King's Research Portal

DOI:

[10.1111/j.1468-0017.2012.01452.x](https://doi.org/10.1111/j.1468-0017.2012.01452.x)

Document Version

Publisher's PDF, also known as Version of record

[Link to publication record in King's Research Portal](#)

Citation for published version (APA):

Shea, N. (2012). Genetic Representation Explains the Cluster of Innateness-Related Properties. *Mind & Language*, 27(4), 466-493. <https://doi.org/10.1111/j.1468-0017.2012.01452.x>

Citing this paper

Please note that where the full-text provided on King's Research Portal is the Author Accepted Manuscript or Post-Print version this may differ from the final Published version. If citing, it is advised that you check and use the publisher's definitive version for pagination, volume/issue, and date of publication details. And where the final published version is provided on the Research Portal, if citing you are again advised to check the publisher's website for any subsequent corrections.

General rights

Copyright and moral rights for the publications made accessible in the Research Portal are retained by the authors and/or other copyright owners and it is a condition of accessing publications that users recognize and abide by the legal requirements associated with these rights.

- Users may download and print one copy of any publication from the Research Portal for the purpose of private study or research.
- You may not further distribute the material or use it for any profit-making activity or commercial gain
- You may freely distribute the URL identifying the publication in the Research Portal

Take down policy

If you believe that this document breaches copyright please contact librarypure@kcl.ac.uk providing details, and we will remove access to the work immediately and investigate your claim.

Genetic Representation Explains the Cluster of Innateness-Related Properties

NICHOLAS SHEA

Abstract: The concept of innateness is used to make inferences between various better-understood properties, like developmental canalization, evolutionary adaptation, heritability, species-typicality, and so on ('innateness-related properties'). This article uses a recently-developed account of the representational content carried by inheritance systems like the genome to explain why innateness-related properties cluster together, especially in non-human organisms. Although inferences between innateness-related properties are deductively invalid, and lead to false conclusions in many actual cases, where some aspect of a phenotypic trait develops in reliance on a genetic representation it will tend, better than chance, to have many of the innateness-related properties. The account also shows why inferences between innateness-related properties sometimes fail and argues that such inferences are especially misleading when applied to human psychology and behaviour because human psychological development is especially reliant on non-genetic inherited representations.

1. Introduction

Innateness claims have a long history in philosophy and have, post-Chomsky, become a central feature of psychology and cognitive science. Studies of what are taken to be innate psychological capacities have proven fruitful in areas such as object permanence, causal inference and basic arithmetical operations in children; object tracking, face recognition and visual perception in adults; and in studies of these and other mechanisms in other species. Yet when philosophers and others have attempted to pin down what it is for a trait to be innate they have found it hard to make coherent sense of the concept. In behavioural ecology the term has been rejected (Bateson, 1991) or superseded by the idea that some traits are programmed in the genes; in molecular developmental biology it is hardly used at all.

Various properties are associated with a trait's being innate, which do not in fact correlate very well in actual cases: its emergence in development being canalized, its being unlearned, its being an adaptation, and so on. The concept of innateness is relied on to underpin inferences between such properties, for example: feature T of psychology is unlearned, hence T is innate, hence T is an adaptation. Many different

The author would like to thank Tim Bayne, David Chalmers, Marian Dawkins, Cecelia Heyes, Russell Powell, Lizzie Schechter and Tobias Uller for discussion and comments on earlier drafts; audiences in Oxford and at the University of Nottingham for helpful discussion; and two anonymous referees for comments on a previous draft. This work was supported by: the Wellcome Trust (grant number 086041), the Oxford Martin School and the John Fell OUP Research Fund.

Address for correspondence: Faculty of Philosophy, 10 Merton Street, Oxford, OX1 4JJ, UK.
Email: nicholas.shea@philosophy.ox.ac.uk

properties feature in inferences of that form. Amongst all the conceptions associated with the concept by ordinary users or in particular sciences, the following nine are central and representative of the range of issues under discussion.¹ Following Mameli (2008),² I call these innateness-related properties or *i-properties*:

i-properties

- (i) Present at birth;
- (ii) Genetically-determined;
- (iii) Invariant: arises in development despite variation in the environment. Active version—developmentally canalized: the trait is the result of mechanisms that buffer its development against environmental variation³;
- (iv) Universal in the species (or typical of members of the species)⁴;
- (v) An evolutionary adaptation;
- (vi) Programmed in the genes;
- (vii) A phenotypic difference caused by a genetic difference, or a genetically heritable trait;
- (viii) Not learned / not the result of a psychological acquisition mechanism. Generalises: not the result of a mechanism of adaptive plasticity (caveat: nor produced by evolutionarily abnormal factors);
- (ix) An internally-organised module responsible for a system of behaviour.

This list represents the leading proposals for giving an analysis of the concept of innateness (i.e. spelling out its definition or meaning—which is not our project).⁵ The problems with that project are well known: no analysis suffices to capture all or most of the intuitive applications of the concept, not even the central ones.⁶

¹ The list is based on the seven different senses identified by Bateson (1991, p. 21) as in use in studies of animal behaviour, supplemented and adapted in the light of the subsequent literature including Griffiths and Gray, 1994; Cowie, 1999; Ariew, 1999; Samuels, 2002; Khalidi, 2002; Griffiths, 2002, 2009 and the comprehensive survey found in Mameli and Bateson, 2006—and indeed much of the literature on innateness cited below.

