Domains, Brains and Evolution

Michael Wheeler[1] and Anthony Atkinson[2]

[1]Department of Philosophy,

University of Dundee,

Dundee, DD1 4HN, UK

Email: mwwheeler@philos.dundee.ac.uk

[2]Psychology Department,

King Alfred's College, Winchester,

SO22 4NR, UK

Email: A.Atkinson@wkac.ac.uk

http://www.wkac.ac.uk/psychology/

The order of authors is arbitrary

Preprint of chapter published in: D. M. Walsh (Ed.), *Naturalism, evolution and mind. (Royal Institute of Philosophy Supplement 49.)* (pp. 239-266). Cambridge: Cambridge University Press.

Abstract

Our aim in this paper is to do some conceptual spring-cleaning. Several prominent evolutionary psychologists have argued that the human cognitive architecture consists in a large number of domain-specific features, rather than, as dissenters claim, a small number of domain-general features. The first difficulty here is that there exists no widely agreed-upon definition of 'domain'. We show that evolutionary psychology has the resources for such a definition: a domain is defined as an adaptive problem, or a set of suitably related adaptive problems. Adopting this definition, we proceed to introduce the distinction between data and algorithms, and to differentiate four conceptions of our cognitive architecture, only two of which, we argue, are viable: (a) general-purpose mechanisms operating on domain-specific information, and (b) special-purpose mechanisms operating on domain-specific information. Typically, evolutionary psychologists argue in favour of (b), as against (a). Following a defence of this position against a recent claim that the process of exaptation makes general-purpose mechanisms evolutionarily plausible, we consider the strongest of the evolutionary psychologists' in-principle arguments for the evolutionary *im*plausibility of general-purpose mechanisms. This argument is based on two requirements: that the human cognitive architecture must (i) be capable of solving all the adaptive problems faced by our ancestors, and (ii) have outperformed all competing designs. Work in artificial intelligence suggests that although requirement (i) might be met by general-purpose mechanisms coupled with domain-specific information, requirement (ii) won't. Nonetheless, we propose (tentatively) that relatively general-purpose mechanisms might result from the operation of multiple, simultaneous, systematically related selection pressures. An examination of this proposal, however, brings into sharp relief the fact that, in many evolutionary scenarios, it simply may not be possible to establish a robust distinction between domain-specific and domain-general features.

I Introduction

According to Darwinian thinking, organisms are (for the most part) designed by natural selection, and so are (for the most part) integrated collections of adaptations, where an adaptation is a phenotypic trait that is a specialised response to a particular selection pressure. For animals that make their living in the Arctic, one adaptive problem is how to maintain body temperature above a certain minimum level necessary for survival. Polar bears' thick coats are a response to that selection pressure (surviving in extreme cold). A thick coat makes a positive difference to a polar bear's fitness, since polar bears with very thin coats left fewer offspring than those with thicker coats. The foundational idea of *evolutionary psychology* is that brains are no different from any other organ with an evolutionary function, insofar as brains too are systems shaped by natural selection to solve adaptive problems. Thus brains have a particular functional organisation because their behavioural effects tend, or once tended, to help maintain or increase the fitness of organisms with those brains. Prominent evolutionary psychologists[1] have endorsed the view that the last time any significant modifications were made by natural selecj 50 0 0 50 1382 .ogy unc with the arguments of the evolutionary psychologists in question, and those of their critics. However, our intention is not to give a 'final' answer to the question of whether our psychology is correctly explained in terms of domain-specific or domain-general features. (Taking on board the thought that some degree of functional specialisation is to be expected in the human mind/brain, perhaps the right question is: how much of our evolved psychology is correctly explained in terms of domain-specific features, and how much in terms of domain-general features? We shall not give a 'final' answer to that question either.) Our aims are more modest. For while we are, in general, highly sympathetic to much of what the evolutionary psychologists say, it seems to us that this particular issue is clouded by the imprecise and indiscriminate use of the most basic terms in which the all-important claims and counter-claims in the dispute are expressed, terms such as 'domain', 'domain-specific' and 'domain-general'. This conceptual imprecision and promiscuity invites confusion and misinterpretation, and results in many supporters and critics talking past each other. Our principal intention in this paper, then, is to attempt some tidying up of the conceptual space encapsulated by the thorny and difficult matter of domain specificity, domain generality, and evolved cognition.[4]

II Domains and Brains

One cannot hope to conduct a proper investigation of whether evolved psychological systems are domain-specific or domain-general in character without first saying what is meant by the term 'domain'. Unfortunately, principled definitions of what constitutes a domain are scarce in psychology. Examples and characterisations are more prevalent.[5] The first task of this section is to say why evolutionary psychology provides us with a good candidate for the sort of principled definition that we need.

A useful starting point is Fodor's influential discussion of domain specificity. Fodor's thought is that certain psychological mechanisms are domain-specific in that they are specialised to operate with distinct and limited sets of inputs or information types. Or as Fodor himself puts it, 'only a relatively restricted class of stimulations can throw the switch that turns it [a mechanism] on'.[6] According to this picture, then, a domain is a 'stimulus domain', that is, a type of stimulus for which humans have a special capacity to respond. Some stimulus types that are good candidates for stimulus domains are faces, non-face objects and language. Certain specific properties of stimuli, such as phonemes, colour, shape and three-dimensional form, are also candidates for stimulus domains, for there is much evidence that we are especially sensitive to these properties, as compared to other properties of the same types of stimuli (e.g., non-speech properties of sounds).

Despite the influential contribution of Fodor's discussion, the concept of a domain is sometimes used in a broader sense than that of a stimulus domain. Hirschfeld and Gelman, for example, outline what they 'take to be a fairly uncontroversial characterisation' of the concept, as follows:

> A domain is a body of knowledge that identifies and interprets a class of phenomena assumed to share certain properties and to be of a distinct and general type. A domain functions as a stable response to a set of recurring and complex problems faced by the organism. This response involves ... perceptual, encoding, retrieval, and inferential processes dedicated to that solution.[7]

In fact it seems there are at least five conceptions of a domain expressed or implied in this quotation.

(1) A specific body of information that constitutes the input to some perceptual mechanism or process, or to a mechanism or process that is more cognitive than perceptual, such as those involved in belief fixation, reasoning, or problem solving.

(2) An information-processing mechanism (or set of such mechanisms) dedicated to producing established responses to a problem (or set of problems) routinely faced by the organism.

(3) The output of an information-processing mechanism, where 'output' is taken to mean not behaviour, but rather an information-bearing representational state that is then delivered as input to some other, cognitively downstream information-processing mechanism.

(4) The behavioural response, or set of responses.

