main features of Baars' GW theory that are relevant to our task. This analysis of GW theory is drawn on in Section 3 to identify how the functioning of conscious processes might be improved to enhance the adaptability and evolvability of humans. The search for potential improvements is aided by developmental research that identifies how adaptability is significantly enhanced in other domains when declarative knowledge is used to assist adaptation. In section 4 we use GW theory and other sources to begin to identify practices and experiences that could enable humans to acquire skills and capacities that would realise the potential to improve the functioning of consciousness.

2. GLOBAL WORKSPACE THEORY

The Global Workspace architecture

According to Baars (1983, 1988, and 1997), the prime functional role of consciousness is to contribute to the development of novel adaptive responses by bringing together new combinations of knowledge, skills, and other resources.

Such a function is of critical importance when an organism encounters new circumstances for which it has no existing adaptive behavioural routine. If the organism has previous experience of particular components of the new situation, its conscious processes can combine the disparate experiences to construct composite mental representations of the novel circumstances and how they will respond to alternative actions. According to GW theory, the recruitment function of consciousness plays a central role in producing these novel representations—it recruits the particular combinations of knowledge and other resources needed to represent actions and situations that have not been experienced previously as a whole.

Baars suggests that the novel combinations of resources that are relevant to solving a particular adaptive problem are recruited through a GW architecture. Whatever occupies the GW at a particular time is broadcast globally throughout the brain. The GW might, for example, be occupied by an adaptive problem, a step in solving a problem, or a perception or issue that is relevant to a current goal. Specialist processors throughout the brain compete to respond to the broadcast as individuals, or in coalitions. The output of the winner of this competition then occupies the GW, and is in turn broadcast. The tasks carried out by the specialist processors or coalitions of processors can be very complex, but are performed outside conscious awareness.

Only outputs of processors that become contents of the GW are conscious. The stream of consciousness is the sequence of contents of the GW.

The competitions between unconscious processors are regulated by goal hierarchies (each goal is a representation of a desired state of affairs, and as the hierarchy is descended, goals become more specific and immediate). As part of the context of the GW, goal hierarchies condition what it takes to win competitions, including those that determine where attention is directed. But the goals and goal hierarchies often do not occupy the GW itself, and on such occasions are not conscious.

By regulating what it is that occupies the GW and therefore what is conscious, goal hierarchies (whether conscious or not) control how consciousness is deployed and what adaptive tasks it is given. Goal states themselves are determined in large part by the hedonic system which endows possible behavioural outcomes with a positive or negative hedonic tone (hedonic tones are generally in the form of valences associated with desires, motivations and emotions (Johnston 1999)). Ultimately, the effect of these processes is to ensure that consciousness is deployed to find creative and adaptive means to serve the ends determined by the hedonic system.

A key feature of such a GW architecture is that it can put together resources that are relevant to an adaptive problem despite the absence of any previous linkage between the resources and the problem, and without any centralised process that 'knows' the location of the resources in the brain, or that 'knows' their relevance to the adaptive problem. And despite this absence of prior knowledge about which resources are relevant, the GW system is capable of rapidly identifying those that are significant through the use of its massively parallel broadcast and competition process (see Shanahan and Baars 2005).

The characterisation of consciousness as a GW process accords well with subjective experience. It is evident that being conscious of an event goes hand in hand with the availability of the event to other resources—if we are conscious of something, we are able to give it attention, think about it, introspect in relation to it, talk about it, feel in relation to it, mull over it, and act on it. When we are conscious of something, the experience is available for other resources to operate on. If we are not conscious of something, the experience is not available to other resources. When we perform some complex automatic behaviour such as walking, the elements of the behaviour are not available to other resources, unless or until something goes wrong that attracts our attention, or we voluntarily direct attention to it.

Furthermore, a strong case can be made out that the GW model can broadly account for the contents of conscious experience—our stream of conscious experience appears to equate closely to the sequence of GW contents that would be expected to arise from the operation of the GW system (Baars 1997; Dennett 2001; and Clark 2005).

Mounting empirical evidence also supports GW theory (for recent reviews see Baars 2002 and 2003).

The serial operation and limited capacity of consciousness

A strength of GW theory is its ability to explain the features of consciousness that restrict its role in producing behaviour. Conscious processing is serial and limited in capacity (Miller 1956; Posner and Snyder 1975). This manifests as an inability to do more than one demanding voluntary task at a time. For example, two stories cannot be followed simultaneously, and an individual can think about how to solve a mathematical problem or how to handle a social dilemma, but not both

at once. The learning of a new complex skill is a slow and error-ridden process, and component skills often have to be mastered separately.

These limitations of conscious action contrast with automatic behaviour. Once a complex skill has been assembled by consciousness, practiced and learnt, it can be performed accurately and quickly with little conscious attention (Anderson 1980; Bargh 1997). However, if an attempt is made to again use thought and analysis to consciously guide the performance of the complex skill, performance deteriorates due to the inability to deal consciously with more than one element of the skill at a time.

The GW model explains these distinctive features of consciousness on the basis that only one internally consistent message at a time can be broadcast throughout the brain and dealt with by the distributed processors. The limited capacity and serial processing of consciousness are a direct consequence of the way in which the GW operates—the GW system will obviously not work effectively if there is ambiguity within a broadcast, or if more than one message is broadcast at any time. Although the GW architecture searches and assesses specialist processors in parallel, the GW and broadcast process is a processing bottleneck.

Complex automatic behaviours are obviously not limited in this way. They owe their speed and accuracy to parallel processing. Outside consciousness, the brain is massively parallel, and many things are being processed at the one time.

It is worth noting here that serial operation and limited capacity are a more significant impediment for tasks that require multiple broadcasts and recruitments through the GW (e.g. where a complex new mental model is being constructed, guided by structured sequences of thought). These limitations are not as serious for tasks where few broadcasts are needed and sequences of thought are unnecessary. For example, if a diversity of specialist processors are already available for adaptive challenges in a particular domain, consciousness can organise and assemble a suitable coalition of processors quickly with few broadcasts and little loading. Where such a coalition is responding to a situation automatically, consciousness can rapidly adjust the response if circumstances change by recruiting new processors to the coalition, again with little call on the limited capacity of the GW (e.g. see Willingham et al. 1989). Nevertheless, the undertaking of tasks that require little conscious involvement, such as the conscious monitoring and adjustment of automatic processes, will still be impeded if consciousness is fully loaded by other tasks.

In general, restricted capacity and serial operation seriously limit the proportion of information processing that can be undertaken by conscious processes. Only a very small amount of incoming sensory information can be processed and responded to consciously (Norretranders 1998). The overwhelming majority of information must be processed and acted upon by massively parallel unconscious processes. Humans can therefore consciously decide or deliberate about only a very small proportion of their behaviour—most must be implemented automatically.

Consciousness is a very limited resource that ideally would be reserved for where it is superior—for developing new adaptive responses to novel circumstances. When new responses have been assembled, learnt and practiced, they can be implemented far more effectively through processes that are automatic and parallel. Throughout an individual's life, an increasing proportion of behaviour is performed automatically with little input from consciousness. Consciousness operates in the main at the growing tip of behaviour where new responses are created. It plays a very narrow role in the determination of the full range of behaviour that is implemented at any moment (Velmans 1991; Bargh 1997).

The view that the role of consciousness in producing behaviour is severely limited does not accord well with subjective experience. Individuals are not generally aware of the extent of their automatic, non-conscious behaviours—their awareness is focused on and fully occupied by the conscious ones. And it does not take much to occupy consciousness fully—it can be fully loaded by an emotional response or a sequence of thought or imagining. Once loaded, there is nothing else in awareness. Due to limited capacity, individuals are necessarily unaware of nearly all the sensory information that impinges on their senses, and that when this information is broadly consistent with past experience, it is being acted upon by automatic and habitual processes that operate largely outside consciousness. Few realise that the breadth of their awareness of their environment and of their own behaviour is as limited as if they were looking through a straw. However, mounting empirical evidence is confirming this general picture of the limited role of consciousness in producing behaviour (Velmans 1991; Bargh and Chartrand 1999; and Wilson 2002). Furthermore, as we will discuss in detail below, this view of the functioning of behaviour can be confirmed by a ‘mindful’ consciousness that has been trained to be unloaded and resistant to distraction by immediate urges and desires.

