

# Content-sensitive inference, modularity and the assumption of formal processing

Mitch Parsell<sup>#</sup>

**ABSTRACT:** *Performance on the Wason selection task varies with content. This has been taken to demonstrate that there are different cognitive modules for dealing with different conceptual domains. This implication is only legitimate if our underlying cognitive architecture is formal. A non-formal system can explain content-sensitive inference without appeal to independent inferential modules.*

## 1. Introduction

Cosmides and Tooby (1989, 1992) have made a sustained case in favor of inferential modularity. One method used to test inferential modularity is whether subjects make different inferences when presented with different contents (Hirschfeld & Gelman, 1994). The Wason selection task, for example, is used to argue for domain specific reasoning because of systematic variations in performance with changes in content. The preference for social presentations of the task is taken to support *social contract theory* (SCT) (Cosmides, 1989; Cosmides & Tooby, 1989, 1992).

According to SCT, the human brain contains an expert system of which is a subroutine for detecting cheaters (Cosmides, 1989; Cosmides & Tooby, 1989, 1992, 1997). Evidence for the theory is supposed to come from the Wason selection data. The Wason task assesses subjects' ability to reason with conditional statements. Extensive research has demonstrated that people deal very badly with this connective, often falling for one of two logically equivalent fallacies: namely, denying the antecedent (if not-p then not-q) or affirming the consequence (if q then p). Interestingly, performance significantly increases when the conditional is put in a social context. Moreover, performance is even better if the task employs *deontic conditionals*: conditional statements describing what a person is obligated or entitled to do (Manktelow & Over,

---

<sup>#</sup> Mitch Parsell, Department of Philosophy, Macquarie University, NSW 2109 Australia, email: mparsell@scmp.mq.edu.au

1995). This appears to support SCT. Systematic variation in inference due to content is explicable if we suppose different contents implement different inferential systems. The preference for social presentations is explicable if we suppose that the social exchange module is inferentially superior to other reasoning systems. But this is not the only possible explanation of the Wason data. Systematic variation in inference due to content need not imply that different contentful domains implement different inferential systems. Systematic variation in inference can occur within a single inferential system if the system is non-formal. Thus only on the assumption of formal processing does content-sensitive inference support distinct inferential modules. But formalism is not the only game in town. As such, the Wason data, independent of a commitment to a formal account of cognition, need not support inferential modularity.

## ***2. Inferential Modularity***

Modern research suggests the mind is modular. The posited specialized systems are extremely varied, including expert modules for “mindreading” (Baron-Cohen, 1995; Leslie, 1987), grammar (Pinker, 1994) and biology (Caramazza & Shelton, 1998; Keil, 1994). Moreover, the nature and extent of modularity is contested. For a neuroscientist modularity refers to the fact that brains are physically structured and different areas perform different information-processing tasks. For a cognitive scientist or evolutionary psychologist a module is an encapsulated mental system evolved to handle specific information of particular importance to the species.

The classic statement of modularity from a cognitive scientist is due to Fodor (1983). Modules are for Fodor (1983) informationally encapsulated, neurologically hardwired and innately specified. They are ‘highly specialized computational mechanisms in the business of generating hypotheses about the distal sources of proximal stimulations’ (Fodor, 1983, p. 47). The critical thing for Fodor (1983, p. 122) is informational encapsulation, which is bi-directional, involving both limited access to and encapsulation from the rest of the mind. There is limited access in that the processing within a module is unavailable to the rest of the mind and there is encapsulation in that modules cannot employ anything other than their own proprietary information sources. According to Fodor modularity is restricted to the periphery of the mind: to input systems, such as vision, and output systems, such as language production and motor

control. General intelligence, including central cognitive processes such as belief fixation and inference, is informationally *unencapsulated*, a system in ‘which any of the available data may be brought to bear on any of the hypotheses that it can entertain’ (Fodor, 1983, p. 122; see also Fodor, 2000, 2002).

Since Fodor’s (1983) classic statement, a number of more radical accounts of modularity have been proposed. Some evolutionary psychologists and cognitive scientists claim the mind is *massively* modular: that the mind consists entirely of evolved modules. According to Cosmides and Tooby (Tooby & Cosmides, 1992; Barklow *et al.*, 1992; Fiddick *et al.*, 2000) even inference is modularized: our cognitive architecture is ‘densely populated with a large number of evolved, content-specific, domain-specific *inference engines*’ (Fiddick *et al.*, 2000, p. 12; emphasis added). These specialized inference systems embody a range of background assumptions about their conceptual domains that allow for the solution of ‘an array of computational problems that plague content-independent systems ... including the combinatorial explosion of possible computations’ (Fiddick *et al.*, 2000, p. 11; see also Tooby & Cosmides, 2000). One of the most discussed proposed modules is SCT. SCT is a computational theory specifying the features an expert system functionally specialized for social exchange should have (Cosmides & Tooby, 1989).