(5) The problem, or set of problems.

Can any of these options be turned into the kind of principled definition of a domain that we are seeking? The answer, we think, is 'yes', and it is evolutionary psychology that shows us how. According to the fifth and final option on the list, a domain is a problem or set of problems routinely faced by an organism. To elaborate this idea into a proper definition, the first thing we need is a theoretically grounded method for determining what exactly constitutes a 'problem'. This is where Darwinian thinking makes its mark. The pivotal notion is that of an *adaptive problem*. Adaptive problems are problems that are specifiable in terms of evolutionary selection pressures, i.e., recurring environmental conditions that affect, or have affected, the reproductive success of individual organisms. So the evolutionary-psychological strategy is to require that the problems that delineate domains be adaptive problems. At a first pass, then, a domain may be a single adaptive problem or a set of adaptive problems. But this definition is immediately in need of attention. It is clearly important to find a way of ensuring that not any old, arbitrarily concocted, rag-bag of adaptive problems will count as a genuine domain. What one wants to say is that only a set of *suitably related* adaptive problems will do, which means saying something useful about what it means for a batch of adaptive problems to be suitably related. For the present we shall place this matter on hold. When we return to it later, we shall discover that it constitutes a pile a trouble.

Of course, there are all sorts of adaptive problems that selection has solved with means other than psychological mechanisms (e.g., we saw earlier that the adaptive problem of surviving extreme cold was solved, in the case of the polar bear, by selecting for a thick coat of fur). So, on the way of talking that we are about to import from evolutionary psychology, there exist non-psychological domains. But if psychological domains turn out to be a subset of domains in general, then we now need a way of drawing the distinction between the non-psychological and the psychological instances. Once one plugs in the fact that most evolutionary psychologists assume an information-processing theory of mind, the obvious move here presents itself. To be distinctively psychological, a domain must be defined by an adaptive problem (or a set of suitably related adaptive problems), posed in the creature's EEA, *for which an information-processing (as opposed to, say, a brutely anatomical) style of solution was appropriate.* For humans, such problems include how to throw in order to injure or kill one's dinner or enemy[8], how to select a mate[9], how to speak and understand language[10], how to engage in and reason about social exchange[11], and how to explain and predict each other's behaviour[12].

Having explicitly introduced the idea of a psychological domain, we shall, from now on, revert to speaking simply of domains. It should be clear enough in context whether we are concerned with psychological domains, non-psychological domains, or the generic notion of a domain that covers both sub-classes.

III Information and Mechanisms

If domains are adaptive problems, then a feature of our cognitive architecture is maximally *domain-specific* just when that feature is dedicated to solving one particular adaptive problem (or one specific set of suitably related adaptive problems). On the other hand, a feature of our cognitive architecture is maximally *domain-general* just when that feature can contribute to the solution of any adaptive problem whatsoever. These definitions describe the end-points of a spectrum of imaginable cases. In many ways, in fact, it makes most sense to think of a system as being relatively domain-specific or relatively domain-general, compared with some other point on the spectrum. A desk-top computer is more domain-general than a pocket calculator, for example, which in turn is more domain-general than an abacus.

So what sort of inner elements might domain-specific or domain-general cognitive features be? As we have already seen, most evolutionary psychologists assume an information-processing theory of mind, such that the adaptationist logic that we have characterised is used to constrain and inform theorising within that information-processing approach. As commonly understood, any information-processing psychology, Darwinian or otherwise, will buy into a distinction between data (bodies of information) and algorithms (processing mechanisms). Thus there are two ways in which a particular architectural feature might be domain-specific: it might be (a) a domain-specific body of information or (b) a domain-specific mechanism. And there are two ways in which a particular architectural feature might be domain-general: it might be (c) a domain-general body of information or (d) a domain-general mechanism.

It is important to note right away that whilst we are not alone in bringing the data/algorithm distinction to the fore in discussions of domain specificity and evolutionary thinking[13], the evolutionary psychologists do not themselves tend to make systematic use of that distinction in order to ask questions about the domain specificity or domain generality of information and mechanisms separately. There is often (although see later) simply an assumption at work that domain-specific mechanisms go hand-in-hand with domain-specific information, and that domain-general mechanisms go hand-in-hand with domain-general information. Thus consider the following, entirely typical quotations:

> [The] human psychological architecture contains many evolved mechanisms that are specialized for solving evolutionarily long-enduring adaptive problems and ... these mechanisms have content-specialized representational formats, procedures, cues, and so on.[14]

> To describe a system as domain-general or content-independent is to say not what it is but only what it lacks: It lacks any specific a priori knowledge about the recurrent structure of particular situations or problem-domains, either in declarative or procedural form, that might guide the system to a solution quickly. It lacks procedures that are specialized to detect and deal with particular kinds of problems, situations, relationships, or contents in ways that differ from any other kind of problems, situation, relationship, or content.[15]

In order to view the full conceptual landscape here, we are going to have to proceed beyond this undifferentiated treatment of data and algorithm, and put some words into the collective evolutionary-psychological mouth. We can begin this process by wielding the conceptual distinction between information and mechanism in order to define a matrix of in-principle possibilities for the design of the human cognitive architecture.

		information	
		domain-general	**domain-specific**
mechanisms	**domain-general**	1. domain-general mechanisms coupled with domain-general information	2. domain-general mechanisms coupled with domain-specific information
	domain-specific	3. domain-specific mechanisms coupled with domain-general information	4. domain-specific mechanisms coupled with domain-specific information

At this point we need to tread carefully, because cognitive psychologists in general (and evolutionary psychologists in particular) employ a range of terminology in this area. For example, as well as the distinction between domain-specific and domain-general architectural features, a contrast is often drawn between *special-purpose* and *general-purpose* architectural features, as well as between *content-dependent (or content-specialised)* and *content-independent (or content-general)* architectural features. It would be an under-statement to say that there is no universal agreement in the psychological literature (evolutionary or otherwise) on exactly how to draw these distinctions, or of how, or even of whether, they line up with each other in systematic ways. Through our eyes, however, the terrain looks like this.

Something is *special-purpose* if it has a particular job to do. Pens are specialised for writing, coffee machines for making coffee, and kangaroo pouches for the carrying of joeys. On the other hand, something is *general-purpose* to the extent that it has many jobs that it can do. For example, there is a piece of folklore about the ingenuity of New Zealanders that involves the countless uses to which no. 8 fencing wire can be put. Although the fencing wire was originally designed solely for use in the construction of fences, in the context of the folklore, the wire is general-purpose. Since information surely cannot be described as 'having a job to do' (as opposed to 'being used by a task-performing mechanism'), in anything other than an artificially contrived sense, it seems natural to confine the terms 'special-purpose' and 'general-purpose' to considerations of mechanism. In addition, we can exploit the thought that 'having a job to do' is equivalent to 'solving a problem', and so (via the idea that domains are to be delineated in terms of adaptive problems) identify (i) the notion of a special-purpose mechanism with that of a domain-specific mechanism, and (ii) the notion of a general-purpose mechanism with that of a domain-general mechanism.