3. THE POTENTIAL TO ENHANCE THE FUNCTIONING OF CONSCIOUSNESS

Now that we have sketched the key processes that comprise consciousness, we can turn to investigating whether these processes function optimally. To what extent could the way in which consciousness operates be changed to enhance the adaptability of humans? Could the significant contribution made by conscious processes to our adaptability be improved further by modifying the way consciousness functions?

We will proceed by examining a major developmental transition in human adaptability, and show that the potential for the transition to increase adaptability has not yet been exploited for the processes that constitute consciousness itself. Karmiloff-Smith (1986 and 1992) and a number of others such as Mandler (1992) and Keil (1989) have referred in differing terms to this transition which I will designate the declarative transition.

Procedural and declarative knowledge

In describing this transition it is useful to consider the knowledge that underpins adaptive behaviours in two categories: procedural and declarative (Anderson 1983). Procedural knowledge is the implicit knowledge that is embodied in skills such as riding a bicycle. The details of such skills are not accessible to consciousness, and are not expressible in linguistic form. We are unable to think about the details of procedural skills. Procedural knowledge is largely specific to the particular skill it underpins, and cannot be used generically for other adaptive purposes (Karmiloff-Smith 1992; Sun 1997).

From the perspective of GW theory, the details of procedural knowledge do not enter the GW and therefore do not become contents of consciousness—procedural skills are behaviours that are implemented automatically and unconsciously through parallel processes. If a skill is underpinned only by procedural knowledge, we have no conscious knowledge of the details of the skill.

In contrast, declarative knowledge is explicit, generic, and easily accessible to consciousness. Declarative knowledge is the knowledge an individual is aware of when consciously modelling alternative behaviours and their impacts on the environment. It is therefore the knowledge associated with thought, deliberation and theorising (Karmiloff-Smith 1992; Sun 1997).

From the perspective of GW theory, declarative knowledge manifests as representations that enter the GW and recruit other explicit representations that are relevant to the modelling of behaviour and its impacts. We are aware of declarative knowledge when it occupies the GW and is experienced as, for example, perceptions, thoughts and images.

These two types of knowledge are acquired differently (Anderson 1983; Keil 1989). Procedural skills in some domains are more or less innately specified (Karmiloff-Smith 1992). Procedural skills can also be acquired by conditioning and other forms of learning, but not insofar as they involve analysis or conscious modelling. As such, procedural knowledge is generally acquired implicitly. However, this does not mean that consciousness is unnecessary for the acquisition of procedural knowledge—as Baars (1988) establishes in detail, and as would be expected from the GW model, the recruitment of resources and therefore conscious awareness is essential for all learning.

In contrast, declarative knowledge is never innate, and its acquisition is generally accompanied by explicit inferences supported by conceptual understanding and analysis.

The declarative transition

In the developmental transition identified by Karmiloff-Smith and others, the adaptive behaviours and skills first acquired by an infant result from the acquisition of procedural knowledge. At different ages for various domains, the child then develops declarative knowledge about the behaviours—the individual gradually develops theories and models of the behaviour and its impacts, utilising schemas, concepts and other representations that are accessible to consciousness. Karmiloff-Smith terms this process in which implicit information eventually becomes explicit knowledge ‘representational redescription’.

The declarative transition produces a major improvement in adaptability in any domain in which it occurs (Karmiloff Smith 1986 and 1992). Procedural skills and behaviours can be very effective in the circumstances in which they are learnt (although they may represent a local optimum, rather than a solution that is universally best). But they are not readily adapted to different circumstances—procedural knowledge is not usable beyond the specific purposes for which it was originally acquired. Procedural skills and behaviours can be adapted only through further costly trial and error—the implicit knowledge that underpins them is not accessible to consciousness, and cannot be readily recombined with other knowledge to adapt to changed circumstances. As put by Pylysyshyn (1980), procedural skills are cognitively impenetrable.

For example, consider an individual who has effective procedural skills for fixing a range of mechanical problems in a particular type of engine, but has no explicit understanding of how the engine works or why the procedures fix the problems. Without an explicit theory of how engines function, the individual would have difficulty in fixing a problem in a different type of engine, and would have to be shown new procedures or develop them by trial and error.

In contrast, behaviours that are underpinned by declarative knowledge can be readily adapted to changed circumstances (Karmiloff Smith 1992). Declarative knowledge is in the form of representations of objects, properties, events, actions and relations that are combined to form a model or theory of the behaviour and its impacts. When circumstances change, these representations can be recombined with other representations to model the changed circumstances and the impact of alternative behaviours. Explicit representations and models developed for other behaviours or even from other domains can be incorporated where this is useful. The feasibility of combining knowledge from different domains increases significantly as common representational approaches are used across domains e.g. as more areas of knowledge are subsumed by science.

The use of declarative modelling can also enable the prediction of the future impacts of behaviour, allowing actions to anticipate future events. In contrast to procedural knowledge, declarative knowledge is usable beyond the special-purpose goals for which it was originally acquired. As a result, declarative modelling enables a significant reduction in the need for costly trial and error when adapting to novel circumstances and future possibilities.

Karmiloff-Smith (1992) traces the declarative transition in a number of domains, including linguistic skills, understanding of the physical world, the ability to understand the psychology of others, and mathematical skills. She shows that the transition unfolds at different times and rates across domains. Karmiloff-Smith also notes that once particular skills or behaviour have been revised and improved using declarative knowledge, they again become automatic and unconscious through a process of proceduralization. Once skills have been revised using the slow and serial processing of consciousness, they can be implemented more efficiently through parallel processing, and the limited resource of consciousness can be utilised more productively elsewhere. However, the 'post-declarative' proceduralized skills and behaviours differ significantly from those that have not undergone the declarative transition. The possession of declarative models of the skills means they can be readily modified and adapted declaratively if circumstances change significantly, or if new declarative knowledge is acquired. This also enhances the capacity to intervene during the implementation of a skill to adapt its use (Willingham et al. 1989; Karmiloff-Smith 1992).

In any particular domain in which a declarative transition unfolds, the serial process of declarative modelling progressively build a range of new resources and other expert processors, including cognitive skills. Once these processors have been built and proceduralized, they perform their specialist functions without loading consciousness—their outputs alone enter consciousness, without the declarative knowledge that went into their construction. The outputs are known intuitively (i.e. they are not experienced as the result of sequences of thought), and complex situations are understood at a glance (Reber 1989). As noted by Dreyfus and Dreyfus (1987), a person who achieves behavioural mastery in a particular field is able to solve difficult problems just by giving them attention—consciousness recruits the solutions directly from the relevant specialist processors.

Evolutionary declarative transitions

Although declarative transitions have been studied mainly in relation to individual development, they are also likely to have been significant in human evolution. Early in the evolution of life, adaptive behaviours were underpinned by procedural knowledge that was specified innately and discovered by genetic trial and error operating across the generations. Later, the processes of classical and instrumental conditioning enabled procedural knowledge to be acquired by trial and error processes operating within individuals (aided, particularly in humans, by imitative learning). It is only in humans that declarative transitions have occurred on any scale, and that declarative knowledge is accumulated and transmitted across the generations through the processes of cultural evolution. This has produced massive improvements in adaptability in the domains where transitions have occurred. Behaviours that were stereotyped and inflexible have become differentiated, variable and context-sensitive. Compare, for example, the diversity and flexibility of roles within a modern industrialised society with those in an early tribal society.

As for declarative transitions during development, evolutionary transitions would also have emerged at different times and at different rates across domains. A number of factors would have contributed to this, including that early declarative modelling would have been capable of exploring only domains with simple interactions with the environment. Furthermore, the unfolding of the transition is intrinsically a step-wise process in some circumstances. This is the

case where the development of declarative models of a domain opens up the possibility of a higher ('meta') level of declarative modelling. The meta-level comprises declarative models about the declarative models of the original domain, enabling the original models to be regulated and optimised declaratively rather than procedurally.

A clear example is provided by the evolution of thinking skills. As with all skills, when thinking first arose, the processes that shaped and structured sequences of thought would have been adapted procedurally. Although the *content* of thought was declarative, the *skills that regulated* the pattern of thought were procedural. Individuals would discover which particular structures and patterns of thought were useful in a particular context by what worked best in practice. They would not have declarative knowledge that would enable them to consciously model alternative thinking strategies and their effects. Thinking skills were learnt procedurally, and there were no theories of thought, or thought about thinking.