A major line of argument in support of modular accounts of inference comes from evolutionary biology. The use of evolutionary theory to support *inferential* modularity is premised on either the impossibility (Carruthers, 2003a) or inefficiency (Cosmides & Tooby, 1992) of evolving a general-purpose inference system. It is claimed that a general-purpose problem-solver would be out-competed by a suite of special-purpose modules. Tooby and Cosmides (1992) have argued that modular processing is likely to be far faster and computationally more efficient. As such, areas of particular evolutionary importance, such as social exchange, would receive selective pressure to produce expert modules. Carruthers (2003a, b) offers a more extreme version of this argument, claiming that holistic computation is not merely inefficient but impossible. If computational psychology is to be comprehensible, Carruthers (2003b) argues that each cognitive module must contain distinctive inferential processors. Notice that both these arguments are essentially the inverse of Fodor’s (1983) pessimism concerning the prospects of computational psychology.

Fodor is famously pessimistic about the prospects of cognitive science given his non-modular account of central cognition. A holistic general cognitive system of the type envisioned by Fodor would exemplify the significance of global factors in belief fixation. But, Fodor (1983, p. 129) argues, no one has any idea how such factors could possibly have their effects. Thus a science of central cognition is likely to be impossible. Both Cosmides and Tooby, and Carruthers simply invert this schema. Implicit in both these arguments and also in Fodor's pessimistic position, is a commitment to classical computationalism. All three positions assume that cognition is literally a form of computation, operating by means of algorithms defined over sentences or sentence-like structures. But, the argument continues, for such a system to avoid computational explosion, and to thereby be tractable, processes need to be local, in the sense of having restricted access to background knowledge in executing their algorithms. While this is true of symbolic systems it is not necessarily the case for non-symbolic accounts of computation. Non-symbolic architectures can successfully implement holistic processes. Chalmers (1990), for example, has trained a highly-distributed connectionist net to transform active into passive phrases by directly mapping distributed representations onto distributed representations. This, and other connectionist models to be discussed later, clearly demonstrates the possibility of successful holistic inference. [1] Arguing for connectionism over symbolism is well beyond the scope of the present paper. My point is merely that the use of evolutionary arguments in support of inferential modularity, when based on the supposed impossibility or inefficiency of general-purpose holistic inference systems, implicitly assumes a symbolic account of cognition. Thus to the extent that connectionism offers a successful non-modular account of inference it removes this as grounds for accepting inferential modularity. It is worth emphasizing that my main goal in this paper is to offer a connectionist alternative to the modular interpretation of the Wason data specifically, not a connectionist reply to evolutionary arguments for modularity in general. In any case, although arguments from evolution are an important motivation for supposing it *plausible* that the mind is modular, they lack the empirical grounding required to support an *empirical* theory. As will become clear, to my mind only the Wason data provides any serious empirical support for inferential modularity. It is to the empirical support for modularity—of which there are at least three forms: developmental, pathological and experimental—that we now turn. [2]

Cognitive development occurs at substantially different paces in different domains: children demonstrate competence in some areas long before they achieve competence in others. Children seem to be able to reason with normative terms such as “should”, “must” and “permissible” long before they can reason with, for example, modal terms such as “possible” and “necessary”. But this early competence alone does not support modularity. In addition it must be established that children do not receive added exposure to normative situations. Or, since children probably do receive more exposure in this area, it must be established that this extra exposure is not substantial enough to explain their earlier competence in the normative domain. Without such control, or a measure of the effect of differential exposure to achieving competence, the different pace of cognitive development is highly questionable as grounds for supporting modularity. But providing such control would seem impossible. As such, the different pace of cognitive development hardly constitutes knockdown evidence in support of modularity.

Pathological data suggests that there is massive dissociation between various cognitive skills. The most important pathological evidence from the current perspective involves social reasoning. For example, social reasoning can be lost independently of the ability to reason about risks and dangers. This may seem to support modularity of social reasoning, but there is reason to suspect dissociation as a premise for modularity. Consider language abilities where we find the most specific dissociations. As Botterill and Carruthers (1999, p. 60) note ‘[p]atients have been found with specific deficits in the naming of living things, or abstract things, or artefacts, or colours, or bodily parts, or people, or fruits and vegetables - in fact, just about any category of item one can think of.’ If dissociation is evidence of modularity and the naming of fruit can be selectively impaired, does it follow that there is a module for fruit naming? Or ‘just about any category of item one can think of’? The answer must be no, but then what principled distinction is there for arguing for modularity in the case of linguistics generally, but not fruit naming specifically? Dissociative evidence, it seems, demonstrates far more than even the most radical proponent of modularity should be willing to grant. Thus one must give up viewing dissociation data as evidence for modularity or become absurdly modular.