In the present context, of course, the mechanisms that are of particular interest to us will be drawn from the class of information-processing mechanisms. And it is here that the contrast between content dependence and content independence comes to the fore, since it allows us to explicate the particular ways in which specifically information-processing mechanisms might be either general-purpose or special-purpose. We can describe as *content-independent* those information-processing mechanisms that are able to operate upon representations regardless of the content of those representations; that is, to put it roughly, it does not matter what the information is about, the processor is able to deal with it. For example, a memory mechanism might store information with many different types of content. Content-independent information-processing mechanisms will be general-purpose (domain-general) in nature. Correspondingly, we can describe as *content-dependent* those information-processing mechanisms for which it does matter what its informational inputs are about; that is, content-dependent subsystems are only able to operate with a specific and limited set of information types or content. Content-dependent information-processing mechanisms will be special-purpose (domain-specific) in nature.

To avoid confusion, we should stress that saying that some mechanism is content-dependent is not the same thing as saying that it necessarily deals in domain-specific information. Similarly, saying that some mechanism is content-independent is not the same thing as saying that it necessarily deals in domain-general information. This is crucial, since we need to ensure that there remains conceptual room for the idea that a general-purpose (domain-general, content-independent) information-processing mechanism might perform its task by accessing and manipulating domain-specific information (possibility 2 in our matrix), and for the idea that a special-purpose (domain-specific, content-dependent) information-processing mechanism might perform its task by accessing and manipulating domain-general information (possibility 3). We can satisfy ourselves that this space remains intact by noting three things. First, where we are talking about content dependence and content independence, these content-related properties of interest are properties of some mechanism. By contrast, where we are talking about the domain specificity or domain generality of some body of information, we are, of course, specifying a property of that information. This indicates that the content-dependent/content-independent and the domain-specific/domain-general distinctions are conceptually distinct. Secondly, if there is such a thing as domain-general information, then that information will have a certain content. But that means that some target mechanism might be keyed into that content, and so will be content-dependent. Finally, since a content-independent mechanism (by definition) doesn't care what the information in which it deals is about, that mechanism might access and manipulate various bodies information, each of which is specific to some particular domain.

In the interests of keeping things as clear as possible, let's now replay our matrix, using the terms 'special-purpose' and 'general-purpose' to replace 'domain-specific' and 'domain-general' in our characterisation of mechanisms. (Since some of the discussion to come draws lessons from domains other than psychological ones, the notions of special-purpose and general-purpose are, for our purposes, preferable to those of content-dependent and content-independent, since the latter pair is restricted to the information-processing, and thus the psychological, cases.)

information

		domain-general	**domain-specific**
mechanisms	**general-purpose**	1. general-purpose mechanisms coupled with domain-general information	2. general-purpose mechanisms coupled with domain-specific information
	special-purpose	3. special-purpose mechanisms coupled with domain-general information	4. special-purpose mechanisms coupled with domain-specific information

So which location in this matrix describes the design of the human cognitive architecture? Let's immediately narrow down the possibilities. As evolutionary psychologists often observe, the idea that cognition might be adaptive given only domain-general information is pretty much a non-starter. However, the point needs to be made with care. Tooby and Cosmides pour scorn on the thought that perfectly general information — information that would apply to every situation in all possible worlds — might have genuine adaptive utility.[16] But a compelling case against the adaptive utility of domain-general information cannot be made on the basis of such an extreme take on the idea of domain generality. The first thing to say here is that the evolutionary strategy (as we have characterised it) is to define domains by reference to the adaptive problems confronting a given species. Thus any notion of perfectly domain-general information needs to be defined in a similarly species-relative way, as information that is about, or relevant to, all the adaptive problems facing the species in question. But now even this idea seems too extreme, since it seems likely that there simply is no item or body of information that really is relevant to *all* the adaptive problems faced by a given species. So, to avoid charges of straw opponents, any interesting argument against the adaptive utility of domain-general information would have to be levelled at some notion of *relatively* domain-general information, information that is relevant to some large number of (species-related) adaptive problems.

Having straightened that out, however, the prosecution-case remains strong. For even this revised notion of domain-general information, as species-specific and relativistic as it is, falls prey to the following principle: as one abstracts progressively further away from contextual detail, it becomes increasingly hard to see how the style of information that emerges could ever figure importantly in satisfying explanations of context-sensitive adaptive intelligence. Thus consider the relatively domain-general item of information that two physical objects cannot take up the same portion of space at the same time. This piece of information might, we suppose, help directly to solve the odd adaptive problem. Moreover, according to some philosophical or cognitive-scientific accounts[17], it might even be granted the status of a kind of tacit background assumption that is made in large numbers of action-related contexts. However, it seems implausible in the extreme that possessing the knowledge that two physical objects cannot take up the same portion of space at the same time is in any way decisive for a creature to achieve adaptive success at, say, foraging, hunting, escaping from predators, or mate-finding. It seems rather that the information that does the

crucial [illegible]

> perform all physiological functions (pump blood, digest food, nourish an embryo, etc.) or that some sort of general-purpose kitchen device could perform all food processing tasks (broil, freeze, chop, etc.). There is no such thing as a 'general problem solver' because there is no such thing as a general problem.[19]

The fan of general-purpose mechanisms should, we think, simply concede that the notion of a *universal* problem solver simply doesn't make sense (cf. the equally troublesome concept of perfectly domain-general information that we ruled out above). However, what surely does make sense is the idea of a relatively general problem solver, a mechanism that works successfully in a large number of domains. So, from now on, we shall understand the term 'general-purpose mechanism' to capture precisely that idea.

IV A Job for the Specialist?