² Mameli and Bateson (2006, p. 178) coined the term ‘i-properties’ for candidates to be a scientific successor to the folk concept (which therefore need not in principle be restricted to conceptions actually associated with the concept by concept users). Mameli (2008, p. 745) defines ‘i-properties’ slightly differently, as properties that regularly feature in inferences to or from the claim that a trait is innate. The term is used in the latter sense in this article.

³ Note, for brevity ‘canalized’ is used as shorthand for canalization against environmental, as opposed to genetic, variation.

⁴ Species-typicality is supposed to be something stronger than a mere statistical tendency as in: ‘humans typically have brown eyes, but green and blue eyes are also quite frequent’. The term is used to mark the fact that even traits that are characteristic of a species usually admit of exceptions in pathological or unusual cases. Here, the shorthand ‘universal’ is used in that sense.

⁵ Stich, 1975; Ariew, 1999; Cowie, 1999; Khalidi, 2002 and Samuels, 2002 offer different analyses or theoretical reconstructions of the concept; Griffiths, 2002; Mameli and Bateson, 2006; Godfrey-Smith, 2007; Mameli, 2008 and Mameli and Bateson, 2011 exemplify pessimism about there being a good analysis. See Griffiths, 2009 for a review.

⁶ Griffiths, 2002; Mameli and Bateson, 2006; Mameli, 2008; Cowie, 2009.

Thus, many seemingly innate traits develop long after birth (e.g. male facial hair) and some learning occurs before birth (e.g. neonates learn the prosodic structure of their native language in utero). No trait is genetically-determined in the sense of being wholly genetically caused; nor are universal, adaptive and seemingly innate traits particularly well insulated against environmental variation—for example, in most primates, development of normal skin structure depends on the ingestion of dietary vitamin C. Adaptations can depend upon learning (plausibly, visuomotor control in primates) and may not be universal or species-typical (e.g. adaptive polymorphisms), nor are all universal traits adaptations (e.g. the human chin). In short, the counterexamples to putative analyses of the concept of innateness proliferate.

The standard counterexamples don't just present an obstacle to a satisfying analysis of the concept of innateness. They also cast doubt on the project of giving any kind of theoretical reconstruction of the property of being innate. The concept of innateness is variously deployed in inferences to and from features (i) to (ix) on the list. Yet the familiar counterexamples show, not just that there is no analysis of the concept that will make most of those inferences deductively valid, but that on any potential theoretical reconstruction, using any concept to make inferences back and forth between features (i) to (ix) will lead to false conclusions in many actual cases.

This article argues that the *i*-properties nevertheless cluster together, in the sense that they tend to be co-instantiated better than chance; at least, they tend to co-occur in non-human animals. Conversely, when a trait lacks one or more of the *i*-properties it tends, better than chance, to lack the others. Mameli (2008) uses 'cluster' in the stronger sense of forming a homeostatic property cluster (Boyd, 1991), and one which supports inferences between *i*-properties with a relatively high degree of confidence. Mameli and Bateson (2011, p. 441) similarly require a cluster to vindicate most of the inferences with a high level of reliability. I claim only that the *i*-properties cluster together in the weaker sense of being co-instantiated together better than chance and for a reason. That claim is still strong enough to be incompatible with Mameli's and Bateson's preferred conclusion, namely that the *i*-properties form a *clutter*—that there is no syndrome which explains why the *i*-properties tend to co-occur and that the idea that they do is an illusion (Mameli and Bateson, 2011, p. 441). This article takes up Mameli and Bateson's challenge to 'give an account of the *i*-properties that constitute the cluster, and of the causal processes that connect such properties and cause them to co-occur' (p. 441).

I should make clear from the outset, however, that I don't think the *i*-properties cluster together well enough that the concept of innateness should be retained. It remains deeply problematic for all the reasons alluded to above. Furthermore, once we have seen the reason why the *i*-properties co-occur in other animals—to the extent that they do—it becomes clear that they will tend to come apart particularly radically in respect of human behavioural traits. That makes innateness claims in psychology especially misleading.

1.1 The Importance Attached to Innateness

The topic has considerable importance given the central role played by innateness claims in both everyday and scientific thinking about the mind. In social and political contexts we use claims of innateness to assess whether outcomes like educational attainment, vocational skills, moral principles, health and happiness can be affected by policies and interventions that we might adopt as a parent, as a teacher, or as a society. The claim that a trait is innate also plays a role in assessing personal responsibility. People typically think of innate characteristics as outside a person's individual responsibility (although still perhaps subject to censure), on a par with characteristics produced by an external cause that is beyond one's control. The everyday thought is roughly that the development of such traits does not occur via the ordinary personal-level processes of believing, wanting, reasoning and forming intentions. They are taken to be part of the psychological background, as a matter of the phenomenology of an individual's psychology—part of the framework within which our psychological development operates. By contrast if a person behaves in a particular way or holds a certain belief because of learning, then she will be aware of the etiology, of the reasons why she reached that outcome—or so the thought goes.