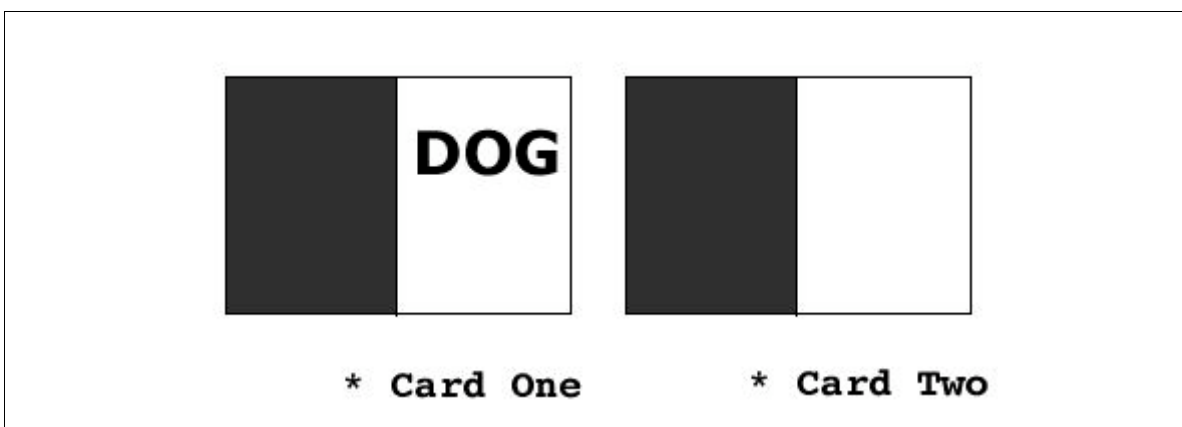
Finally, and most importantly, there is some experimental evidence in support of inferential modularity. Most critically of all, there is the Wason data in support of SCT.

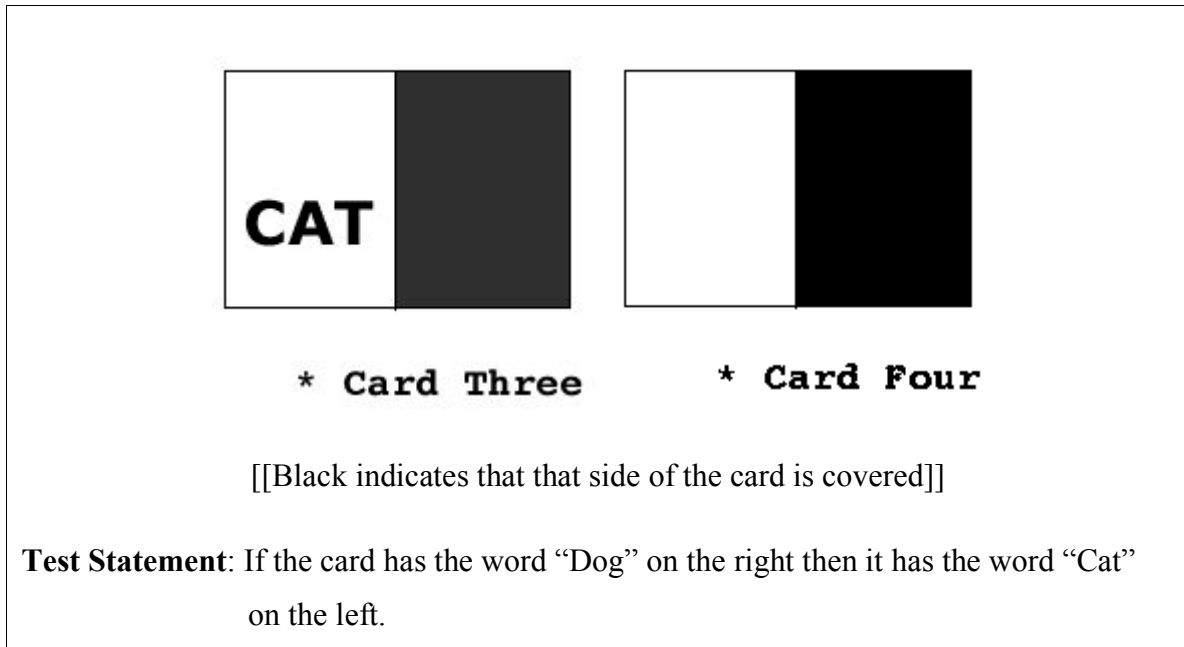
Recall the evolutionary argument for the plausibility of inferential modularity: namely, modular inference is likely to be faster and more efficient and, thereby, there is likely to be selective pressure to produce modular systems for biologically significant tasks. For humans, social exchange is of critical importance. Thus it seems plausible to posit an expert system devoted to the social domain. Moreover, as there is particular evolutionary advantage in the detection of cheaters, we would expect people to have a special ability to detect cheaters. Thus SCT predicts that people will be better at tasks involving the detection of cheaters than logically equivalent non-social tasks. The Wason selection task has been extensively employed to test this cheater detection hypothesis and the prediction has been verified in a number of studies (Fiddick *et al.*, 2000; Cosmides, 1989; Cosmides & Tooby, 1992).

### ***3. The Wason Task as Evidence for SCT***

The Wason selection task tests the ability to reason with the 'if p then q' locution. The task was developed by Wason (1966, 1968) and has been extensively studied by Wason and Johnson-Laird (Johnson-Laird & Wason, 1970; Wason & Johnson-Laird, 1972). [3] As an illustration of the selection task consider the following:

You are presented with four cards, each with one side covered. Your task is to decide which, if any, of the four cards you need to completely uncover to decide *decisively* if the test statement is true of *every card*.