With the clarifications of the previous section in place, the core interpretative question can now be posed. Do evolutionary psychologists expect our innate cognitive architecture to be a matter of special-purpose mechanisms coupled with domain-specific information, or of general-purpose mechanisms coupled with domain-specific information? There is much evidence in favour of the former. For example, Tooby and Cosmides claim that:

> the human mind consists of a set of evolved information-processing mechanisms … many of these mechanisms are functionally specialized to produce behavior that solves particular adaptive problems … to be functionally specialized, many of these mechanisms must be richly structured in a content-specific way.[20]

And in a passage, the bulk of which we have reproduced once already, they state that:

> the human psychological architecture contains many evolved mechanisms that are specialized for solving evolutionarily long-enduring adaptive problems and ... these mechanisms have content-specialized representational formats, procedures, cues, and so on. … [they are] richly content-sensitive evolved mechanisms.[21]

How does such a view become plausible? There are certainly experimental studies which suggest that special-purpose mechanisms are indeed in operation in certain domains. Here is the flagship example: The Wason selection task requires subjects to detect violations of conditional rules, for which logically correct answers require the correct application of modus ponens and modus tollens. In the crucial experiments, the rules are presented either in abstract form (using letters and numbers), or with content that reflects a familiar context. It is a widely accepted interpretation of the robust results from studies of this task that subjects perform much better on the problem whose content reflects a familiar context, and is not presented merely in abstract form.[22] But this interpretation is problematic, because some familiar examples seem to cause as much difficulty as those with abstract symbols. In a series of studies, Cosmides, Tooby, and others have shown that subjects are good at the Wason selection task, if the example requires them to enforce a social contract.[23] Spotting cheaters in social exchanges (those who accept a benefit without paying the required cost) is an ability which would have an obvious selective advantage in the human EEA, where many of the most serious adaptive challenges took place during social interactions with conspecifics. So far this looks like another good argument for the presence of domain-specific information (about social exchange) which, for all we know, may be manipulated by some sort of general-purpose reasoning algorithm. What tells against this interpretation, to some extent anyway, are experiments involving the 'switched social contract' task.[24] In this task, the contractual information is presented in such a form that, if subjects apply our paradigm cases of general-purpose rules, modus ponens and modus tollens, they would give logically correct answers that violate the cost-benefit structure of social contracts. But if subjects apply rules of inference specialised for cheater detection, then they would give answers that fit the cost-benefit structure of social contracts, but which are logically incorrect. Overwhelmingly, subjects give answers consistent with the latter prediction, which indicates that special-purpose rules may well be involved (although it presumably remains possible that some presently unknown general-purpose rules other than those of the propositional calculus might be in operation).

By suggesting that special-purpose mechanisms are operative in particular contexts, experimental investigations such as the cheater detection studies are, of course, compelling and important. But one might wonder whether there are any good, in-principle evolutionary arguments for the claim that our cognitive architecture *must* consist in large numbers of special-purpose mechanisms rather than some small number of general-purpose mechanisms.

According to Shapiro and Epstein, the answer is 'no'.[25] They suggest that general-purpose mechanisms could well have evolved via what is, in effect, a process of *exaptation*. In Gould and Vrba's original explication of this concept, there are, strictly speaking, two varieties of exaptation.[26] In the first, an established biological feature that has no evolutionary function is later selected for by evolution to perform some adaptive task. In the second, an existing biological feature that has been selected for by evolution, and thus already has an evolutionary function, is later co-opted by selection to perform some new function. It is clear that to get the link between exaptation and the evolution of general-purpose mechanisms off the ground, one needs to focus on the second of these varieties of exaptation, and also add the rider that the previously evolved function must be maintained in addition to the newly acquired function; but with those qualifications in place, the following argument does have some prima facie appeal:

> Why develop new eyes with which to find kin when one can use old eyes that have served well in the detection of predators? Likewise, why develop a new process for recognising kin when one has already in place a process that allows avoidance of kin when searching for mates? We should not be surprised if natural selection has recruited extant cognitive capacities for new purposes rather than going to the trouble of developing new capacities every time a novel problem comes along.[27]

This argument trades on an historical conception of domain generality, according to which any mechanism that is selected as the solution to a particular adaptive problem counts initially as a special-purpose mechanism, but becomes increasingly general-purpose in character if and as it is exapted to perform other tasks. Of course, 'having been exapted' is not, on its own at least, a sufficient condition for a mechanism to count as general-purpose, since a mechanism that is selected for one task and then exapted for one other is surely not yet general-purpose, in any interesting sense. So the idea must be that there is some (no doubt vague) point beyond which a repeatedly exapted mechanism will uncontroversially receive its general-purpose spurs. (Once one takes this fact on board, Shapiro and Epstein's once-exapted kin recognition mechanism still looks suspiciously special-purpose, relatively speaking, but we can let that pass.)

The problem is that, even with this tidied-up historical conception of domain generality in play, Shapiro and Epstein's argument misses its target, because when evolutionary psychologists use the term 'general-purpose mechanism' (or its equivalent), they simply do not mean 'suitably exapted mechanism'. They are appealing instead to another sense of general-purpose according to which any suitably (and so repeatedly) exapted mechanism that we might come across was, in fact, always general-purpose in character. On this view, the process of exaptation merely permits an intrinsically general-purpose mechanism to fulfil its potential. For this idea to make sense, what is needed is a strategy for characterising what it means to be general-purpose that does not appeal to any actual history of evolutionary use. Such a strategy is precisely what we find in evolutionary psychology. Evolutionary psychologists typically illustrate the notion of a general-purpose mechanism by reference either to (a) what Tooby and Cosmides call the Standard Social Science Model (SSSM)[28], or to (b) certain reasoning/learning architectures already established within cognitive science that are taken to be consonant with SSSM. Roughly, SSSM is the view that human beings are born with no innate psychological structure except for a small number of widely applicable learning devices that provide a capacity for absorbing culturally transmitted information and skills. Cognitive-scientific systems that evolutionary psychologists tend to cite as paradigmatic general-purpose mechanisms include the General Problem Solver[29] and mainstream connectionist networks. These are systems that can be applied to a wide class of information-processing problems, assuming that the problem-space in question can be defined within the particular representational format that is appropriate for the system being deployed.[30] So when evolutionary psychologists speak of general-purpose mechanisms, what they typically mean is the sort of intrinsically general-purpose mechanism that might be identified using (a) and (b). And what they typically say is that general-purpose mechanisms of this character are evolutionarily implausible *in principle* (and thus could not provide the raw material for exaptation).

At this point we need to make explicit certain explanatory constraints that are already implicitly at work in the arguments presented by evolutionary psychologists. By taking an evolutionary approach, the psychologist inherits two requirements on any proposed hypothesis about the design of our cognitive architecture.[31] The first is what we shall call the *Minimal Requirement*. An estimated 99.9% of all species that have existed on Earth are now extinct[32], yet we as a species still exist. Thus we can confidently state that whatever design our cognitive architecture has, it must have been successful in solving the adaptive information-processing problems that our hunter-gatherer ancestors faced. From this it follows straightforwardly that for some postulated design even to be a reasonable candidate for the right story here, that design must, in principle, be able to solve all and each of those adaptive problems. The second requirement is what we shall call the *Darwinian Requirement*. The logic of Darwinism dictates that if a number of solutions to an adaptive problem (in this context, designs for cognitive systems) are present in the population, then selection will favour the superior performers. (What counts as superior will depend on the particular evolutionary scenario, but will typically involve criteria such as speed, efficiency, and reliability.) In the present context, the key implication of the Darwinian Requirement is that any information-processing system that is proposed as part of our evolved psychological architecture must have been superior to any competing system that was actually present in the human EEA (and which therefore could, potentially, have been the object of selection).