The development of declarative knowledge *about* thinking strategies enabled existing thinking skills to be improved and to be adapted to meet new requirements. In particular, it eventually enabled an explicit understanding of what constituted rational and scientific thought, and where and why it was superior. The declarative transition that enabled thinking about thinking was a major transition in human evolution that occurred on a wide scale only within recorded history (see Turchin (1977) and Heylighen (1991) who examine this shift within the framework of metasystem transition theory). This transition remains an important milestone in the development of individuals, and broadly equates to the achievement of what Piaget referred to as the formal operational stage (Flavell 1985). However, many adults still do not reach this level (Kuhn et al. 1977).

The unfolding of declarative transitions in human evolution is also likely to mirror the overall trend that has been identified in human development. Developmental studies show that throughout the lives of individuals, a progressively wider range of behaviours are adapted with the assistance of declarative modelling. Importantly, this trend also applies to the internal processes that adapt behaviour—cognitive skills are increasingly adapted using declarative modelling. As Kegan (1982) puts it in relation to development, what was subject progressively becomes an object of consciousness. Humans have acquired the capacity to think and talk about more and more, including about their cognitive processes. As human evolution has unfolded, the great potential of declarative modelling to improve adaptation has been realised over a progressively wider range of domains, and this trend can be expected to continue.

We now turn to identifying some key areas in which the declarative transitions are yet to occur to any extent, and where substantial potentials therefore exist for further improvements in adaptability. We will focus on the processes that constitute consciousness. Clearly, humans as yet have limited declarative knowledge about the central processes that produce consciousness, and little capacity to model and adapt these processes with the assistance of declarative knowledge.

Drawing on the GW model, we will now examine two sets of processes that are important in constituting consciousness, and look at the potential for declarative transitions to increase their adaptability and the adaptability of consciousness.

Processes that establish goals

As discussed previously, the goals that are served by behaviour are determined by the patterns of positive and negative reinforcement that constitute the hedonic system, including positive and negative affect associated with emotions. Hedonic tone generally attaches to all perceptions (including to internal representations of possible events), denoting their salience and relevance

(Zajonc 1980). The hedonic system determines what it is that we desire, like and dislike, need, want and are motivated to do.

Although the hedonic system is the central determinant of what individuals do from moment to moment in their lives, it is not adapted to any extent with declarative knowledge. Individuals generally do not choose deliberately what it is that they like or dislike, what they desire, or what they are motivated to do. They largely take these as given. They cannot change at will the impact of their desires, motivations and emotions on their behaviour, even when they see that the influence is maladaptive.

Instead, desires and motivations have been determined procedurally. Initially they were established and tuned by natural selection so that they were effective proxies for evolutionary success. Selection favoured desires and motivations that tended to produce behaviour that was advantageous in evolutionary terms.

These elements of the hedonic system can also be adapted somewhat during the development of the individual, but by classical conditioning rather than by processes that use declarative knowledge. Classical conditioning will establish an association between a hedonic tone that is innately elicited by particular stimuli and other stimuli that regularly co-occur with the innate stimuli.

Although humans do not yet use declarative modelling to establish and adapt their ultimate goals to any extent, increasingly individuals are developing some capacity to resolve conflicts between different desires and motivations declaratively. This is particularly the case when longer term desires and motivations conflict with more immediate impulses, such as those associated with strong emotions. Developmental studies suggest that early in life the demands of socialisation ensure that individuals acquire procedural skills that enable them to defer acting on immediate impulses where this conflicts with longer-term goals (Luria 1973). Increasingly this procedural knowledge is now being supplemented by declarative knowledge. The recent expansion of research into emotional regulation, the proliferation of self help literature and the emergence of education programs that teach emotional intelligence all appear to be part of a collective declarative transition in this area (see, for example, Ochsner and Gross 2005).

But here we will not focus on this transition which enables declarative knowledge to be used to resolve conflicts between elements of the hedonic system. Instead we will focus on a transition which would enable all goals to be established and adapted declaratively, irrespective of whether conflicts are involved—i.e. a transition that would free individuals from the dictates of the hedonic system. If such a transition were accomplished in full, conflicts between goals would be resolved effortlessly.

The potential adaptive benefits of the conscious use of declarative knowledge to set goals are particularly significant given the central role played by goals in determining what gains access to the GW. As we saw earlier, the goals reinforced by the hedonic system strongly influence what we are conscious of, where attention is directed, what is recruited by consciousness and what voluntary actions we initiate. Plans, thoughts, proposed behaviours and objectives that are not relevant to active desires and motivations are unlikely to gain access to the GW and influence behaviour.

As a consequence, humans have difficulty in consistently pursuing objectives that are decided with the assistance of reason, but which are not reinforced by the hedonic system. The pursuit of 'reasoned' objectives is disrupted whenever goals which are supported by the hedonic system cause the GW to recruit sequences of thoughts, plans and behaviours that conflict with the

reasoned objectives. Much of the time, passions dominate reason (see, for example, Dilman 1984, and Ellis 2005).

The hedonic system severely constrains how an individual might respond to a given set of circumstances, as well as what the individual actually experiences. The hedonic system's domination of recruitment to the GW means that consciousness will recruit only from specialist processors and other resources that are relevant to whatever hedonic goals are active. The individual will be incapable of responding in other ways. Furthermore, attention will be given only to those aspects of the individual's environment that are relevant to active goals. Both behaviour and conscious experience is restricted to a small subset of what is logically possible.

However, the declarative transition in this domain would enable reasoned objectives to be pursued without distraction, and for all conflicting goals to be aligned with them. Any conflict between longer-term and more immediate goals, including immediate goals associated with strong emotional responses, could be effortlessly resolved. For example, an individual could decide at will to 'turn the other cheek' even when criticized entirely without justification.

The transition would also enable consciousness to draw on all the individual's resources to develop effective responses, not just those dictated by active goals supported by the hedonic system. It would enable an individual who is experiencing strong desires or emotion to nonetheless respond from the whole of herself, not from just a limited part.

This transition would also be very significant from a wider evolutionary perspective. Potentially it would enable individuals to pursue evolutionary ends directly, rather than continue to pursue proxies that may no longer be relevant. It would free individuals from the dictates of proximate goals established by past evolution, permitting them to transcend their biological and cultural past (Stewart 2001). The transition would enable individuals to use declarative modelling to build resources to assess the evolutionary consequences of alternative behavioural strategies, implementing those that maximise future evolutionary success (Stewart 1997a, 1997b and 2000 discuss in detail the goals that would be consistent with 'future evolutionary success').

The use of declarative modelling to identify and pursue evolutionary advantage is potentially a much more successful strategy in evolutionary terms than the pursuit of proxies for evolutionary success. The proxies have been selected by past circumstances that may change considerably in the future, and can adapt only through selection processes operating across the generations. Evolution produced organisms that are controlled by hedonic systems only because it was unable to immediately produce organisms that could use mental modelling to pursue evolutionary success directly.

However, accomplishment of this declarative transition would not necessarily mean that individuals would pursue evolutionary goals directly. The transition would enable individuals to pursue whatever goals they choose, whether or not the goals are congruent with future evolutionary success. In itself, the transition does not go further than to open up the possibility of conscious evolution.

Processes that regulate the use of sequences of thought and images

We will now consider processes that determine when consciousness is occupied by structured sequences of thought and images, such as thinking, planning, worrying, fantasising, reverie or other chains of declarative contents. These structured sequences take consciousness 'off line' by loading its limited capacity for the duration of the sequence. While a sequence unfolds, it largely precludes the recruitment by consciousness of resources relevant to other adaptive needs.

Declarative modelling appears to have little input into whether an individual engages in a structured sequence of thought in particular circumstances. It is not something that individuals generally deliberate about or have detailed theories about. We do not, for example, decide whether to engage in a sequence of thought on the basis of declarative knowledge about what is optimal for our adaptability in particular circumstances. Nor is it something that is usually under voluntary control. In fact, we cannot voluntarily stop thought for extended periods as a simple experiment demonstrates—if we attempt to remain aware of the second hand of a watch while we stop thought, we find that sequences of thought will soon arise and fully load consciousness, ending our awareness of the second hand. Nor are we able to drive a car along a familiar route and remain fully aware of the driving process for an extended period without engaging in sequences of thought. Furthermore, when individuals are not engaged in a demanding task, rather than remain on-line and aware of what they are doing, they are highly likely to be engaged in reverie, fantasy or other sequences of thought and images (Antrobus et al. 1970; Smallwood and Schooler 2006).

Whether an individual is likely to engage in sequences of thought in particular circumstances is determined largely by procedural knowledge acquired by operant conditioning. Conditioning will tend to entrench engagement in sequences of thought in whatever circumstances the sequences are reinforced by the hedonic system. This will occur even if the hedonic reinforcement is a poor indicator of the actual adaptive utility of the thought—e.g. in the case of much fantasising.