**Figure 1:** An Illustration of the Wason Selection Task

Extensive research has demonstrated that most people perform very badly on the task. Typically, only around 10% of subjects get the test right. In our illustration, the correct response is to want to see card one and card four. Obviously, if card one does not have the word “Cat” on the left it fails the test statement. Equally, though less obviously for most people, if card four does have the word “Dog” on the right it breaks the conditional. Wanting to see card three would be a case of affirming the consequent, whereas wanting to see card two would be a case of denying the antecedent.

Performance on the selection task can be greatly improved with variations in design. One of the earliest identified improvements occurs with *realistic* relations between the antecedent and the consequence of the conditional. Johnson-Laird *et al.* (1972) found, in an experiment using the example of stamps, that if the test statement was ‘if it is sealed then it has a 5d stamp on it’ there was a dramatic increase in logical performance. More recently, Ahn and Graham (1999) have shown subjects improve if the task is given thematic content. Sperber *et al.* (1995) have shown improvements with the right pragmatic context, when the existence of p and not-q is present as highly unusual (for example, machine malfunction). In all these situations performance can increase up to 60%. But the most studied improvements are those which occur with social presentations of the task. Brown and Moore (2000) have recently found subjects perform significantly better on tasks assessing altruism and cheater detection than parallel task in

other domains. Frydman *et al.* (1999) have demonstrated that even 11-yr-olds can perform well on the task provided that the rules take the form of permissions or obligations. One of the most employed social conditional links drinking alcohol with age. Griggs and Cox (1982), for example, provide the following conditional: ‘If a person is drinking beer, then the person must be over 19 years of age’.

It may be tempting to explain improvement on social presentations of the task in terms of the subject’s familiarity with such situations. But logically correct responses can be elicited in *any* presentation that involves a law represented in deontic social terms. This is even true of fictitious laws set in non-existent cultures (Cosmides, 1989). Consider the following example: ‘If a man eats cassava root, then he must have a tattoo on his face’. Subjects perform just as well on this cassava-tattoo law as the culturally familiar drinking age law (Cosmides, 1989). Thus it appears that culturally specific knowledge cannot be the source of better performance. This has led to the claim that only modular theories of inference are capable of explaining improvements in the selection task given social presentations: that to explain such systematic results ‘one is forced to invoke content-specifized inferential machinery, including social contract algorithms’ (Fiddick *et al.*, 2000; p. 5).

Before moving on to my criticism of the Wason data as a grounds for accepting inferential modularity, it is worth re-emphasizing that Cosmides and Tooby’s case for SCT is not restricted to the Wason data. Indeed, they are explicit in their claim that multiple independent convergent lines of evidence should be used to support theories of modularity (Fiddick *et al.*, 2000). Nevertheless, they do see the Wason data as a powerful form of support for the SCT. As such, removing the Wason data as grounds for SCT removes some of the most powerful evidence for SCT. Moreover, as I have attempted to argue, only the Wason data provides any firm *empirical* foundation for inferential modularity.

#### ***4. A Connectionist Alternative to Inferential Modularity***

Cosmides and Tooby are right to claim that the Wason data supports SCT *if* one assumes a symbolic account of cognitive life. But symbolism is no longer the only game in town. Connectionist nets offer a competing explanation of cognition and inference that is at least worthy of serious consideration. Moreover, connectionists generally avoid modularity and, as such, connectionism seems a plausible place to look for a non-modular

account of the Wason data. [4]

Classical cognitive modeling is premised on the Physical Symbol Systems hypothesis (Newell & Simon, 1976). Symbolic systems achieve intelligent behavior by the rule-governed manipulation of physical items (*symbols*) that can be non-semantically individuated. Processing is purely formal (non-semantically) in the precise sense that inferential rules apply according to the physically instantiated properties of the symbol structures (namely, their syntactic form). To enable meaningful inference syntax mirrors semantics. The meaning of a symbol is systematically related to its syntactic form. Thus it is possible to instantiate syntactic processes that respect content without appeal to content. On this model of intelligent behavior, systematic variation in inference due to content cannot result from systematic variation in the inferential *process*, as the process of inference is not sensitive to content. Thus to explain content-sensitive inference it is necessary *on the assumption of formal processing* to suppose inference is modular. But some connectionist networks are non-formal.

Connectionist nets come in a variety of different flavors, with most not significantly different from formal models. Local and semi-distributed networks process information in an essentially symbolic, and by extension formal, manner. [5] Highly-distributed nets, on the other hand, are non-symbolic and, by extension, non-formal. To appreciate this, consider the property of *deep context-sensitivity*. An item is context-sensitive if it varies depending on the context in which it occurs. A semantic item is context-sensitive if the *meaning* associated with the item varies with different situations. An item is *deeply* context-sensitive if it does not have a *fixed* content independent of context. A representational *system* is deeply context-sensitive if it has no content bearers that carry a fixed meaning independent of context. This is exactly the situation with highly-distributed connectionist nets. Content only arises as a response to the total input, as the items attributed representational content—patterns of node activation—only arise when input is received. Moreover, the content can only be precisely fixed with reference to both the network's specific learning environment and exact computational structure. [6]