Evolutionary psychologists present a range of arguments designed to establish that general-purpose mechanisms will fail to meet the Minimal Requirement.[33] We shall concentrate on what is, we think, the strongest argument on offer, namely that there are many adaptive problems that could not, *in principle*, be solved by any general-purpose mechanism, because such a mechanism would be paralysed by a vicious combinatorial explosion.[34]

At the very core of the argument from combinatorial explosion is what, in artificial intelligence (AI), is called the frame problem.[35] As characterised by Fodor the frame problem is ‘the problem of putting a “frame” around the set of beliefs that may need to be revised in the light of specified newly available information’.[36] Put another way, the difficulty is how to retrieve and revise just those beliefs or other representational states that are relevant in some particular context of action. Since a general-purpose reasoning system will have access to an entire system of beliefs etc., there is a real danger here of a kind of computational paralysis due to the enormity of the processing-load. It might seem that the obvious move is to deploy heuristics that decide which of a system’s representations are relevant in a particular scenario. Unfortunately, this kite won’t fly. The processing mechanisms would still face the problem of accessing just those relevancy heuristics that are relevant in the current context. And it is not merely that some sort of combinatorial explosion or infinite regress threatens here (which it does — how do we decide which relevancy heuristic is, in the present context, relevant?). The deeper concern is that we have no really good idea of how a computational process of relevance-based update might work. As Horgan and Tienson point out, the situation cannot be that the system first retrieves an inner structure (an item of information or a heuristic), and then decides whether or not it is relevant, as that would take us back to square one.[37] But then how can the system assign relevance until the structure has been retrieved?

Fodor argues that cognitive science has made progress in understanding input systems (e.g., perceptual systems) only because those systems are 'informationally encapsulated': they are systems for which the class of relevant information is restricted. In contrast, he argues, central cognitive systems, whose primary job is the fixation of beliefs, are unencapsulated, such that (i) information from anywhere in the agent's entire stock of beliefs is potentially relevant to fixing any specific belief, and (ii) the confirmation of an individual belief is relative to the structure of the whole belief system.[38] On Fodor's analysis the frame problem is symptomatic of this situation. But now a way to repel the frame problem seems to present itself. If central cognitive systems were like Fodorian input systems, and were thus informationally encapsulated rather than informationally unencapsulated, then they wouldn't succumb to the frame problem. And that, in part, is the picture of central cognitive systems that evolutionary psychologists have.

We say 'in part' because there is more to be said. It seems clear that the kind of system that Fodor identifies as unencapsulated must, by its very nature, contain general-purpose (i.e., content-independent) psychological mechanisms, so that any body of information present in the cognitive architecture as a whole might potentially be accessed and processed. However, it is far less clear that a commitment to informational encapsulation forces us to jump one way or the other on the issue of whether the psychological mechanisms concerned will be special-purpose or general-purpose. Any special-purpose (content-dependent) psychological mechanism will be encapsulated; but so, it seems, will an intrinsically general-purpose psychological mechanism that, through the opportunistic process of design by selection, has been allowed access only to a restricted body of information. If this analysis is right, then adopting the thought that central cognitive systems will need to be informationally encapsulated in order to defeat the frame problem (and thus, in this context, to satisfy the Minimal Requirement) does not yet decide the question of whether the mechanisms involved will need to be special-purpose or general-purpose, or, indeed, whether that is a question to which an *in-principle* answer can even be given. From what we have seen so far, the evolutionary psychologist will claim that distinctively special-purpose mechanisms, coupled with suitably organised bodies of domain-specific information, will be required. But if informational encapsulation really is the key here, then why, in principle, couldn't mechanisms which are intrinsically general-purpose in character, but which (as a matter of evolutionary fact) have access only to the appropriate body of domain-specific information, be capable of meeting the challenge?

As it happens, it is possible to be more positive. There already exists work in AI which suggests that if the designer can pre-define what is relevant in a particular scenario, then a system featuring a combination of informational encapsulation, general-purpose mechanisms, and domain-specific information will not be completely incapacitated by the frame problem. Consider Shakey[39], a famous AI-robot that inhabited an environment consisting of rooms populated by static blocks and wedges, in which it confronted a problem-domain defined by tasks of the form, 'move block A to some specified new location'. (We might treat this problem-domain as analogous to a simple adaptive domain in the natural world, although of course it is structured by the human designer rather than selection.) Shakey used visual input plus domain-specific 'innate' knowledge about its constrained world of rooms and static blocks, in order to build a model of that world in a set of first-order predicate calculus representations. These representations were then delivered to a central planning system based on the General Problem Solver (one of our paradigm cases of a general-purpose mechanism), which proceeded to determine a sequence of appropriate actions. The fact is that Shakey actually worked (albeit very slowly, more on which below); and so have many related AI-systems. It is true that these recorded successes might, in many ways, be attributable to the extreme simplicity and static nature of the relevant operating environments.[40] Nevertheless, such modest victories surely count as empirical evidence that general-purpose mechanisms are not always computationally paralysed by the frame problem, and thus that the Minimal Requirement might, at least sometimes, be met by general-purpose mechanisms using domain-specific information.

But now what about the Darwinian Requirement? For general-purpose mechanisms to meet this constraint, it would have to be the case that they are not, in general, outperformed in the efficiency stakes by special-purpose mechanisms. Tooby and Cosmides claim that special-purpose mechanisms will be the more efficient. As they put the point, 'domain-general, content-independent mechanisms are inefficient, handicapped, or inert compared to systems that also include specialized techniques for solving particular families of adaptive problems'.[41] In response, Samuels complains (a) that no clear argument is provided for this specific claim, as opposed to the more general claim that systems with domain-specific features of some kind will be more efficient than those without, and (b) that 'it is far from clear that anyone knows how such an argument would go'.[42] Samuels may be right; but perhaps something other than an argument will take us at least part of the way here. Some quite recent empirical work in AI-oriented robotics suggests that robots featuring architectures in which special-purpose mechanisms are employed alongside domain-specific information are typically faster and more robust than robots in which general-purpose mechanisms are used. For example, in the *subsumption architecture*, as pioneered by Brooks and his colleagues[43], individual behaviour-producing mechanisms called 'layers' are designed to be individually capable of (and to be generally responsible for) connecting the robot's sensing and motor activity, in order to achieve some particular, ecologically relevant behaviour. Starting with layers that achieve simpler behaviours (such as 'avoid obstacles' and 'explore'), these special-purpose mechanisms are added, one at a time, to a debugged, working robot, so that overall behavioural competence increases incrementally. In a 'pure' subsumption architecture, no detailed messages are passed between the parallel-running multiple layers. Indeed, each layer is completely oblivious of the layers above it, although it can suppress or *subsume* the activity of the layers below it. In practice, however, the layers are often only semi-independent.