The sense that innate traits develop by a route that bypasses personal-level awareness is clearly part of the popular idea that some behavioural dispositions are 'hard-wired'. With that idea the confusions multiply further. The idea is supposed to be that, if a psychological trait is universal, then it is probably hard-wired, and not under the control of personal-level choice and reasoning. Furthermore, if we discover a basis in the brain for the trait (as if there might not be one), then that confirms that the behaviour has been hard-wired into us by our genes. The result is a capacity that is insulated from reason and that we can't do anything about, even through society-wide environmental interventions. The term 'hard-wired' lends a pseudo-scientific credibility to this misguided tangle of commitments.

The concept of innateness is also relied on in science to underpin inferences between various of the *i*-properties (Knobe and Samuels, in progress), probably in different ways in different sciences (Linguist *et al.*, 2011, p. 452). These inferences also go wrong, for the same reasons, most likely in slightly different ways in different disciplines (Mameli and Bateson, 2006).

So a concept which is important in both scientific and everyday use misleads: it underpins bad inferences to false conclusions. But even if a concept of innateness which rolls together the *i*-properties is so deeply confused that it ought to be eliminated, we still need to examine whether the central forms of reasoning in which innateness claims are deployed have any merit. For everyday reasoning, we need to know whether the genuine concerns which have often been expressed using the concept of innateness can be articulated in other ways. Armed with an understanding of why the *i*-properties cluster together to the extent that they do, we can better address society's questions about the efficacy of interventions and of the role of the will in individual development. And for scientific purposes, we want

to know whether the *i*-properties cluster together better than chance, to understand why, and thereby to better characterise when exceptions will occur.

2. Motivating the Approach

2.1 A Role for Genetic Representation

Attempts to give a satisfactory account of innateness repeatedly come up against the ‘interactionist consensus’ (Sterelny and Griffiths, 1999, pp. 97–100)—the well-substantiated view that the development of all behavioural traits depends very substantially on both genetic and non-genetic causes. This point was injected forcefully into ethology by the comparative psychologist Daniel Lehrman in a highly influential paper (Lehrman, 1953). At that point the dominant view, led by Konrad Lorenz, was that many animal behaviours (instincts) were innate in the sense of being determined by the animal’s heredity and developing entirely independently of experience (Lorenz, 1937). Hard-wiring was also part of the consensus, for example: the eminent ethologist Niko Tinbergen saw instinctive behaviours as acts ‘the neuromotor apparatus of which belongs, in its complete form, to the hereditary constitution of the animal’ (Tinbergen, 1942—his view prior to Lehrman’s intervention).

Lehrman criticised the instinct theorists on the grounds that evolved behavioural traits do not arise independently of experience, but develop by many ‘devious’ paths in which causes endogenous to the zygote interact in complex ways with aspects of the environment, including an animal’s experience. Theorists have been searching for a defensible understanding of innateness ever since. No convincing alternative has emerged to replace the simple, but false, idea that the development of some traits is determined by the animal’s genetic heredity and is independent of experience.

Another, older idea is that, for some traits, the *capacity* to develop the trait is internal, and present at birth, even if the trait itself is not. The interactionist consensus also makes trouble for this idea. If having the capacity to develop a trait at birth is compatible with the trait’s development being causally dependent on the environment in intricate ways, it looks as if the infant is born with the capacity to develop every trait which is a causally possible outcome of development. At least, theorists have so far failed to identify a sense in which infants are born with the capacity to develop only some, but not all the traits that they will eventually go on to acquire by some means or other.

Lehrman’s critique led Lorenz to change his view and formulate a quite different proposal, identifying innate traits as those whose development is caused by genetic information (Lorenz, 1965). That is a modern form of the idea of being born with the capacity for some but not all traits—that the capacity to develop certain traits is in the genes, or genetically coded. Thus, being genetically coded is a further candidate to provide a theoretical explication of the concept of innateness. According to the later Lorenz, deprivation experiments show that the information

needed to make a trait adaptive did not come from the environment, so must instead derive from genetic information.

Lehrman was happy to concede that some traits are coded in the genes; but he thought that does not tell us anything about whether their development will be insulated from environmental influences:

It seems to me, then, that although the idea that behavior patterns are 'blueprinted' or 'encoded' in the genome is a perfectly appropriate and instructive way of talking about certain problems of genetics and of evolution, it does not in any way deal with the kind of questions about behavioural *development* to which it is so often applied (Lehrman, 1970, p. 136).

Lehrman was reacting against the prevailing assumption that a behaviour's being coded in the genes told you all you needed to know about its development. Lickliter and Berry (1990) homed in on this idea and christened it the 'phylogeny fallacy'. The fallacy is to assume that facts about the evolutionary history of a trait should 'lead to the belief that the process of development is thereby somehow explained or understood, eliminating the need for any further investigation or research' (1990, p. 354). Lehrman's arguments forcefully made the point that a trait's being genetically selected and coded in the genes does not in any way exclude the possibility of its development depending in intricate ways on the environment.