The content of the hidden-layer of a distributed net is usually understood in one of two ways. First, the hidden-layer is sometimes taken as representing the input constituents (Goschke & Koppelberg, 1991). Second, the content of the hidden-layer is

often derived from a grouping operation, such as some form of cluster analysis (Sejnowski & Rosenberg, 1987) or a nearest neighbor-based metric (Elman, 1990). Neither view allows a precise understanding of context. To achieve this it is necessary to base an understanding of content on a relation sensitive to the particular context in which representations are developed. Boden and Niklasson (2000) propose a relation that is sensitive to the hidden-layer's particular weight configuration, which is determined by the training set and the computational structure of the network. The only context supplied to a network is provided in the training set. As nets begin life with randomly assigned weights, there is no content prior to training, so all content depends on the training environment. As such, all representations reflect the context of their particular learning environments. But the representational economy of a net can be altered in two ways. First, the learning environment can be changed. Second, the representational economy can be altered by the addition (or subtraction) of processing primitives (nodes). Thus content is only definable with reference to both the training set and the precise computational structure of the network. Complex expressions will qualify as context-dependent as they are a direct result of the exact weight configurations formed from the *specific* learning environment. This context-sensitivity includes the formation of representations (compositional context) and the processes applied to the representations, as both are products of the learning environment. All hidden-layer content is context-dependent since activation within this layer is a combined function of the input and the hidden-to-output layer weight configuration. The content of the input-layer depends on the content of complex expressions as the learning procedure may affect the input-to-hidden layer weights resulting in a different content being extracted from the input. Thus highly-distributed connectionist nets are context-dependent at the input, hidden and output layers. This entails inference is content-sensitive, as it is the total input, together with the learning environment and exact computational structure, which provides the context to the network. As such, content-sensitive inference, even within the one module (network), is to be expected.

Highly-distributed nets have been successfully employed as inferential models. There are a large number of connectionist models of inference, encompassing the entire gamut of connectionist architecture types (see Browne & Sun (2001) for an exhaustive survey of the literature). Highly-distributed nets are generally favored over local and

semi-distributed systems, as both these lack expressive power and are unable to avoid cross talk between multiple conjunctive concepts. Nevertheless, distributed models face difficulties. There is the problem of explaining systematicity, which has long been recognised as a major challenge by both critics (Fodor & Pylyshyn, 1988) and supporters (Smolensky, 1988) of the connectionist paradigm. Recent research suggests that tensor product networks may offer reasonably strong systematicity of inference (Halford *et al.*, 1998; Phillips, 1998; Phillips & Halford, 1997). But such architectures suffer from computational explosion: specifically, there is an exponential increase in the number of elements required for representational binding (for details see Browne & Sun, 2001, p. 1346). As such, they are unsuitable as an alternative to modular theories of inference. But the problem of computational explosion can be overcome, without giving up the expressive power required to explain systematicity, if novel recursive representations are employed. One prominent example of a system that employs recursive representation is the Recursive Auto-Associative Memory (RAAM) networks developed by Pollack (1990). The integration of process and data in the RAAM architecture enforces a structural organisation of context-dependent representations. RAAM and other architectures with recursive representations have been successfully employed to deal with context-sensitive language processing.

Modeling human language abilities has long occupied a central place in connectionist research (see McClelland & Rumelhart, 1981; Rumelhart & McClelland, 1982; McClelland & Elman, 1986; Elman & McClelland, 1986; Pollack, 1991; Norris, 1994; Hare & Elman, 1995). Recently there have been a number of studies into context-sensitive languages. Chalup and Blair (1999) have shown that second-order nets trained with back-propagation can learn and demonstrate good generalization abilities within a finite subset of context-sensitive languages. Gers and Schmidhuber (2001) have demonstrated the same abilities with Long-Short Term Memory (LSTM) nets and Boden and Wiles (2000, 2002) have done the same with second-order Sequential Cascade Networks (SCN: Pollack, 1991). Forming Global Representations with Extended Back-Propagation (FGREP) nets have been used to form representations appropriate to multiple inference tasks (Miikkulainen & Dyer, 1988). These and many other studies besides show that certain connectionist architectures have the resources to deal adequately with complex inferential tasks that require context-sensitivity. Indeed, they establish that

distributed nets can deal with context in language tasks and (multiple) inference tasks more generally. This places distributed nets in an ideal position to explain the Wason data without positing inferential modularity.

There is only one connectionist model of the Wason task available in the literature. This model, due to Leighton and Dawson (2001), clearly demonstrates the critical importance of context to connectionist inference. Leighton and Dawson's (2001) model consists of three networks corresponding to three possible responses to the selection task: (i) the p card, (ii) the p and not-q cards and (iii) the p and q cards. The three networks required different hidden-layer resources and training epochs to achieve their respective tasks. The p net required 3 hidden-layer units and 83 epochs to converge to a solution. The p and not-q net required 8 hidden units and 414 epochs to converge on a solution. Finally, the p and q network required 8 hidden units and 115 epochs to converge to a solution.