The use of declarative modelling to assist in determining when it is optimal to engage in sequences of structured thought can significantly improve human adaptability. This is primarily because of the extremely limited capacity of consciousness. When the GW is occupied by a structured sequence of thought, images, or other declarative modelling, consciousness is loaded, and perceptions and other representations cannot easily access the GW and recruit other resources. Consciousness is therefore unavailable for any adaptive purpose other than the goal that is being addressed by the sequence of declarative modelling. This holds unless a perception arises that is of such salience that it out-competes the next element in the sequence of thought and gains access to the GW (for example, we can be brought back on-line when driving a car if the driver of another vehicle does something dangerous or unexpected).

Thus while a sequence of structured thought loads consciousness, an individual is not able to respond creatively to whatever is encountered. The individual is unaware of what is occurring in the immediate environment, and consciousness is unable to recruit resources that interpret events and develop responses. Responses to events can only be automatic and habitual, triggered by perceptions that are not interpreted, assessed or adapted with the assistance of consciousness. Nor are conscious processes available to monitor and evaluate the implementation of automatic responses, and to intervene and modify them as necessary where this is possible (e.g. where they have been constructed at least in part with structured thought or other declarative knowledge). If, as appears to be the case, humans are almost continually engaged in sequences of thought and imagining, adaptability is seriously impaired for extended periods. (For evidence and discussion about the propensity of task-unrelated thought to impair the effectiveness of task performance, see Giambra (1995) and Smallwood et al. (2004)).

Apart from crowding out effective responses to other adaptive challenges, thinking is unlikely to develop optimal responses even for the particular adaptive challenge that it is addressing. Each step in a sequence of thought that occupies the GW will recruit only from resources in the brain that can provide the next step in the chain of thought. The sequence can draw only on resources and specialist processors whose outputs are thoughts. The sequence will not have access to the other non-declarative resources in the brain, no matter how useful and ‘wise’ they might be for the particular challenge. For example, the sequence will not have the benefit of input from procedural

knowledge such as intuitions and feelings, except to the extent that these have thoughts as outputs. Access is also prevented to the complex specialist processors that have been put together previously with the assistance of declarative modelling and which now operate largely outside consciousness, having been proceduralized. (For a recent discussion of empirical and anecdotal evidence that rational thought impedes access to intuition and creativity, see Claxton (2000)).

Furthermore, conscious processes such as deliberation, chains of reasoning and other forms of declarative modelling are serial and therefore too slow to develop complex responses in real time. Adaptability is better served by leaving consciousness free to recruit from existing expert processors and resources (including those that have been put together by declarative modelling during periods set aside for planning and contemplation). Consciousness will then be fully loaded only momentarily as it coordinates and recombines expert processors (largely without thought), and as perceptions are assessed and responses are developed outside consciousness by expert processors. Because the loading is not on-going, nothing prevents its interruption at any time. As a result, consciousness will continually have the opportunity to recruit other resources and serve other purposes.

Furthermore, as Baars (1997) emphasizes, consciousness does not have to use thought to recruit relevant resources—all that is required is for consciousness to be directed at the world, and it will effortlessly organise relevant resources to, for example, learn, find patterns, solve problems, and develop behavioural responses. Awareness itself, without thought, can bring all our accumulated knowledge and skills to bear on an adaptive challenge. Sequences of thought cannot. This is why, as Perls (1969) noted, ‘attention in and of itself is curative.’

In summary, the individual should be on-line as much as possible when responses to external events might be required, and off-line only when responses are unnecessary i.e. during periods of contemplation and reflection. When the individual is actively dealing with the world, the loading of consciousness by sequences of declarative modelling will impair adaptability. Even when declarative modelling is being used during periods of contemplation, adaptability would be enhanced if it was interspersed with periods when consciousness is unloaded, enabling all relevant resources to be accessed.

When individuals are not actively engaging with the world, adaptability benefits significantly from periods spent off-line using declarative modelling to construct novel adaptive behaviours and cognitive skills. This is essential for the advancement of the declarative transition. It enables the construction and accumulation of complex specialist processors that can then be relied upon when the individual is actively engaging with the world and when there is insufficient time for serial declarative modelling to construct appropriate responses. The processors can be recruited and recombined by consciousness to generate appropriate behaviour without the need for structured sequences of thought. And their implementation can be monitored and adjusted by relatively unloaded consciousness, without thought. Given that most behaviour is necessarily automatic, there are obvious advantages in the systematic use of consciousness to build the best automatic processors possible.

For all these reasons, when we actively engage with the world, cognition can be more effective in the absence of sequences of thought. However, few would confirm this on the basis of their subjective experience. In part this is because we rarely have direct experience of engaging with the world while our consciousness is unloaded by sequences of thought and imagining. But it is also because we are not directly aware of unconscious information processing in our brain, no matter how complex and effective it is. We are only aware of processing that occupies the GW and becomes contents of consciousness, such as declarative modelling organised by sequences of thought. It is difficult to accept that the cognition we know about (and are continually involved in)

is often less effective than cognition we are not even aware of. But our belief in the importance of thought has a deeper origin: unless we think something, we don't know that we know it. More significantly, if we don't continually think about who we are and why we are acting, we don't know who we are and why we act. In an important sense, if we stop thinking, we cease to know that we exist (Morin 2005).

However, the reluctance to accept that most thought lacks utility dissipates once individuals develop the capacity to engage with the world with unloaded consciousness. This enables them to obtain direct experience of the effects on their adaptability of engaging with the world in this way. In the next section we will turn to examining how such a capacity can be developed

The development by humanity of a capacity to use declarative modelling to consciously regulate the use of sequences of thought would also be significant from a wider evolutionary perspective—it would substantially enhance human evolvability. This transition would be the third in a series of important developments in the evolution of thought processes. The first was the organisation of thoughts into structured sequences. The emergence of organised sequences enabled thoughts to collectively contribute to developing a response to an adaptive issue. This linking together of individual broadcasts (and recruitments) to form ordered sequences enabled conscious processes to build representations and models containing many elements.

Importantly, when combined with a working memory capability, the emergence of structured sequences of thought enabled consciousness to work around the limitations imposed by the serial nature of the GW architecture. By controlling and co-ordinating the contents of the GW over extended periods, structured thought could put together complex sets of resources to model alternative behaviours and their consequences, despite seriality (of course, the emergence of symbolic thought itself was an earlier step in working around the limited capacity of consciousness—as pointed out by Miller (1956), it enabled symbols to occupy consciousness in the place of the complex concepts that the symbols stood for).

As discussed earlier, the second major development in the evolution of thinking was the use of declarative modelling to consciously structure the content of sequences of thought to improve their effectiveness e.g. by ensuring thought processes are rational.

However, these significant improvements in adaptability came at a major cost. As we have seen, the structured sequences of thought that the improvements relied upon restrict access to other resources, including those put together previously by thought. The potential benefits of overcoming this restriction increase rapidly as the range of specialist processors and other resources built by declarative modelling increases. The significance of a capacity to regulate engagement in thought is that it allows this potential to be realised. It enables access to these specialist processors when they are needed (e.g. when actively engaging with the world), while allowing the benefits of structured thought to be realised at other times when access is not as necessary (e.g. during periods of contemplation).

4. REALISING THE POTENTIAL TO ENHANCE THE FUNCTIONING OF CONSCIOUSNESS

What capacities must individuals acquire if the declarative transition is to occur for the processes that set goals and regulate sequences of thought? Obviously, individuals would need to acquire declarative knowledge about what goals to pursue and about when it is optimal to have consciousness unloaded by sequences of thought. But this would not be enough. Merely having declarative knowledge about the optimal outputs of conscious processes will not change those processes. For the transition to occur, individuals will need to develop the capacity to modify the

processes so that they produce the outputs identified by declarative modelling. In particular, to implement alternative goals, individuals would have to be able to prevent pre-existing goals that are supported by the hedonic system from continuing to control behaviour by regulating recruitment to the GW. And to prevent sequences of thought from loading consciousness, individuals would need to be able to stop sequences of thought from gaining access to the GW and perpetuating themselves.

How could an individual develop these capacities? What sources of declarative knowledge are available to guide their acquisition?