The Lehrman quotation suggests something stronger, however: that being genetically coded cannot tell us anything about how development will unfold. Griffiths and Gray (2005) take a similar line, accepting that an informational treatment of (some) heredity systems is useful for thinking about evolution, but rejecting its relevance to understanding development. One thought is that being genetically coded is a historical property, being a matter of what a gene was selected for, and that historical properties are causally inert. Of course it would be wrong to think that a trait's evolutionary history was continuing to make an impact on development that could somehow trump current causes. Distal causes have their effects through more proximal causes, if they are to have any effect at all. But it is not a fallacy to think that facts about evolution and selection can predict and explain patterns in and features of the traits we see today. I agree that such facts do not predict that genetically evolved traits will develop independently of the environment—Lickliter and Berry were right that it is incorrect to presuppose that genetic coding excludes interactionist development—but I argue below (section 3) that there are some predictions about how development is likely to unfold that can generally be drawn from the fact that an aspect of the adaptively-relevant information carried by a trait is genetically represented.

Other authors don't think it is ruled out that genetic information could play a role in explicating innateness, but doubt that there is a satisfactory account of genetic information that can do the job:

In the absence of a good account of the general notion of genetic information, identifying innateness with the genetic coding of phenotypes is trading one confused notion for another and is therefore no progress at all (Mameli and Bateson, 2006, p. 159).

Godfrey-Smith (2007) positively rejects the idea that the genome carries (semantic) information about phenotypes. But, interestingly, he takes that to be a major argument in favour of a critical or deflationary treatment of innateness. These objections suggest that, if there were a good notion of genetic information, that property could help us to understand innateness. The present article takes up that challenge in the light of recent theoretical work on genetic information (Shea, 2007a, 2011, forthcoming a; Bergstrom and Rosvall, 2011). This research builds on earlier proposals for a teleosemantic treatment of genetic information (Sterelny *et al.*, 1996; Maynard Smith, 2000; Jablonka, 2002),⁷ and follows Jablonka and Lamb's suggestion that a function-sensitive notion of information is a profitable way to understand the special features of inheritance systems (Jablonka and Lamb, 2007, p. 382). However, in sharp contrast to Lorenz, I do not conclude that genetic coding salvages the concept of innateness.

2.2 Relating This Approach to the Concept of Innateness

Our project is to see how the *i*-properties relate. That is a question about properties in the world, not about anyone's concepts. Granted, it is the rich history of work on scientific and folk concepts of innateness that has shown there to be a question here at all. But the present article is advancing a thesis about the relations between the *i*-properties themselves. It remains neutral on questions about the nature of the concept or concepts of innateness used by the folk and in various sciences. So I am not claiming that in using the concept of innateness the folk are thinking about genetic representation, or trying to keep track of the effects of genetic representations without realising it. Nor that genetic representation gives the meaning of innateness as it is used in one or another of the sciences. It is not an accident that genetic coding has emerged in scientific discourse as one possible theoretical reconstruction (e.g. in the debates canvassed above), but in my view that is because there is a real but exception-plagued connection between the property of being genetically coded and the other *i*-properties. That is a connection in the world. Our project neither appeals to, nor sets out to explain, facts about how the *i*-properties are related in the minds of either scientists or the folk.

Although it is not the concern of this article, some might want to deploy this account in the debate about the concept or concepts of innateness. A theoretical

⁷ Teleosemantic accounts of genetic information remain controversial (Griffiths, 2001; Moss, 2003; Sarkar, 2003; Stegmann, 2005); the current article is not intended as a full-scale defence of teleosemantic treatments, but instead aims to show that useful explanatory work can be done with the account of genetic representation found in Shea, 2007a.

reconstruction can appeal to a property that is outside the ken of any user of the concept (e.g. the property of being H_2O as a theoretical reconstruction of the pre-modern concept of water). But if genetic representation were offered as a theoretical reconstruction of the concept of innateness, it would have to be an account that showed why concept users (scientists or the folk) could get away with relying on the concept when there is so much scope for inferences amongst the i-properties to go awry, especially in its application to human behaviour. In this respect it would be more like explaining how chemists once got by with the concept of phlogiston. Even if such a theoretical reconstruction were available, it is a further question whether anybody's innateness concept in fact refers. That involves thorny issues in the theory of reference. I don't need to get into those issues here because I argue that people should stop using the concept in any event, in both its everyday and scientific incarnations, irrespective of whether any of these uses succeeds or fails to refer. In my view, so much confusion is now associated with the concept (or concepts) of innateness that it would be better eliminated.

3. Inherited Representation and Adaptive Match

This section sketches an account of genetic representation that I have developed elsewhere (Shea, 2007a, 2011, forthcoming a) and shows how it gains some explanatory purchase on developmental questions. It allows us to explain how an adaptive match between a feature of an organism and an aspect of its environment can be the outcome of individual development.