These results indicate that selecting two cards is more difficult than selecting a single card regardless of the presence of negations (selecting p and q is as difficult as selecting p and not-q). This contradicts some suggestions that the selection of not-q is inherently difficult (see, for example, Braine, 1978). Leighton and Dawson (2001) argue that the not-q response is extremely difficult to generate in abstract, but in meaningful version of the task the not-q response is easier to generate as subjects can rely on their *everyday* expertise. If the deeply context-sensitive account of inference I have been developing in this paper is true then it is not merely everyday expertise that mediates the ability to generate the not-q response in a concrete context. Rather the ability to generate the not-q response is mediated by the context provided by the non-abstract setting *regardless of familiarity* with the *specific* situation. Even unfamiliar or imaginary situations can and are interpreted as social exchanges. The cassava-tattoo law is clearly an example of an exchange between social players, whether one knows what a cassava is or not. In a highly-distributed network interpreting a situation as a social exchange can effect the interpretation of the conditional statement regardless of whether the exchange is familiar or not. Thus within in highly-distributed net, even when the context is imaginary, different interpretations of the context can effect the meaning of the 'if p then q' locution and thereby significantly effect the processing of the system. As such, it is possible to explain the preference for social exchange even with *unfamiliar or imaginary*

contexts without positing a separate inferential module for social exchange.

Leighton and Dawson (2001) have clearly demonstrated the mediating factor in producing the correct response to the selection task is highly susceptible to contextual effects, indicating that the specific context is likely to be the major determinant of selection task performance. Highly-distributed connectionist models provide an ideal means of including just this critical factor in inferential systems. Moreover, I have suggested that distributed systems employing novel recursive representations, such as RAAM architectures, are especially promising as models of inferential tasks requiring context-sensitive processing. Indeed, that such architectures can offer an explanation of content-sensitive processing without the need to posit modular inferential systems.

## ***5. Conclusion***

Humans are good at reasoning about social situations. Better than they are at reasoning about logically parallel situations in other semantic domains. And this does not seem to be a result of the familiarity of the social domain, as the social superiority effect is retained even under experimentally controlled situation involving imaginary exchanges. One possible explanation of this is that an inferentially superior module handles social exchange. But this is not the only possible explanation. Indeed, one need only posit such modularity if one supposes inferential processing to be independent of context. Symbolic models are insensitive to this form of context. As such, if one supposes cognition to be symbolic, systematic variation in inferential performance as a result of variation in content does support inferential modularity. But symbolism is not the only plausible account of cognitive life. Connectionism offers a competing account of human cognitive performance.

Highly-distributed connectionist nets are context-sensitive at the input, hidden and output layers. All representations within a highly-distributed net reflect the specific learning environment of the network. As such, the particular interpretation of a situation, even if the situation is unfamiliar or imaginary, can impact the interpretation of conditional statements and, thereby, alter the processing of the network. This carries significant implications for using the Wason data as grounds for accepting SCT. Highly-distributed nets can explain systematic variation in inferential performance as a result of content without the need to posit separate inferential modules. This includes contexts that

are unfamiliar or even imaginary. Indeed, merely interpreting a situation as a social exchange can result in significantly altered performance within a highly-distributed net.

In sum, systematic variation in reasoning on the Wason selection task with variation in content *does not* support inferential modularity without assuming formal processing. This is not to deny that the human mind *may be* modular. Nor that we could have good evidence in support of modularity. The claim is only that the Wason data does not support modularity without assuming symbolism.

### ***Notes***

[1] Of course, Fodor (most importantly in Fodor & Pylyshyn, 1988) has raised what he believes are strong independent grounds for rejecting connectionism as an account of human cognitive processing. Similarly, though far less persuasively, Carruthers (2003a, b) suggested that classical computationalism is the only serious game in town. Cosmides and Tooby, however, do not offer any reason for their (implicit) commitment to symbolism.

[2] My discussion of the evidence for modularity closely follows Carruthers (2003b). Fiddick *et al.* (2000) also mention studies in economic theory, which appears to lack serious empirical support.

[3] Strictly, the first experiment by Wason using the ‘if p then q’ locution to demonstrated substantial reasoning errors was performed in 1964. This experiment did not, however, employ a *selection* task, but required subjects to answer a simple modal question.

[4] Though *most* connectionists avoid modularity there are a variety of modular connectionist networks in use. The most commonly employed are manually designed expert systems (see Husken *et al.*, 2002; Jordan & Jacobs, 1995; Sharkey, 1996). Miikkulainen and Dyer (1989, 1991) have also done extensive work on modular connectionist models of language processing.

[5] In a local network, the processing primitives (nodes of the network) and the representational primitives coincide. (The coincidence of representational and processing primitive defines what it is to be symbolic). In a distributed network, the processing and representational primitives do not coincide, as content is attributed to patterns of nodes. Nevertheless, as each node is only associated with a single content, processing proceeds

in an essentially symbolic manner. For more detail see Clark (1993).

[6] The discussion of context-sensitivity that follows is based on the recent work of Boden and Niklasson (2000).

### **References**

Ahn, W. and Graham, L.M. (1999). The impact of necessity and sufficiency in the Wason four-card selection task. *Psychological-Science* 10, 237-242.

Barkow, J., Cosmides, L., & Tooby, J., (1992). *The adapted mind: evolutionary psychology and the generation of culture*. New York: Oxford University Press.