Knowledge and practices of religious and contemplative traditions

In general, the development of these sorts of capacities has not been a goal of western empirical psychology and it has not devised practices and procedures for developing them. However, various religious and contemplative traditions more or less explicitly seek to produce key aspects of these capacities in their adherents. In particular, they pursue capacities that appear synonymous with an ability for individuals to free themselves from the dictates of goals established by their hedonic system. They seek, for example, to have their adherents ‘resist temptation’ and ‘turn the other cheek’ (Christianity), free themselves from all desires (Buddhism), and experience equanimity in the face of pleasure or pain (Hinduism).

Most traditions also attempt to develop capacities that enable individuals to consciously regulate their engagement in sequences of thought and imagining. For example, nearly all contemplative traditions seek to produce a mode of functioning that is now commonly referred to as ‘being in the present’. The key characteristic of such a mode appears to be that consciousness is not loaded by sequences of thought. It is referred to as ‘being in the present’ because the individual is ‘on-line’ and able to attend to stimuli arising at each moment, rather than off-line engaged in thoughts that are nearly always about the past or future.

The contemplative traditions teach a wide range of practices and techniques that are explicitly intended to develop these capacities. Broadly, the practices can be understood as operating to disengage the pre-existing processes that enable the hedonic system to control behaviour and that perpetuate sequences of thought, replacing them with new processes that embody the desired capacities.

It is worth noting here that disengagement from automatic processes has also been necessary in past evolution whenever new adaptive capabilities evolve. For example, the emergence of operant learning required the disengagement of innate reflexes, and the evolution of off-line declarative modelling required the disengagement of both learned and innate responses to environmental stimuli (see Barkley (2001), and the discussion by Nagarjuna (2005) of the need to break the ‘stimulus-response loop’ to allow voluntary control of motor behaviour).

Providing a detailed survey and analysis of these practices is beyond the scope of this article. Instead, we will identify some of the key elements that are common to many of the practices and techniques, and show how their reported effects are consistent with the framework developed here. In doing so I will draw on general accounts of the practices of Buddhism (Sogyal Rinpoche 1992; Surya Das 1997; and Bucknell and Kang 1997), Hinduism (Mascaro 1962; Mascaro 1965; and Goleman 1988), the Gurdjieff system (Nicoll 1996; Ouspensky 1949) and on surveys of other contemplative traditions (Wilber 1995 and 2000; Combs 2002). I will also draw on the extensive work being done to integrate key practices such as mindfulness meditation into empirical clinical psychology, and to interpret their injunctions and effects within the theories and models of western

scientific psychology (for recent reviews see Breslin et al. 2002; Kabat-Zinn 2003; Baer 2003; Lau and McMain 2005; Walsh and Shapiro 2006; and Shapiro et al. 2006).

The central elements in many of these practices are:

- (i) to rest attention on a stimulus that is not associated with any goal and that does not recruit any sequences of thought or other resources (an 'inert' stimulus); and to
- (ii) bring attention back to the stimulus whenever it is realised that attention has become to be occupied with thoughts or feelings.

The traditions emphasise that the entire practice, including the realisation by the practitioner that consciousness is occupied by thoughts or feelings, and the movement of attention back to the inert stimulus, is to be undertaken non-deliberatively and non-judgementally. In contrast to cognitive behavioural therapy, thoughts and feelings are not disputed, challenged or evaluated (Lau and McMain 2005). To do so would train further engagement with thoughts and feelings, not disengagement. If elaborative thoughts or judgements arise, these too are disengaged from, and attention returned to the inert stimulus.

A very wide range of internal and external phenomenon can serve as the inert stimulus. One of the most common recommendations is to focus attention on sensations of the breath. Other foci of attention recommended by various contemplative traditions include external objects, visualised objects, internal or external sounds (including chanting and mantras), other physical and mental sensations (including resting attention on awareness itself), repetitious cognitive tasks such as counting or prayer, and goalless emotional states such as reverence, love or feelings of surrender.

The essential effects of this core practice are to train processes that:

- (i) disengage attention from sequences of thought and from the recruitment of resources by the hedonic system; and that
- (ii) maintain consciousness free from on-going loading and from domination by the hedonic system.

In circumstances where these processes are trained, the individual will have the capacity to choose when consciousness is loaded by sequences of contents, and will be able to access any resources to devise suitable responses, not just those dictated by the hedonic system.

Standard learning theory is capable of fully explaining the ability of the core practice to train these capacities (see, for example, Breslin et al. 2002; and Baer 2003).

Most contemplative traditions supplement a core practice of this type with a range of other activities and teachings that enhance its effectiveness, including:

- (i) Adherents are taught to undertake the core practice initially in stimulus-controlled circumstances that require less skill and motivation to be effective. For example, the adherent might withdraw from the distractions and challenges of ordinary life for a period in a retreat, pilgrimage, or monastery (or at least withdraw for short periods each day), and the practice is undertaken as formal sitting meditation in a quiet place with a physical posture that is consistent with un-distracted alertness.

(ii) Particularly in traditions that seek to produce these capacities in the midst of everyday life, steps are subsequently taken to generalise the skills learnt through this initial practice to a progressively wider range of stimulus conditions. This is done to ensure the capacities are not trained only in the limited circumstances in which they are learnt—e.g. on the meditation cushion. The capacities are generalised by performing the core practice while engaged in a progressively wider range of normal daily activities.

(iii) A number of traditions include practices and teachings that have the effect of reducing the difficulties associated with extending the core practice to daily life. For example, they may reduce the likelihood that adherents experience strong negative emotions by promoting acceptance, forgiveness, surrender and an attitude of loving kindness towards all others.

(iv) Some traditions encourage adherents to perform the core practice while putting themselves in situations that they would otherwise avoid because, for example, the situations would produce negative affect (e.g. the Gurdjieff system). Without intentionally experiencing such situations, individuals would not have the opportunity to practice disengagement from the impact of aversive stimuli and negative affect, and their behaviour could continue to be controlled by them. These approaches seem to be founded on similar principles to the desensitization and related therapies developed by clinical psychology (e.g. see Breslin et al. 2002).

(v) Traditions generally provide adherents with a set of narratives and beliefs that motivate the performance of the practices, including by pointing to benefits they supposedly will bring. Claimed benefits range from internal peace and bliss to various forms of eternal life. Traditions are also often accompanied by teachings that question the efficacy of continual involvement in thought, as well as the wisdom of a life spent in the pursuit of self-centred desires and motivations. These teachings are directed at overcoming reluctance to continually disengage from sequences of thought and from responses dictated by feelings and emotions. Until individuals gain some proficiency in disengaging from thought and stilling the mind they are unlikely to discover for themselves the limitations of thought—they will be unable to acquire direct knowledge about the lack of utility of most thought while their consciousness is continually loaded by sequences of thought and imagining.

(vi) Mindfulness meditation (also known as insight meditation and self-observation) is a particularly effective practice common to a number of traditions. It combines the core practice with processes that enable the acquisition of knowledge about the efficacy of thought and the operation of the hedonic system. In mindfulness meditation, perceptions of thoughts and feelings can serve as the inert stimulus. Attention is given to these perceptions as they arise from moment to moment, non-deliberatively and non-judgementally. Whenever the meditator realises that consciousness has come to be occupied with thoughts or feelings, rather than only with inert perception of them, the thoughts or feelings are noted (again non-deliberatively and non-judgementally), and attention is taken back to them as objects of inert perception. Alternatively, attention can be returned to another inert stimulus such as the breath.

This practice facilitates the accumulation of resources and knowledge about the efficacy of thinking and feelings without conflicting with the central objective of the practice—to disengage from thought and from recruitment dictated by the hedonic system. This knowledge is accumulated non-deliberatively and non-judgementally during the practice itself, but can be further analysed deliberately and judgementally during periods set aside for this purpose when the practice is not being undertaken. This process will eventually produce procedural resources that can be recruited and used by consciousness for adaptive purposes without loading consciousness for extended periods.

Global Workspace theory and the capacities developed by religious and contemplative traditions

It is informative to compare the reported effects of the use of these types of practices with the effects that would be predicted by GW theory, as developed throughout this paper. This comparison does not have to be restricted to the impact of the practices on behaviour and cognition. It can also compare the actual experiences reported to be produced by the practices with those predicted by the model. This is possible because the GW model makes some predictions about what will be experienced in particular modes of conscious functioning. It does this by predicting what will occupy the GW in the various modes, and by proposing that these contents of the GW equate to the contents of consciousness.

This is not the place to attempt a full and comprehensive comparison. However, we will illustrate the potential utility of such an approach by making some brief comparisons here.