Although, from a certain perspective, the tasks performed by subsumption-based robots remain quite simple (e.g., navigating to a light source without bumping into obstacles[44], or collecting soft-drink cans from around the MIT labs[45]), the approach has scored notable successes in achieving real-time behaviour in everyday, unconstrained, dynamic environments. It is thus a potentially important observation that whilst Shakey was not completely incapacitated by the frame problem, it was painfully slow, and its environment was (unlike the environments of most animals) essentially static. So to the extent that it is permissible, in the present context, to generalise from the achievements of these relatively simple robots, there is some empirical support for the following claim: even though general-purpose mechanisms might have been able to solve the adaptive problems faced in the human EEA, they will have been less efficient and robust than any competing special-purpose mechanisms that were available to evolution at the time, and thus will have been selected against.

V Contradictions, Crabs, and a Crisis

So far we have characterised evolutionary psychology as being committed to the view that the human cognitive architecture is essentially a collection of special-purpose mechanisms operating on domain-specific information. However, it is time to complicate matters, and record the fact that statements of a less dogmatic kind can often be found nestled in the evolutionary-psychological literature, statements which suggest that special-purpose and general-purpose mechanisms might peacefully co-exist, and which at least allow for the possibility that adaptive success might sometimes be explained by general-purpose mechanisms operating on domain-specific information. For example, in a recent paper, Tooby and Cosmides remind us that, in previous work, they have

> made some serious and extended arguments about why many more cognitive mechanisms [than those underpinning vision and language] will turn out to include domain-specific features alongside whatever domain-general principles of operation they have.[46]

And in an earlier paper, Cosmides and Tooby summarise their view by saying that our minds consist primarily of a 'constellation of specialized mechanisms that have domain-specific procedures, or operate over domain-specific representations, or both'.[47] So it seems that evolutionary psychologists do (sometimes at least) treat the human cognitive architecture as containing a mix of general-purpose and special-purpose mechanisms. But one surely has to wonder whether the evolutionary psychologists we have been considering have the warrant to occupy such a position. After all, we have seen them launch *in-principle* arguments against the evolutionary plausibility of general-purpose mechanisms, so it is difficult to avoid the thought that they are in real danger of simply contradicting themselves here by conceding the evolutionary plausibility of such mechanisms. At first sight, however, we are rather better placed than these evolutionary psychologists to take the mixed mechanism option seriously. We found the best in-principle argument against the evolutionary plausibility of general-purpose mechanisms to be somewhat less then compelling, and ended up by confining the case 'merely' to empirical evidence from research in AI which suggests that general-purpose mechanisms will not meet the Darwinian

Requirement. Moreover, perhaps we have been too hasty in suggesting (implicitly) that the results of that empirical research can be generalised across all evolutionary niches. It is time, then, to explore the possibility that there exist styles of evolutionary niche which are distinctive in that they provide conditions in which the Darwinian Requirement may be met by recognisably general-purpose mechanisms. The phenomenon of diffuse co-evolution[48] generates a prime example of the kind of conditions we have in mind.

Diffuse co-evolution occurs when a group of traits in one population (which may contain several species) drives evolutionary change in a trait in a second population, via the imposition of selection pressures that act in parallel. For example, the hard shells of many crustaceans are an evolutionary response to a range of shell-breaking mechanisms deployed by a number of different predatory species. If each individual shell-breaking strategy (the deployment of pincers, crushing techniques, and so on) can be treated as a distinct adaptive domain (more on this soon), one may conclude that the crustacean shell is a relatively general-purpose response that works in a variety of different adaptive domains.

As a first pass, then, one might think that general-purpose mechanisms may result from the operation of multiple simultaneous selection pressures. But this idea is in immediate need of tightening up. Just about every organism faces multiple, simultaneous selection pressures, but only a (probably quite small) subset of these would likely have led to the development of general-purpose mechanisms. To introduce a selective force in favour of a general-purpose mechanism, the selection pressures concerned would have to be *related* in some *systematic* way. It's not easy to give a principled definition of 'systematically related', but its meaning is at least intuitively clear. On the one hand, it is rather unlikely that a mechanism would have evolved that was capable of, say, protecting the crab from attack *and* controlling its leg muscles or pumping blood around its body. But on the other hand, it seems clear that (a) attacks by one type of weapon can count as one selection pressure, and (b) attacks by different types of weapons can count as multiple selection pressures which are systematically related.

The proposal we wish to extract from the example of diffuse co-evolution has now emerged: where multiple, simultaneous, systematically related selection pressures are in operation, the adaptive response may well be a general-purpose mechanism. (Although diffuse co-evolution provides a paradigmatic case of this process, the *co*-evolutionary nature of the scenario is not strictly crucial to out point.)[49] Our further suggestion (which owes much to a discussion of diffuse co-evolution by Bullock[50]) is that there might well be analogous selective forces in favour of general-purpose cognitive mechanisms. These forces may occur, we suggest, just where an organism confronts multiple, simultaneous, systematically related selection pressures that require an information-processing style of solution. Note that the general-purpose character of these mechanisms would not approach that of the mechanisms evolutionary psychologists typically rally against, that is, mechanisms identified by reference either to the SSSM or to reasoning/learning architectures such as the General Problem Solver and mainstream connectionist networks (see section IV). Perhaps a good term for what might be achieved is *cross-domain adaptation*.[51]

As mentioned above, one constraint on the evolution of such relatively general-purpose, cross-domain adaptive mechanisms is that they must meet the Darwinian Requirement. But, at first sight, this seems to introduce a difficulty with the present proposal, namely that in any one of its domains of application, we might expect a cross-domain adaptive mechanism to have been out-performed by any suitably aimed special-purpose mechanism that was present in the relevant EEA. With respect to that particular domain, therefore, the cross-domain adaptive mechanism would fail to meet the Darwinian Requirement. In fact, what is at fault here is not the argument for cross-domain adaptation, but the idea that the Darwinian Requirement must be met for each individual selection pressure (e.g., each predatory strategy). After all, while pincer cutting machinery may out-score a hard shell in the context of attack by pincers, it will do much worse than the hard shell against some crushing strategy. It is no good running rings round the pincer monster if you get flattened by the crushing beast five minutes later. Thus it is surely plausible that where the relevant selection pressures are multiple, simultaneous, and systematically related, the selective advantages bestowed by relative efficiency, relative reliability, and so on, will accrue at the level of the whole set of selection pressures.