By way of analogy, consider the adaptive match between the defensive armour grown by the water flea *Daphnia pulex* and the presence of predators. The flea only grows the armour if it detects in the water a particular chemical, which is a sign of predators.⁸ It has a developmental mechanism that is designed to detect an informational signal in its environment and react with an appropriate outcome. The mechanism of adaptive plasticity in *Daphnia pulex* has a structure found in many simple systems that make use of representations. It contains a *consumer mechanism* which has the function of reacting to a range of different signals (chemical, no chemical) with a range of different outputs (armour, no armour). For each output, there is a specific condition that obtains when it performs its function in the way that accounts for its having been selected. For the production of armour, the condition is that there are predators present. The signal to which the consumer mechanism reacts (the chemical) correlates with that condition. When a mechanism has these features, teleosemantics furnishes a legitimate sense in which the signal to

⁸ In fact the armour can also be triggered by exposure of the mother. This maternal effect is also a mechanism of adaptive plasticity (albeit transgenerational), carrying information for the same reasons (a 'detection-based effect': Shea, Pen and Uller, 2011; Shea, forthcoming b), and so is omitted for ease of exposition.

which the consumer system reacts is a representation, at least of a low level variety, having correctness conditions, or satisfaction conditions, or both (Millikan, 1984; Shea, 2007b).

These intentional idioms do not imply that the water flea has any understanding of the signal. It is just reacting to a chemical concentration. Considered merely causally, the cases where it 'reads' information in its environment are just a subset of all the causal factors that play a role in development. They are a subset with a particular evolutionary function. Identifying them as such delivers additional explanatory purchase. Analogously, the evolved functions of an organ like the heart are a proper subset of its causal-role functions (e.g. producing characteristic sounds), and having the extra property of being an evolved function makes a predictive and explanatory difference. One explanatory payoff of recognising that the fleas are representing the presence of predators comes when we ask how the fleas manage to match their armour (on / off) adaptively to the environment (predators / no predators). That match is not accidental. It is the result of a mechanism designed to detect, in the course of development, information about predators. The explanandum concerns the source of information, calling for an explanans in terms of reading information from the environment (in this low-level, but more than merely causal, sense of 'read').

Less obviously, the same structure applies to genomes. The DNA in a zygote is a signal that is transmitted between generations of organisms (Bergstrom and Rosvall, 2011). The mechanisms of DNA expression are a consumer system for zygotic DNA. They take different genotypes as input and, as a result of all the complex interactions of development, produce different phenotypes as output (including 'extended' phenotypes). Crucially for our purposes, the mechanisms of DNA replication, expression and repair have evolved to perform the function of transmitting phenotypes down the generations. Once a particular genotype has been selected in virtue of some particular phenotype to which it gives rise, producing that phenotype in response to that genotype becomes one of the functions of the mechanisms of DNA expression. For example, a genotype G^* might be selected over other genotypes G in a population because, against the background of the range of genetic and environmental conditions present at the time of selection, expression of G^* tended to produce organisms with thicker hair (fitter because the environment was cold, say). Then, producing thick hair in response to G^* would be one of the functions of DNA expression. Furthermore, only if the environment was cold would expression of G^* perform its function in an evolutionarily normal way (i.e. in the way which accounted for its selection). In our simple example, the result of genetic selection is that the frequency of G^* over G increases; so G^* comes to correlate with the colder environment. Those are all the pieces needed for there to be representational contents in simple systems. The result of selection on genes for their phenotypic effect—an uncontested verity of evolutionary theory—is that genes carry representational contents, of both the indicative and the imperative variety (Shea, 2007a). In our example G^* in the zygote represents: *the environment is cold, produce thick hair*.

This account shows how the process of natural selection can give rise to semantic information (i.e. representational content), and not just correlational information (a correlation between types, e.g. of the Shannon variety). It can also apply to other mechanisms of inheritance, but not automatically. If there are to be genuine representation consumers in the temporally-extended processes of inheritance, then we need to appeal to something more than the existence of selective functions. No version of teleosemantics will be apposite—and this has not been sufficiently appreciated—unless a mechanism of inheritance not only happens to act as a locus for transmitting phenotypes down the generations, but also has the evolutionary function of doing so. Only then is there a real consumer that has the function of responding to a range of different representations (e.g. transmitted genotypes) with a range of different outcomes (phenotypic traits). That is a demanding constraint.⁹ Not every case in which non-genetic phenotypic effects are heritable will satisfy it.

For example, membrane structures in the zygote act as templates to construct the membranes of daughter cells (Jablonka and Lamb, 2005). Even if changes in these features were heritable, it seems unlikely that the whole system has been selected in order to be a locus for the transmission of phenotypic variants. By contrast, if some form of chromatin marking plays a substantial role in germ line as well as somatic cell inheritance (Jablonka and Raz, 2009), then there may well have been selection on that mechanism for its role in transmitting phenotypes down the generations. Transmission of behavioural phenotypes by imitation in humans is another candidate for having the demanding (meta-) function, and so qualifying as a system for transmitting representations between generations (Shea, 2009). I use ‘inherited representation’ for representations, carrying correlational information as a result of natural selection, which are transmitted down the generations by inheritance mechanisms with the *selective function* of transmitting phenotypes down the generations.¹⁰ Genetic representation is one species of inherited representation.