Overcoming the domination of recruitment by the hedonic system

Some of the reported outcomes of using the practices appear to directly correspond to a mode of functioning which is free from the dictates of the hedonic system. Common descriptions of these outcomes include:

- (i) Equanimity in the face of both pleasurable and painful experiences;
- (ii) Non-attachment to the consequences of actions, whether the consequences are positive or negative;
- (iii) Death of the 'ego' (the term ego is not used in the Freudian sense, but rather refers to the psychological system which continually seeks to satisfy immediate and often self-centred impulses, desires and motivations, such as those associated with the maintenance of social status and image);
- (iv) 'Objective consciousness' and 'impersonal enlightenment', in which individuals cease to have their experience of the world and their relationship to it determined by how it impacts on their personal desires and motivations. For the first time, they see the world 'as it really is'.

Descriptions of the experiences associated with these capacities demonstrate that they do not involve repression of impulses and emotional responses (Sogyal Rinpoche 1992). Not only do the traditions report that emotions and feelings continue to be experienced, they can in fact be experienced more vividly. But they no longer dominate behaviour. For example, a practitioner who is experiencing a strong emotion will not be completely absorbed in the experience, even if acting upon the emotion. The practitioner is free to consider and implement alternative responses at any time.

This is consistent with a mode of functioning in which emotions and feelings no longer dictate what resources are recruited to the GW. In such a mode, although emotions and feelings may be active, consciousness is free to recruit resources for other goals and purposes, including goals not supported by the hedonic system. The individual is able to remain aware of other internal and external stimuli and act on them. Feelings and emotions continue to enter consciousness as they momentarily occupy the GW, and the individual is free to take them into account when deciding what goals to pursue, but they no longer unilaterally control recruitment to the GW. In this mode consciousness is able to recruit resources *about* emotions and feelings, not just resources dictated by them. As a result, feelings and emotions are experienced as objects of perception that rise into

consciousness. They may be given further attention, thought about or acted upon, as decided by the individual.

Researchers exploring the potential of mindfulness meditation as a therapeutic intervention have often described mindfulness as a mode in which feelings and emotions are learnt to be accepted and tolerated non-judgementally (see, for example, Breslin et al. 2002,). Feelings and emotions arise, but no longer need to be acted upon. Again this is consistent with the interpretation developed here: when an individual can tolerate feelings as they arise, and can just ‘be’ with the feelings without reacting to them, the hedonic system no longer has the capacity to dominate recruitment to the GW. Some contemplative traditions suggest that a similar outcome can be achieved through practices that involve surrender or sacrifice.

It is worth distinguishing the practices under consideration here from other approaches that seek to align goals supported by the hedonic system with longer-term goals. The core practice outlined above operates to transcend the dictates of the hedonic system. In contrast, these other approaches attempt to achieve their objectives by modifying the goals supported by the hedonic system, rather than by transcending them. For example, neuro-linguistic programming advocates the use of various techniques to establish emotions and motivations that are consistent with longer-term goals (O’Connor and Seymour 1995), and Ellis (2005) proposes using meta-emotion and voluntary imagery to elicit emotions that are congruent with goals established by reason. In these cases, the ‘replacement’ emotions and motivations continue to dominate recruitment to the GW and the hedonic system is not transcended.

Regulation of sequences of thought

Some of the reported outcomes of using the practices appear to directly correspond to a mode of functioning in which consciousness remains unloaded by sequences of thought and imagining for extended periods. Common descriptions of these outcomes include:

(i) Awakening, silent witnessing, being ‘in the now’, and being ‘in the present’—in these modes of functioning, the individual is reported as being more or less continually aware and able to respond to both internal and external events as they arise in the moment, rather than being absorbed in thought, imagining, emotions and feelings, and responding habitually (from the GW perspective, these are modes in which the capacity of consciousness is unloaded by sequences of thought and free from hedonic control, and therefore is continually on-line and able to recruit from all the available resources in the brain while developing responses to events. Entering these modes is analogous to awakening from a dream because it enables the individual to draw on a much wider range of resources to decide how to respond to a given event. From the wider perspective afforded by these resources, the individual realizes that the previous (normal) mode of functioning was blind to many relevant factors. This is similar to awakening from a dream and realising that responses made during the dream were blind to the wider context which is available to consciousness upon awakening.).

(ii) Having enhanced access to intuition, wisdom, and all of the individual’s knowledge (not just knowledge represented by thought), including the ability to ‘see’ complex patterns at a glance. These forms of knowledge emerge into consciousness when the mind is silent and still (and when, from the GW perspective, the ability of consciousness to access relevant resources is not crowded out by thought, and is not limited to those resources that contribute the next thought in a sequence).

(iii) Experiencing consciousness as more spacious and perceptions as more detailed and vivid, while the mind is silent, still and peaceful (and, from the GW perspective, when more of the

limited capacity of consciousness is available for perception). Some religious traditions interpret these experiences as the presence of God or the divine.

(iv) Acting out of silence and un-knowing, acting spontaneously, and experiencing phenomenon as empty—in these modes of functioning, perceptions, realisations and actions are not accompanied by interpretation, analysis, plans, rehearsals or other sequences of thought (from the perspective of GW theory, if the GW is not occupied by thought and other forms of declarative knowledge, and it is operating mainly by recruiting the outputs of unconscious specialist processors, the individual will not be conscious of how intuitions or judgements are arrived at, and perceptions will not be accompanied by continual commentary and thought. Intuitions, judgements and perceptions will just appear in consciousness.).

As individuals develop these capacities and are able to exercise them in progressively wider ranges of circumstances, they will be able to maintain consciousness that is unloaded by sequences for longer periods. The individual will be less likely to go off-line as a consequence of the occupation of the GW by a sequence of thought and consciousness is less likely to be fully occupied with the pursuit of narrow goals established by the hedonic system. During the periods when consciousness is unloaded, all the resources of the brain will be accessible, and the individual can act from the whole of herself. Consciousness will continually be poised and free to deal with whatever arises, moment by moment.

From the perspective of the contemplative traditions, individuals will be awakened, in the present, or in the witness state for longer periods and in a wider range of circumstances. In the terminology of the Gurdjieff system, individuals will increasingly experience themselves as developing a soul—they will begin to experience themselves as a continual awareness that is unbroken by thoughts, feelings and external stimuli (these will continue to enter consciousness, but will no longer load it or dominate recruitment for extended periods). Initially, this experience of on-going awareness will be interrupted as sequences of thoughts and emotions arise that load consciousness and dominate recruitment, but the experience will increasingly extend for longer periods. In the terminology of Wilber (1997), subject permanence emerges.

From the perspective of Kegan (1982), individuals will spend less time embedded in thoughts and feelings (consciousness will cease to be prevented from considering alternative responses due to loading or domination of recruitment—in the language of the contemplative traditions, the individual will cease to be attached to or identified with the thoughts or feelings). As a consequence, thoughts, feelings and reactions to them can be ‘objects of consciousness’. Events (whether internal or external) can be objects of consciousness only when they cease to elicit sequences of thought or other responses that load consciousness or otherwise dominate recruitment over time. Once an event ceases to elicit such a response, it can be given ‘bare attention’, and consciousness is free to recruit resources *about* the event, or to shift attention to other events or responses. For example, once thoughts or emotions are objects of consciousness we can relate to them in the same way as we relate to inert objects in our external environment—we are free to choose whether to respond to them, including whether to give them further attention. Mindfulness researchers appear to be describing the same phenomenon when they refer to taking a ‘decentered perspective’ or achieving ‘cognitive distancing’ in relation to thoughts and emotions (Hayes et al. 1999; Teasdale et al. 1995).

Consistent with Kegan’s general characterisation of the trajectory of development, what was subject (thoughts, feelings and emotions) increasingly becomes object. In Kegan’s terminology, individuals who have developed these capacities will *have* thoughts, feelings and emotions; thoughts, feelings and emotions will no longer *have* them.

Most of the contemplative traditions suggest that the movement from subject to object in these domains tends to occur in stages. After surveying a wide range of traditions, Wilber (1995) suggested that traditions generally agree on at least three broad stages that he termed the psychic, subtle and causal. From the perspective being developed here, these appear to equate broadly to stages in which:

- (i) The physical environment can be an object of consciousness, and the individual is able to maintain subject permanence while experiencing the external environment, but only in the absence of stimuli that strongly elicit thoughts or feelings (e.g. the individual is able to be 'in the present' when contemplating the natural environment).
- (ii) Thoughts (as well as the physical environment) can be objects of consciousness, and the individual is able to maintain subject permanence while thoughts arise, but only in the absence of stimuli that elicit strong emotions (the individual is able to be 'in the present' except, for example, in affectively-charged social interactions).
- (iii) Emotions, desires and other feelings (as well as thoughts and the physical environment) can be objects of consciousness, and the individual is able to maintain subject permanence while desires and emotions arise.