Baron-Cohen, S. (1995). *Mindblindness: an essay on autism and theory of mind*. Cambridge, MA: MIT Press.

Boden, M. and Niklasson, L. (2000), Semantic Systematicity and Context in Connectionist Networks. *Connection Science* 12, 111-142.

Boden, M. & Wiles, J. (2000). Context-free and context-sensitive dynamics in recurrent neural networks. *Connectionist Science* 12, 197-210.

Boden, M. & Wiles, J. (2002). On learning context free and context sensitive languages. *IEEE Transactions on Neural Networks* 13, 491-493.

Botterill, G. & Carruthers, P. (1999). *The Philosophy of Psychology*. Cambridge, MA: Cambridge University Press.

Braine, M.D.S. (1978). On the relation between the natural logic of reasoning and standard logic. *Psychological Review* 85, 1-21.

Brown, W.M. & Moore, C. (2000). Is prospective altruist-detection an evolved solution to the adaptive problem of subtle cheating in cooperative ventures? *Evolution and Human Behaviour*, 21, 25–37.

Browne, A. & Sun, R. (2001). Connectionist Inference Models. *Neural Networks* 4, 1331-1355.

Caramazza, A., & Shelton, J. (1998). Domain-specific knowledge systems in brain: the animate-inanimate distinction. *Journal of Cognitive Neuroscience* 10, 1-34.

Carruthers, P. (2003a). Is the mind a systems of modules shaped by natural selection? In C. Hitchcock (ed.). *Contemporary Debates in the Philosophy of Science*. Oxford: Blackwell.

Carruthers, P. (2003b). Moderately massive modularity. In A. O'Hear (ed.). *Mind and Persons*. Cambridge, MA: Cambridge University Press.

- Chalmers, D.J. (1990). Syntactic transformations on distributed representations. *Connectionist Science 2*, 53-62.
- Chalup, S. & Blair, A.D. (1999). Hill climbing in recurrent neural networks for learning the a"b"c" language. In *Proceedings of the 6th International Conference on Neural Information Processing*, 508-513.
- Clark, A. (1993). *Microcognition: Philosophy, Cognitive Science and Parallel Distributed Processing*. Cambridge, MA: MIT Press.
- Cosmides, L. (1989). The logic of social exchange: has natural selection shaped how human reason? Studies with the Wason selection task. *Cognition 31*, 187-276.
- Cosmides, L., & Tooby, J. (1989). Evolution psychology and the generation of culture, Part II. Case study: a computational theory of social exchange. *Ethology and Sociobiology, 10*, 51-97.
- Cosmides, L., & Tooby, J. (1992). Cognitive adaptations for social exchange. In J. Barkow, L. Cosmides & J. Tooby (Eds.). *The adapted mind*. New York: Oxford University Press.
- Cosmides, L., & Tooby, J. (1997). Towards an evolutionary taxonomy of treatable conditions. *Journal of Abnormal Psychology, 108*, 453-464.
- Cosmides, L. & Tooby, J. (2000). Consider the source: The evolution of adaptations for decoupling and metarepresentation. In D. Sperber (ed). *Metarepresentations: A multidisciplinary perspective*. New York: Oxford University Press.
- Elman, J.I. (1990). Finding structure in time. *Cognitive Science, 14*, 179-211.
- Elman, J. & McClelland, J. (1986). Exploring lawful variability in speech waveforms. In S. Perkell & D. Klatt (eds.). *Invariance and Variability in Speech Processing*. Seattle: Erlbaum.
- Fiddick, L., Cosmides, L. & Tooby, J. (2000). No interpretation without representation: the role of domain-specific representations and inferences in the Wason selection task. *Cognition, 77*, 1-79.
- Fodor, J.A. (1983). *Modularity of Mind*. Cambridge, MA: MIT Press/Bradford Books.
- Fodor, J.A. (2000). Why are we so good at catching cheaters? *Cognition, 75*, 1-79.
- Fodor, J.A. (2002). Reply to Beaman. *Cognition, 83*, 221.
- Fodor, J.A. & Pylyshyn, Z.W. (1988). Connectionism and cognitive architecture: a critical analysis. *Cognition, 28*, 3-71.