Armed with this account of genetic representation, we can take the explanation of adaptive match from cases like the water flea’s defensive armour and apply it to cases where the adaptive match is due to a history of selection. The rough idea is that genetic representation can explain how it is that sometimes an animal knows what to do (i.e. produces species-typical adaptive behaviour) without having to learn (M. S. Dawkins, 1995, pp. 55–69 presses this question). For example, African horsefly larvae of the genus *Tabanidae* engage in corkscrew burrowing behaviour before pupating in wet mud. The corkscrew burrow is designed to protect the pupa from cracks that will form if the mud dries out and cracks. We can focus on an individual larva and ask what accounts for the adaptive match between the burrow it builds and its future environment (cracked mud). If the standard

⁹ This is to agree with Godfrey-Smith (2007) that the real existence of consumers is a very demanding constraint, but to disagree about whether it can be met.

¹⁰ Shea, Pen and Uller (2011) and Shea (forthcoming b) call these ‘selection-based effects’.

understanding of the burrowing behaviour is correct, then the burrowing tendency arises as a result of expression of genes that have been selected for producing that behaviour, where natural selection on phenotypes has increased the frequency of the burrow-producing phenotype. The question about adaptive match in an individual case is answered as follows: developmental processes in that individual larva relied on a genetic representation that carried representational content which instructed a particular phenotypic outcome (the burrowing behaviour) and indicated that the environment would be conducive to that phenotype (the mud would crack).

As with *Daphnia* detecting chemicals, considered causally genes are just one of very many factors involved in developing the burrowing behaviour in *Tabanidae*. That these particular causal interactions are also a matter of reading instructional and indicative content in the genes does not imply any spooky mind-involving processes or a special kind of causation. What it does is to identify a proper subset of causal factors in development that are apposite for addressing a particular explanandum: how did the (individual) larva achieve an adaptive match between its behaviour and the environment? That is an informational question, which we can answer by recognising that some of the causes in development are also representations.

The account of genetic representation outlined here is set out and argued for more fully in Shea (2007a). The account of representation by the genome developed there and relied upon here does not depend at all on how the idiosyncratic complexities of development unfold. It requires correlations between genotypes and phenotypes at the time of selection, and can then treat development as a black box. In Shea (2007a) I went further and conceded the Lehrman point quoted above—that genetic representation cannot tell us anything about individual development. Shea (forthcoming a) withdraws that concession and argues that, where development involves reading genetic representations (in our sense), some conclusions about development do tend to follow. For instance, that makes it unlikely that there will be a mechanism of adaptive plasticity designed to pick up on the same kind of information in the course of individual development. The present article endorses that view. While agreeing with Lickliter and Berry (1990) that it would be a big mistake to think the phylogenetic facts (genetic representation) exclude the possibility that development depends in intricate ways on the environment, I do think genetic representation gives rise to some generalisations about development, as argued for in the next section.

To recap, an informational question about individual development (adaptive match to the environment), which is sometimes explained by the individual detecting information in the course of development (e.g. in *Daphnia*), is also sometimes answered by observing that information generated by a process of natural selection over many generations is read in the course of the development (of an individual) so as to give rise to a phenotypic feature which adaptively matches an aspect of the individual's environment.

4. Explaining the Cluster

This section argues that, where some aspect of a trait develops in reliance on a genetic representation it will tend, better than chance, to have many of the *i*-properties (especially in non-human cases). As a result, even without any analytically-guaranteed connection, using an innateness concept to make inferences back and forth amongst the *i*-properties is more likely to lead to true conclusions than if the concept were entirely defective. The discussion is initially restricted to inferences based on genetic representation, widening later to include other forms of inherited representation.

4.1 Universality / Species-Typicality

Where a trait develops in reliance on a genetic representation it will tend to be universal or species-typical. That is simply because only a small relative fitness advantage is enough to drive a trait to fixation in a large population (Bromham, 2008, p. 144). We must distinguish here between adaptive phenotypic plasticity and genetic polymorphisms. Evolution sometimes selects for a genetic polymorphism, which means that there are non-universal traits that are gene-based adaptations, the development of which is relatively insensitive to variations in developmental environment. In plants, suites of such local adaptations can give rise to ecotypes—genetic variants adapted to the ecology of their local environment. Similarly, genetically-selected quantitatively variable traits are important exceptions to the generalisation about universality.

A source of genetic polymorphisms in sexually-reproducing species is heterozygote advantage, in which case the phenotype of the heterozygote may be an adaptation (e.g. against malaria) while the phenotype of the homozygote is not (e.g. sickle cell disease). The difference between homozygous and heterozygous carriers of the sickle cell gene is a genetic difference, but not one that has representational content. The sickle cell gene carries the information that the environment has a high rate of malaria and instructs production of a protective phenotype. The homozygote carries no more information than the heterozygote, and suffers from the disadvantage of causing sickle cell disease. By contrast, if we look at the contrast between human groups that do and do not carry the sickle cell gene, carrying the gene does carry the information that you are likely to be in a malarial area. So this is a case where a non-universal characteristic develops in reliance on a genetic representation. Another case is where frequency-dependent selection leads a population to contain a mix of two phenotypes, each of which is adaptive in the context of the other (at a certain frequency). Despite such exceptions, given that gene-based evolution has produced the robust patterns of similarity and difference of characteristic traits observed across species and other taxa, as catalogued by systematics over hundreds of years, a trait's developing in reliance on a genetic representation clearly lends some inductive support to the conclusion that it is likely to be universal in the species.