Within the framework developed by Turchin (1977) and Heylighen (1991), the development of the capacities represented by each of these stages constitutes a metasystem transition.

The culmination of this developmental (and evolutionary) trajectory from subject to object appears to be a mode of functioning referred to by a number of the contemplative traditions as the non-dual mode. It represents an end point because the movement from subject to object is complete—all phenomenon (internal as well as external) are potentially objects of consciousness that can be adapted consciously. Metzinger (2003) suggests that in this mode, consciousness ceases to be continually loaded by what he refers to as the phenomenal model of the self (the PMS) and the phenomenal model of the intentionality relation (the PMIR). The PMS is a representation of the self, including representations of the boundaries between the self and non-self and of the location of the self in the world. The PMIR is a representation of the self in-interaction-with and having-intentions-towards the world (see also Clark 2005). Once the PMIR and the PMS can be objects of consciousness, the individual can experience phenomenon without the PMIR and PMS either loading consciousness or being experienced as part of the subject. As a consequence, the individual in the non-dual mode ceases to experience the world from a perspective centred on the self, no longer experiences boundaries between self and the world, and ceases to experience phenomenon in terms of their relationship to her intentions. Without any privileged perspective, all is experienced 'as one' and no phenomenon, including those associated with the individual, are seen as more central or relevant than others.

The unfolding of the declarative transition

The progressive development of these capacities opens the way for the declarative transition to proceed. As more becomes object, it can be optimised using the superior potential of consciousness aided by declarative modelling. In particular, the acquisition of these capacities will provide individuals with the opportunity to consciously regulate when they will engage in sequences of thought and to implement whatever goals they choose, irrespective of whether the goals are supported by the hedonic system.

But these are potentials only. The extent to which individuals actually take up the possibility of optimising these functions depends on whether they accumulate the relevant declarative

knowledge and use it to build appropriate models and specialist processors. An extremely large and complex body of declarative knowledge needs to be acquired to model (and therefore ‘understand’) the effects of pursuing existing goals in various circumstances, the consequences of modifying the goals in particular ways, and the impacts and effects of adopting alternative goals. And this knowledge cannot begin to be acquired to any extent until the capacities described here have been developed. The capacity to remain in the present while engaging with the world is needed to enable the self-observation and mindfulness that is essential for the acquisition of this declarative knowledge (and comprehensive self-knowledge in general).

Individuals who develop these capacities without acquiring the requisite declarative knowledge will not accomplish the declarative transitions and therefore will not realise the full potential to enhance the functioning of consciousness. Contemplative traditions provide many examples of individuals who develop these capacities without going on to acquire the relevant declarative knowledge (e.g. the ‘silly saints’ of Gurdjieff (Ouspensky 1949)). From the perspective of the declarative transition, enlightenment is a means to an end, not an end in itself.

4. CONCLUSION

Religious and contemplative traditions have accumulated a substantial body of declarative and procedural knowledge about how to modify the functioning of consciousness. This knowledge has the potential to significantly increase human adaptability and evolvability. By enabling conscious processes to be modified, it opens the way for declarative modelling to optimise the functioning of consciousness. This would enable the full capacity of consciousness to discover new adaptations to be used to adapt and enhance consciousness itself. However, the development of these capacities would not just significantly enhance adaptability—it would also change what occupies the GW through time, and therefore what an individual is conscious of. It would change the experience of what it is to be a human being.

However, the explanations and interpretations developed by the contemplative traditions to account for the practices they use and the experiences and capacities they produce are pre-scientific. Their theories have not been disciplined by the scientific method. In particular, they have unnecessarily introduced a plethora of theoretical entities unknown to empirical psychology and science that have little predictive value. Nor have the theories taken advantage of the powerful models and understandings embodied in standard learning theory, clinical psychology, information processing models of cognition, and other areas of cognitive neuroscience.

It can be expected that these deficiencies will be corrected as the knowledge accumulated by the contemplative traditions is integrated into the framework and practice of scientific psychology. As this integration proceeds, it is likely that far more powerful models of the phenomenon and associated processes will be developed, and that these in turn will enable more effective practices and interventions to be developed. Such a re-interpretation of pre-scientific declarative and procedural knowledge in the light of the conceptual frameworks and models of other domains is an important step in the declarative transition in any domain. In particular, it enables discoveries, models and understandings from other areas to be applied to the domain in question, and vice versa. Science has been the key vehicle for this process in the most recent 400 years of human evolution. This paper is a contribution to the early stages of this interpretation and integration process for consciousness (see Walsh and Shapiro 2006 for a recent overview of progress).

The integration of the discoveries of contemplative traditions with scientific psychology can be expected to greatly assist and accelerate the unfolding of this declarative transition across humanity in general. The successful accomplishment of the transition would open up adaptive possibilities of great evolutionary significance. It would, for example, provide humans with the

possibility of choosing to pursue evolutionary goals directly, rather than continuing to pursue proxies for evolutionary success. It would also enable these goals to be pursued more creatively and successfully. Furthermore, as the transition extends to more aspects of consciousness, humans would increasingly be able to choose to adopt particular modes of consciousness to match the needs of different circumstances, just as we now can choose to adopt particular physical postures to match the needs of different physical tasks.

5. REFERENCES

- Anderson, J R 1980, *Cognitive science and its implications*, Freeman, San Francisco.
- Anderson, J R 1983, *The architecture of cognition*, Harvard University Press, Cambridge, MA.
- Antrobus, J S, Singer, J L, Goldstein, S Fortgang, M 1970 Mindwandering and cognitive structure, *Trans N Y Acad Sci*, vol. 32, no. 2, pp. 242-252.
- Baars, B J 2003, How brain reveals mind: neuroimaging supports the central role of conscious experience, *Journal of Consciousness Studies*, vol. 10, no. 9-10, pp. 100-114.
- Baars, B J 2002, The conscious access hypothesis: origins and recent evidence, *Trends in Cognitive Sciences* vol. 6, no. 1, pp. 47-52.
- Baars, B J 1997, In the theatre of consciousness, *Journal of Consciousness Studies*, vol. 4, no. 4, pp. 292-309.
- Baars, B J 1997, *In the theater of consciousness*, Oxford University Press, New York.
- Baars, B J 1988, *A cognitive theory of consciousness*, Cambridge University Press, Cambridge, MA.
- Baars, B J 1983, Conscious contents provide the nervous system with coherent, global information, in R J Davidson and G E Schwartz, eds., *Consciousness and self-regulation*, Plenum Press, New York
- Baer, R A 2003, Mindfulness training as a clinical intervention: a conceptual and empirical review, *Clinical Psychology: Science and Practice*, vol. 10, no. , pp. 125-143.
- Bargh, J A 1997, The automaticity of everyday life, in R S J Wyer, ed., *Advances in social cognition*, vol. 10, pp. 1-61, Erlbaum, Mahwah, N.J.
- Bargh, J A and Chartrand, T L 1999, The unbearable automaticity of being, *American Psychologist*, vol.54, no. 7, pp. 462-479.
- Barkley, R 2001, The executive functions and self-regulation: an evolutionary neuropsychological perspective, *Neuropsychology Review*, vol. 11, no. 1, pp. 1-29.
- Breslin, F C, Zack, M, and McMMain, S 2002, An information-processing analysis of mindfulness: Implications for relapse prevention in the treatment of substance abuse, *Clinical Psychology: Science and Practice*, vol. 9, no. 3, pp.275-299.
- Bucknell, R and Kang, C (eds.) 1997, *The meditative way: readings in the theory and practice of Buddhist meditation*, Curzon, Richmond, Surrey.
- Clark, T W 2005, Killing the observer, *Journal of Consciousness Studies*, vol. 12, no. 4-5, pp. 38-59.
- Claxton, G 2000, *Hare brain tortoise mind: how intelligence increases when you think less*, Ecco Press, New York.
- Combs, A 2002, *The Radiance of being: Understanding the grand integral vision; Living the integral life*. Paragon House, St Paul, MN.
- Das, S 1997, *Awakening the Buddha within*, Bantam Books, Sydney.
- DeHaene, S and Naccache, L 2001, Toward a cognitive neuroscience of consciousness: basic evidence and a workspace framework, *Cognition*, vol. 79, no. 1-2, pp. 1-37.
- Dennett, D 2001, Are we explaining consciousness yet? *Cognition*, vol. 79, no. 1-2, pp. 221-237.
- Dennett, D 1991, *Consciousness explained*, Little, Brown, Boston.
- Dilman, I 1984, Reason, passion and the will, *Philosophy*, vol. 59, pp.185-204.
- Dreyfus, H and Dreyfus, S 1987, *Mind over machine: the power of human intuition*, The Free Press, New York.