Unfortunately we have just stepped on to a very slippery slope indeed. We have suggested that if a group of multiple, simultaneous selection pressures (such as the different predatory strategies facing the crab) are systematically related, then an accurate account of what structures and traits have evolved, and why, would seem to require that selective advantage be specified at the level of the entire set of selection pressures. But if that is so, then why isn't the set of selection pressures in question itself just a bona fide adaptive domain, making the mechanism that evolved in response (the crab's shell, for example) special-purpose? In this context, recall our earlier observation that a domain may be a single adaptive problem, or it may be a set of *suitably related* adaptive problems. This observation was accompanied by the warning that once we sorted out what 'suitably related' might mean, there would be trouble. We are now in a position to see why. It is surely plausible that at least one way in which a collection of adaptive problems might be suitably related to each other, so as to constitute a domain, is by meeting the dual requirements of operating in parallel and being systematically related. But that means that the set of different predatory strategies facing the crab constitutes an adaptive domain, and thus that the crab's shell is a special-purpose mechanism. Indeed one might conceivably go further: if certain multiple simultaneous selection pressures are indeed systematically related in the way we have described, why doesn't that make their joint operation precisely a kind of single, although higher-order and complex, selection pressure in its own right, one to which the crab's shell is a special-purpose response?

At this juncture you might well be thinking, 'So much the worse for Wheeler and Atkinson's attempt to identify a style of evolutionary niche in which the Darwinian Requirement may be met by a recognisably general-purpose style of mechanism'. If that were the extent of the trouble here, then, indeed, not much would be at stake. But now consider this. It is certainly beginning to look as if the specification of what counts as a domain, and therefore any subsequent judgement of domain specificity or domain generality, will be relative to our description of the evolutionary scenario. In one sense, this may come as no great surprise. As any survey of the appropriate academic literature will indicate, it is common for cognitive and brain scientists to be somewhat promiscuous in their deployment of the term 'domain', in line, we suppose, with its pre-theoretical, everyday usage. Perceptual and cognitive tasks at various levels of description are said to be domains (and therefore cognitive/neural mechanisms at various levels of description are said to be domain-specific or special-purpose). For example, language, vision and audition are often delineated as domains. But so are more fine-grained tasks, such as word recognition, object recognition and face recognition. And sometimes even finer-grained psychological problems seem to count, problems such as phoneme segmentation and line orientation. What our investigation here suggests is that, from an evolutionary perspective, the promiscuity present in the scientific literature may well be warranted; which is just fine, unless, that is, one feels that there ought to be room for a genuine debate over whether the features of our evolved psychological architecture should be categorised as domain-specific or as domain-general. The fact that such judgements are, if we are right, relative to some adopted level of description suggests that the best one can say in this debate is that once a level of description has been agreed, a particular mechanism, or a particular body of information, will turn out to be in one category or the other. Nevertheless, there always remains the possibility of switching to another level of description, at which we may well get an entirely different result.

The consequence just highlighted might not seem so bad, if one concentrates on cases where there exist plausible criteria for deciding, in a principled fashion, between competing levels of description. But what do we say about cases where no such criteria exist? Think again about the crab's shell. On the face of it, there doesn't seem to be anything to choose between our two competing interpretations of this evolutionary scenario. The interpretation that finds the target mechanism to be special-purpose doesn't seem to be

This work grew out of discussions on evolutionary psychology between the two of us while we were both based in the Department of Experimental Psychology, University of Oxford. During this time, AA was partly supported by St. Edmund Hall, Oxford, and MW was supported by a Junior Research Fellowship at Christ Church, Oxford, with additional assistance from the McDonnell-Pew Centre for Cognitive Neuroscience. Our thanks to all these institutes for their support. AA also wishes to thank the Department of Psychology, University of Yorp0 50 36 0 Tm (i) T0 0 Tm (,) T6 50 0 0

Footnotes

[1] For example: L. Cosmides and J. Tooby, 'From Evolution to Behavior: Evolutionary Psychology as the Missing Link', in The Latest on the Best: Essays on Evolution and Optimality, ed. J. Dupre (Cambridge, MA: MIT Press, 1987), pp.227-306; L. Cosmides and J. Tooby, 'Beyond Intuition and Instinct Blindness: Toward an Evolutionary Rigorous Cognitive Science', Cognition, 50 (1994), 41-77; D. Symons, 'On the Use and Misuse of Darwinism in the Study of Human Behavior', in The Adapted Mind: Evolutionary Psychology and the Generation of Culture, ed. J.H. Barkow, L. Cosmides, and J. Tooby (New York: Oxford University Press, 1992), pp.137-159.

[2] See, e.g., J. Tooby and L. Cosmides, 'The Psychological Foundations of Culture', in The Adapted Mind: Evolutionary Psychology and the Generation of Culture, ed. J.H. Barkow, L. Cosmides, and J. Tooby (New York: Oxford University Press, 1992), pp.19-136; Symons, op. cit. note 1.

[3] See e.g., R. Samuels, 'Evolutionary Psychology and the Massive Modularity Hypothesis', British Journal for the Philosophy of Science, 49 (1998), 575-602; L. Shapiro and W. Epstein, 'Evolutionary Theory Meets Cognitive Psychology: A More Selective Perspective', Mind and Language, 13 (1998), 171-194.

[4] In what follows we shall often write as if evolutionary psychologists speak with one voice (roughly, that of Cosmides and Tooby). This is, of course, a distortion of what is a thriving research endeavour driven forward by internal disputes and differences of approach. At the very least, however, the position we describe is widely held within the field, and is championed vigorously by some of the discipline's leading theorists.

[5] L.A. Hirschfeld and S.A. Gelman, 'Toward a Topography of Mind: An Introduction to Domain Specificity', in Mapping the Mind: Domain Specificity in Cognition and Culture, ed. L.A. Hirschfeld and S.A. Gelman (Cambridge and New York: Cambridge University Press, 1994), pp.3-35.

[6] J.A. Fodor, The Modularity of Mind (Cambridge, MA: MIT Press, 1983), p.48.

[7] Op. cit. note 3, p.21.