One type of case that produces variation in traits should not be considered an exception, however. In cases of adaptive phenotypic plasticity, different individuals develop different phenotypes, with each tending to be adaptive in the environment in which it is found. The *Daphnia* water fleas vary in their armour. Each variant is adaptive in its particular environmental circumstances (predators / no predators). Neither outcome is universal or species typical (in the relevant sense). Nevertheless, adaptive plasticity does not sever the connection between adaptation and universality, since individuals with different phenotypes share a mechanism of phenotypic plasticity—a mechanism that detects which kind of environment the individual is in and causes it to develop the adaptively-appropriate phenotype. The mechanism of adaptive plasticity is universal to the species. The investment needed to develop a mechanism of adaptive plasticity itself is only adaptive because (i) there is a significant chance of living in a predator-rich environment; (ii) there is a significant chance of living in a predator-free environment; and (iii) the costs of defence are such that it is worthwhile to incur them in (i) but to forego them in (ii). Genetic representation is not contributing specific information about the local environment (predator-rich vs. predator-free), but it is contributing the information encapsulated in the mechanism of adaptive plasticity, namely: there may well be predators and, if such-and-such chemical is present, there are more likely to be predators, so produce armour; if not, not, so don't.

Sometimes discussions of the connection between adaptiveness and universality make the mistake of identifying the potentially adaptive trait at the wrong level of grain. If a phenotypic polymorphism derives from the operation of a mechanism of adaptive plasticity, then that shared mechanism will often be a gene-based adaptation. If so, the information encapsulated in the mechanism of adaptive plasticity derives from a genetic representation, so this is a case where reliance on genetic representation gives rise to a universal characteristic: not the final outcome, but the possession of the mechanism of developmental plasticity.

The inference in the opposite direction, from universality to developing on the basis of a genetic representation, only has any force when an adaptive aspect of a trait is under consideration. Even then there will be exceptions. Consider what would happen if the density of predators increased so that there were always predators in the water flea's developmental environment, with defensive armour becoming universal. Armour would then be a universal trait which adaptively matched the environment (predators) but which was due to a developmental process designed to pick up on environmental information (the chemical signal). In practice, genetic assimilation will militate against the perpetuation of such cases (Gilbert and Epel, 2009, pp. 381–4). When the range of environments narrows so that a mechanism affording adaptive plasticity to other environments is no longer needed, there will be selection pressure against maintaining the mechanism and selection pressure in favour of canalizing the phenotype against environmental variation. Nevertheless, it is by no means guaranteed that a universal adaptation develops in reliance on a genetic representation. The problem is particularly acute with human psychological traits, since

human development is supremely adaptively plastic, even with respect to features of the environment that are very widely shared (e.g. the belief that water is wet).

4.2 Developmental Invariance

Traits that develop in reliance on a genetic representation will tend to have a certain kind of developmental invariance, that is, they will tend to develop across the range of environments that were present when the trait was selected. For natural selection to have operated in the first place the phenotypic variant under selection has to have had a reasonably high degree of heritability, which is to say that the genotype-phenotype correlation must have been robust across a reasonably wide range of the developmental environments encountered by the organism during selection. To the extent that current or experimentally induced variations in the developmental environment fall within the range encountered during selection of the trait, the selected outcome will tend, better than chance, to be produced despite such variations.

Once we go outside the parameters encountered in selective environments, there is no reason to expect developmental invariance. To take up an earlier example, development of normal skin structure in primates is invariant in the face of wide variation in diet (e.g. from berries to tropical fruit), but does not develop in environments that differ from the historical range and contain no vitamin C in the diet (e.g. a diet of pure junk food). So a trait that is developmentally invariant across the range of environments in which it was selected may be very sensitive to variations that occur amongst current environments.

The extent to which an adaptive trait is now developmentally invariant depends upon the extent to which past selective environments varied in ways that are causally-relevant to its development, and the extent to which causally-relevant variations in currently-experienced environments exceed those limits. For many traits of non-human species the range of relevant variation in historical and current environments is sufficiently similar that traits which develop in reliance on a genetic representation will continue to develop across a wide range of current developmental environments. However, where the environment has changed beyond the historical parameters there will be exceptions, and those exceptions will be especially important in humans, since the range of environments in which humans develop has now been altered radically in some respects.

A trait is *canalized* against variations in the environment if the way it develops is buffered against environmental variations or if there are alternative pathways or backup mechanisms for its development which ensure that the same outcome is achieved despite differences in the environment. One large-scale example is thermoregulation, which canalizes a whole host of developmental pathways against variations in the temperature of the developmental environment. Canalization is an active variety of developmental invariance.

Adaptation also explains developmental canalization (Ariew, 1999). A gene may be selected initially despite not producing beneficial outcomes in all circumstances. Some environments may lead to a developmental outcome $P+$ that is

fitness-enhancing and some ('hindering' environments) to an outcome P− that is fitness-reducing. The gene will be selected if the fitness benefits of P+ outweigh the fitness costs of P−. However, once the gene has been selected, there remains a selection pressure to produce P+ in the hindering environments as well. So mechanisms may be selected which allow P+ to develop across a broader range of environments. Development of a trait can thereby come to be canalized against environmental variation (against the variation encountered in its evolutionary history). A trait will be canalized across variations in the current environment to the extent that those variations overlap with the variation encountered in its selective history.