- Ellis, R D 2005, The roles of imagery and meta-emotion in deliberate choice and moral psychology, *Journal of Consciousness Studies*, vol. 12, no. 8-10, pp. 140-157.
- Flavell, J H 1985, *Cognitive development*, 2nd ed., Prentice-Hall, Englewood Cliffs, NJ.
- Giambra, L M 1995, A laboratory based method for investigating influences on switching attention to task unrelated imagery and thought, *Consciousness and Cognition*, vol. 4, no. 1, pp. 1-21.
- Goleman, D 1988, *The meditative mind – the varieties of meditative experience*, G. P. Putnam's Sons, New York.
- Hayes, S C, Strosahl, K D, and Wilson, K G 1999, *Acceptance and commitment therapy: an experiential approach to behaviour change*, Guilford Press, New York.
- Heylighen, F 1991, 'Cognitive Levels of Evolution: from pre-rational to meta-rational', in *The Cybernetics of Complex Systems – Self-organisation, Evolution and Social Change*, F. Geyer (ed.), Intersystems, Salinas California, pp.75-91.
- Johnston, V 1999, *Why we feel: the science of human emotions*, Perseus, Cambridge, MA.
- Kabat-Zinn, J 2003, Mindfulness-based interventions in context: past, present and future, *Clinical Psychology: Science and Practice*, vol. 10, no. 2, pp. 144-156.
- Kanwisher, N 2001, Neural events and perceptual awareness, *Cognition*, vol. 29, no.1-2, pp. 89-113.
- Karmiloff-Smith, A 1992, *Beyond Modularity: A Developmental Perspective on Cognitive Science*, MIT Press, Cambridge MA.
- Karmiloff-Smith, A 1986, From meta-processes to conscious access: evidence from children's metalinguistic and repair data, *Cognition*, vol. 23, no. 2, pp. 95-147.
- Kegan, R 1982, *The evolving self*, Harvard University Press, Cambridge, MA.
- Keil, F 1989, *Concepts, kinds and cognitive development*, MIT Press, Cambridge, MA.
- Kuhn, D, Langer, J, Kohlberg, L and Haan, N S 1977, The development of formal operations in logical and moral judgment, *Genetic Psychology Monographs*, vol. 95, pp. 97-188.
- Lau, M A and McMain, S F 2005 Integrating mindfulness meditation with cognitive and behavioural therapies: the challenge of combining acceptance and change-based strategies, *Can J Psychiatry*, vol. 50, no.13, pp 863-869.
- Luria, A R 1973, *The working brain*, Basic Books, New York.
- Mandler, J 1992, How to build a baby, *Psychology Review*, vol. 99, no. 4, pp. 587-604.
- Mascaro, J (trans.) 1962, *The Bhagavad Gita*, Penguin, Harmondsworth.
- Mascaro, J (trans.) 1965, *The Upanishads*, Penguin, Harmondsworth.
- Metzinger, T 2003 *Being no one: the self-model theory of subjectivity*, MIT Press, Cambridge, MA.
- Miller, G A 1956, The magical number seven, plus or minus two: some limits on our capacity for processing information, *Psychological Review*, vol. 63, no. 2, pp. 81-87.
- Morin, A 2005, Possible links between self-awareness and inner speech, *Journal of Consciousness Studies*, vol. 12, no. 4-5, pp.115-134.
- Nagarjuna, G 2005, Muscularity of mind: towards an explanation of the transition of unconscious to conscious. Retrieved November 22, 2006 from the Cogprints archive at <http://cogprints.org/4352/01/mom.pdf>
- Nicoll, M 1996, *Psychological Commentaries on the Teaching of Gurdjieff and Ouspensky*, Samuel Weiser, York Beach Maine.
- Norretranders, T 1998, *The user illusion*, trans. J. Sydenham, Viking, New York.
- Perls, F 1969, *Gestalt therapy verbatim*, Real People Press, Lafayette, CA.
- Ochsner, K N and Fross, J J 2005 The cognitive control of emotion, *Trends in cognitive sciences*, vol. 9, no. 5, pp. 242-249.
- O'Connor, J and Seymour, J 1995, *Introducing neuro-linguistic programming*, Thorsons, Hammersmith, London.
- Ouspensky, P 1949, *In search of the miraculous*, Harcourt Brace Jovanovich, New York.

- Posner, M I and Snyder, C R 1975, Attention and cognitive control, in *Information processing and cognition*, R L Solso, ed., Erlbaum, Hillsdale, NJ.
- Pyllyshyn, Z W 1980, Computation and cognition: issues in the foundations of cognitive science, *Behavioural and Brain Sciences*, vol. 3, pp. 111-132.
- Reber, A 1989 Implicit learning and tacit knowledge, *Journal of Experimental Psychology: General*, vol. 118, no. 3, pp. 219-235.
- Sogyal Rinpoche, 1992 *The Tibetan Book of Living and Dying*, Random House, Sydney.
- Shanahan, M and Baars, B 2005, Applying global workspace to the frame problem, *Cognition*, vol. 98, no. 2, pp. 157-176.
- Shapiro, S L, Carlson, L E, Astin, J A, and Freedman, B 2006, Mechanisms of mindfulness, *Journal of Clinical Psychology*, vol. 62, no. 3, pp. 373-386.
- Smallwood, J and Schooler, J W 2006, The restless mind, *Psychological Bulletin*, vol. 132, no. 6, pp. 946-958.
- Smallwood, J, Davies, J B, Heim, D, Finnigan, F, Sudberry, M V, O Connor, R C and Obonsawain, M C 2004, Subjective experience and the attentional lapse: task engagement and disengagement during sustained attention, *Consciousness and Cognition*, vol. 4, pp. 657-690.
- Stewart, J E 2001, Future psychological evolution, *Dynamical Psychology*, retrieved October 6 2006 from <http://www.goertzel.org/dynapsyc/2001/PsyEv.htm>
- Stewart, J E 2000, *Evolution's Arrow: the direction of evolution and the future of humanity*, Chapman Press, Rivett, Canberra.
- Stewart, J E 1997a, 'Evolutionary Progress', *Journal of Social and Evolutionary Systems*, vol. 20, no. 4, pp. 335-62.
- Stewart, J E 1997b, 'Evolutionary transitions and artificial life', *Artificial Life*, vol. 3, no. 2, pp. 101-120.
- Sun, R 1997, Learning, action and consciousness: a hybrid approach toward modelling consciousness, *Neural Networks*, vol. 10, no. 7, pp. 1317-1331.
- Teasdale, J D, Segal, Z V, and Williams, J M G 1995, How does cognitive therapy prevent depressive relapse and why should attentional control (mindfulness) training help? *Behaviour Research and Therapy*, vol. 33, pp. 24-39.
- Turchin, V 1977, *The Phenomenon of science: A cybernetic approach to human evolution*, Columbia University Press, New York.
- Velmans, M 1991, Is human information processing conscious? *Behavioral and Brain Sciences*, vol. 14, pp. 651-726.
- Walsh, R and Shapiro, S L 2006, The meeting of meditative disciplines and western psychology: a mutually enriching dialogue, *American Psychologist*, vol. 61, no. 3, 227-239.
- Wilber, K 2000, *Integral Psychology*, Shambhala, Boston.
- Wilber, K 1997, *The eye of spirit*, Shambhala, Boston.
- Wilber, K 1995, *Sex, Ecology, Spirituality*, Shambhala, Boston.
- Willingham, D, Nissen, M, and Bullemer, P 1989 On the development of procedural knowledge, *Journal of Experimental Psychology: Learning, Memory, and Cognition*, vol. 15, pp. 1047-1060.
- Wilson, T D 2002, *Strangers to ourselves: discovering the adaptive unconscious*, Harvard University Press, Cambridge, MA.
- Zajonc, R B 1980, Feeling and thinking: preferences need no inferences, *American Psychologist*, vol. 35, pp. 151-175.