[8] W. Calvin, The Throwing Madonna: Essays on the Brain (New York: McGraw-Hill, 1983).

[9] D.M. Buss, 'Mate Preference Mechanisms: Consequences for Partner Choice and Intrasexual Competition', in The Adapted Mind: Evolutionary Psychology and the Generation of Culture, ed. J.H. Barkow, L. Cosmides, and J. Tooby (New York: Oxford University Press, 1992), pp.249-266.

[10] S. Pinker and P. Bloom, 'Natural Language and Natural Selection', Behavioral and Brain Sciences, 13 (1990), 707-727.

[11] L. Cosmides and J. Tooby, 'Evolutionary Psychology and the Generation of Culture, Part II. Case Study: A Computational Theory of Social Exchange', Ethology and Sociobiology, 10 (1989), 51-97; L. Cosmides and J. Tooby, 'Cognitive Adaptations for E13 Tf 50 0 0 50 131 0 Tm ()Tj 50 0 0 50 55 0 Tm (ha)4 Tj 50 0 0 50 69

Age, ed. D. Michie (Edinburgh: Edinburgh University Press, 1979), pp.242-70.

[18] Although we have played along with a way of talking according to which certain general background assumptions, such as 'two physical objects cannot take up the same portion of space at the same time' and 'natural light always comes from above', are treated as items of information that the system possesses, we are not remotely convinced that this is the best way to conceptualise matters. If such knowledge is not internally represented in a form which makes it accessible to and/or manipulable by the relevant psychological mechanisms, but rather is tacitly embedded in the basic operating principles of those mechanisms, then it is far from clear that it should be regarded as *information* that the system possesses at all.

[19] Symons, op. cit. note 1, p.142.

[20] Tooby and Cosmides, op. cit. note 2, p.24.

[21] Op. cit., p.34.

[22] For example: R.J. Bracewell and S.E. Hidi, 'The Solution of an Inferential Problem as a Function of the Stimulus Materials', Quarterly Journal of Experimental Psychology, 26 (1974), 480-488; R.A. Griggs and J.R. Cox, 'The Elusive Thematic-Materials Effect in Wason's Selection Task', British Journal of Psychology, 73 (1982), 407-420; R.A. Griggs and J.R. Cox, 'The Effects of Problem Content and Negation on Wason's Selection Task', Quarterly Journal of Experimental Psychology, 35A (1983), 519-533; P. Pollard, 'The Effect of Thematic Content on the "Wason selection task" ', Current Psychological Research, 1 (1981), 21-29.

[23] For example: L. Cosmides, 'The Logic of Social Exchange: Has Natural Selection Shaped How Humans Reason? Studies with the Wason Selection Task', Cognition, 31 (1989), 187-276; Cosmides and Tooby, op. cit. note 11; G. Gigerenzer and K. Hug, 'Domain-Specific Reasoning: Social Contracts, Cheating, and Perspective Change', Cognition, 43 (1992), 127-171.

[24] Cosmides, op. cit.; see also Cosmides and Tooby (1992), op. cit. note 11.

[25] Op. cit. note 3.

[26] S.J. Gould and E.S. Vrba, 'Exaptation: A Missing Term in the Science of Form', Paleobiology, 1:8 (1982), 4-15.

[27] Shapiro and Epstein, op. cit. note 3, p.176.

[28] Op. cit. note 2.

[29] A. Newell and H.A. Simon, 'GPS — A Program tha

op. cit. note 2, pp.100-112.

[36] Op. cit. note 6, pp.112-13.

[37] T. Horgan and J. Tienson, 'A Nonclassical Framework for Cognitive Science', Synthese, 101 (1994), 305-45.

[38] Op. cit. note 6.

[39] N.J. Nilsson, 'Shakey the Robot', S.R.I. AI Centre Tech. Rept. no.323 (1984).

[40] For discussion, see R.A. Brooks, 'Intelligence Without Reason', in Proceedings of the Twelfth International Joint Conference on Artificial Intelligence San Mateo, California (Morgan Kauffman, 1991), pp. 569-95.

[41] Tooby and Cosmides, op. cit. note 2, p.111.

[42] Samuels, op. cit. note 3, p.588.

[43] See, e.g., Brooks, op. cit. note 40.

[44] N. Franceschini, J.M. Pichon and C. Blanes, 'From Insect Vision to Robot Vision', Philosophical Transactions of the Royal Society, B, 337 (1992), 283-94.

[45] J.H. Connell, 'A Colony Architecture for an Artificial Creature', MIT AI Lab Memo 1151 (1989).

[46] J. Tooby and L. Cosmides, 'Evolutionizing the Cognitive Sciences: A Reply to Shapiro and Epstein', Mind and Language, 13 (1998), 195-204, p.198.

[47] Cosmides and Tooby (1994), op. cit. note 1, p.94.

[48] D.H. Janzen, 'When is it Co-Evolution?' Evolution, 34 (1980), 611-612.

[49] In cases of exaptation (as described earlier in this paper), the multiple selection pressures at issue will certainly be systematically related, and, in some instances of the phenomenon, may even end up acting in parallel. However, by hypothesis, those selection pressures will be introduced in a sequential or additive fashion. At first, selection will generate a trait as a response to a single

adaptive problem. Given the Darwinian Requirement, we would typically expect the nature of the evolutionary response to be tailored in highly specific ways to the nature of the adaptive problem. Subsequently, selection may tinker with that already-established trait, in order to meet later adaptive challenges. But, in our view, it is to the extent that the selection pressures in question operate in parallel *from the beginning* of the relevant period of evolutionary design-history that we should expect to see some sort of intrinsically general-purpose mechanism as the outcome.

[50] S. Bullock, Evolutionary Simulation Models: On Their Character and Application to Problems Concerning the Evolution of Natural Signalling Systems. DPhil Thesis, University of Sussex (1998).

[51] Our thanks to Gary Brase for suggesting this alternative term.

[52] There are without doubt some close connections between the problem that we have aired here (roughly, that choosing between levels of evolutionary description may sometimes be arbitrary), and what Sterelny and Griffiths have identified recently as evolutionary psychology's 'grain problem'; see K. Sterelny and P.E. Griffiths, Sex and Death: An Introduction to Philosophy of Biology (Chicago: University of Chicago Press, 1999). (As these authors explain it, the grain problem is that it is very difficult, and in many cases perhaps impossible, to match distinct adaptive problems with distinct phenotypic features that have evolved to solve them.) In the end, it may even be that our two problems are in fact equivalent. If so, then what one line of thought in the present paper has achieved is the mapping out of a significantly different route to, and a significantly different elaboration of, that single problem.