- Frydman, O., Light, P. and Alegria, J. (1999). Pragmatic determinants of children's responses to the Wason Selection Task. *Psychologia*, 42, 59-68.
- Ger, F.A. & Schmidhuber, J. (2001). LSTM recurrent networks learn simple context free and context sensitive languages. *IEEE Transactions on Neural Networks*, 12, 1333-1340.
- Goschke, T. & Koppelerberg, D. (1991). The Concept of Representation and the Representation of Concepts in Connectionist Models. In W. Ramsey, S. Stich and D. E. Rumelhart (eds.). *The Philosophy of Connectionist Theory*. Seattle: Lawrence Erlbaum.
- Griggs, R., & Cox, J. (1982). The elusive thematic-material effects in Wason's selection task. *British Journal of Psychology*, 73, 407-420.
- Hare, M. & Elman, J.L. (1995). Learning and morphological change. *Cognition*, 56, 61-98.
- Halford, G. Wilson, W. & Phillips, S. (1998). Processing capacity defined by relational complexity: implications for comparative, developmental, and cognitive psychology. *Behavioral & Brain Sciences*, 21, 803-831.
- Hirschfeld, L., & Gelman, S. (1994). *Mapping the domain: domain specificity in cognition and culture*. New York: Cambridge University Press.
- Husken, M., Igel, C. & Toussaint, M. (2002). Task-dependent evolution of modularity in neural networks. *Connection Science*, 14, 219-229.
- Johnson-Laird, P.N. & Wason, P.C. (1970). Insight into a logical relation. *Quarterly Journal of Experimental psychology*, 22, 49-61.
- Johnson-Laird, P. N., Legrenzi, P., & Sonino-Legrenzi, M.S. (1972). Reasoning with a sense of reality. *British Journal of Psychology*, 63, 395-400.
- Jordan, M. I. & Jacobs, A. (1995). Modular and hierarchical learning systems. In M.A. Arbib (ed.). *The Handbook of Brain Theory and Neural Networks*. Cambridge, MA: MIT Press.
- Keil, F. (1994). The birth and nurturance of concepts by domain: the origins of living things. In L. Hirschfeld, & S. Gelmann (Eds.). *Mapping the domain: domain specificity in cognition and culture*. New York: Cambridge University Press.
- Leighton, J.P. & Dawson, M.R.W. (2001). A parallel distributed processing model of Wason's selection task. *Journal of Cognitive Systems Research*, 2, 207-231.

- Leslie, A. (1987). Pretense and representation: the origin of "theory of mind." *Psychology Review*, 94, 412-426.
- Manktelow, K., & Over, D. (1995). Deontic reasoning. In S. Newstead, & St. B. T. Evans (Eds.). *Perspectives on thinking and reasoning: essays in honour of Peter Wason*. Hove: Lawrence Erlbaum.
- McClelland, J. & Elman, J. (1986). The TRACE model of speech perception. *Cognitive Psychology*, 18, 1-86.
- McClelland, J.L. & Rumelhart, D.E. (1981). An interactive action model of context effects in letter perception: Part 1. An account of basic findings. *Psychological Review*, 88, 373-407.
- Miikkulainen, R. & Dyer, M.G. (1988). Forming global representations with extended backpropagation. *Proceedings of the IEEE Second Annual International Conference on Neural Networks*, 285-292.
- Miikkulainen, R. & Dyer, M. G. (1989). A modular neural network architecture for sequential paraphrasing of script-based stories. In *Proceedings of the International Joint Conference on Neural Networks*, 49-56.
- Miikkulainen, R. & Dyer, M.G. (1991). Natural language processing with modular PDP networks and distributed lexicon. *Cognitive Science*, 15, 343-399.
- Newell, A., & Simon, H. (1976). Computer science as empirical enquiry. In J. Haugeland (Ed.). *Mind Design*. Cambridge, MA: MIT Press.
- Norris, D. (1994). Shortlist: a connectionist model of continuous speech recognition. *Cognition*, 52, 189-234.
- Phillips, S. (1998). Are feedforward and recurrent networks systematic? Analysis and implications for a connectionist cognitive architecture. *Connection Science*, 10, 137-160.
- Phillips, S. & Halford, G.S. (1997). Systematicity: psychological evidence with connection implications. *Proceeding of the 19th Annual Conference of the Cognitive Science Society*.
- Pinker, S. (1994). *The language instinct*. New York: Morrow.
- Pollack, J. (1990). Recursive Distributed Representations. *Artificial Intelligence*, 46, 77-105.
- Pollack, J. (1991). The induction of dynamical recognizers. *Machine Learning* 7: 227.

- Rumelhart, D.E. & McClelland, L.J. (1982). An interactive activation model of context effects in letter perception: part 2. The context enhancement effect and some tests and extensions of the model. *Psychological Review*, 89, 60-94.
- Sharkey, A. (1996). Modularity, combining and artificial neural nets. *Connection Science*, 9, 3-10.
- Sejnowski, T. and Rosenberg, C. (1987). Parallel networks that learn to pronounce English text. *Complex Systems*, 1, 145-168.
- Smolensky, P. (1988). On the proper treatment of connectionism. *Behavioral & Brain Sciences*, 11, 1-23.
- Sperber, D., Cara F., & Girotto, V. (1995). Relevance theory explains the selection task. *Cognition*, 57, 31-97.
- Tooby, J. & Cosmides, L. (1992). The psychological foundations of culture. In J. Barkow, L. Cosmides & J. Tooby (eds.). *The Adapted Mind*. New York: Oxford University Press.
- Wason, P.C. & Johnson-Laird, P. (1972). *Psychology of reasoning: structure and content*. Cambridge MA: Harvard University Press.
- Wason, P.C. (1964) The effects of self-contradiction on fallacious reasoning. *Quarterly Journal of Psychology*, 16, 30-34.
- Wason, P.C. (1966). Reasoning. In B. Foss (ed.) *New Horizons in Psychology*. Penguin: London.
- Wason, P.C. (1968). Reasoning about a rule. *Quarterly Journal of experimental Psychology*, 20, 273-281.