Invariance in the face of environmental variation is a degenerate form of a norm of reaction: a mapping from environments to phenotypes, given a particular genotype. A genotype may have a much more complex norm of reaction. A common example is the level of activity of an enzyme produced by a particular gene (Lewontin, 1974, p. 407; Nijhout, 2001). The enzyme will have different levels of activity at different temperatures. In organisms without thermoregulation, the effects of the enzyme will vary continuously across variations in the temperature of the developmental environment. (A genetic variant may produce a related enzyme with a different temperature optimum.) The gene for the enzyme will tend to be selected if, across the range of evolutionary environments, the fitness benefits of the different phenotypes produced are on average positive (averaging according to the frequency with which those environments are encountered, modified to take account of population structure in some cases). So natural selection can explain why development in a species has the reaction norm it does (e.g. Lind and Johansson, 2007; Pen *et al.*, 2010). Fitness benefits at one or more points on the norm of reaction may outweigh the fitness costs of the phenotypes which result at other points. Where the norm of reaction is such that there are two or more phenotypic outcomes with each conferring a fitness advantage in its particular environment by a different route, then we are back to a case of adaptive plasticity, as discussed above.

In other cases a norm of reaction has been selected because it produces an adaptive mix between different phenotypes. Sex ratios are an example. In many crocodilians, sex determination depends on the incubation temperature of the egg. Either different ratios have been adaptive in different environments, or the variability of temperature in selective environments has been sufficiently wide to ensure that this mechanism produces an adaptive sex ratio. In such cases, there is genetic selection for a mechanism to produce a particular overall sex ratio. The genome carries an instruction about the probability with which each sex should be produced (e.g. 50% chance of producing a male). In the course of individual development, temperature is not being consumed as a representation of some adaptively-relevant environmental parameter; rather, it is part of the causal basis on which a given sex ratio is achieved. Such cases serve to weaken the connection between genetic representation and developmental invariance, but do not undermine it completely.

The inference in the opposite direction, from the developmental invariance of an adaptive aspect of a trait in current environments, to its developing in reliance on a genetic representation, will go through in many cases. As with universality, there will be exceptions when the current environment has become narrower or otherwise altered to be restricted to only one of the range of environmentally-cued phenotypic outcomes on which selection originally acted. Those exceptions will be especially important where cultural processes can lead learnt outcomes to arise across a wide range of current developmental environments (e.g. the belief that water is liquid).¹¹

4.3 Evolutionary Adaptation

By construction, any trait that develops in reliance on a genetic representation is an evolutionary adaptation. As noted above, not all evolutionary adaptations need be transmitted by DNA, or indeed by any other inheritance system, so the inference in the other direction is not watertight. Nevertheless, given the undisputed importance of genetic evolution, especially outside humans, the inference from evolutionary adaptation to developing in reliance on a genetic representation is strong. As before, care is needed about the level of grain: it is the aspect of a trait whose adaptive match to the environment is explained by a history of selection about which we can say that it develops in reliance on a genetic representation.

A potential objection is that inferences can be made between adaptation and the other *i*-properties directly—so why go via genetic representation? The informational treatment has the advantage of allowing us to predict and explain features of the developmental process. Those explanatory payoffs are set out in detail in Shea (forthcoming a) and there is not space to repeat them all here, but Section 3 above highlighted one explanatory question that is distinctively informational: asking where in the course of individual development the information derives that accounts for development eventuating in a trait, a feature of which adaptively matches the environment. An explanation in terms of genetic representation is more direct than one that can appeal only to genetic adaptation. Secondly, by underwriting parallels between genetic and environmental sources of information in development, it can explain the existence of trade-offs between the two sources (Leimar *et al.*, 2006). Even the later Lorenz (1965) was committed to a dichotomy: the information required for an adaptive match must derive either from the environment or from inheritance but not both. A merit of our treatment of genetic representation is that it shows how both genetic and environmental information may contribute to a particular adaptive match.¹² A third advantage is that it allows us

¹¹ This is to agree with Griffiths and Machery (2008) that human behavioural adaptations may well not be canalized against environmental variation.

¹² The possibility of such a mix was one of Haldane's criticisms of the later Lorenz's account of innateness (Griffiths, 2004).

to see the connection between poverty of the stimulus arguments—which concern the sources of information available to development—and the other *i*-properties, as we shall see next.

4.4 Learning / Poverty of the Stimulus

One way to test whether an adaptive match between trait and environment develops in reliance on environmental information is to examine whether the trait develops in the same way in the absence of any relevant cues from the environment—i.e. that a poverty of the stimulus argument can be made about an aspect of the trait. The poverty-of-the-stimulus test is compatible with the interactionist consensus and is not aimed at assessing the extent of causal interaction with the environment. The development of a trait may be the result of rich causal interdependence between a series of internal and external factors, but if the relevant information (e.g. that the mud the animal is burrowing in is likely to dry up and crack) is not detectable in the environment given the organism's sensory and learning mechanisms, then we can infer that that aspect of the trait (e.g. the pattern of burrowing behaviour) developed in reliance on a genetic representation. Cross-fostering experiments, for example, are designed to uncover whether development of a trait depends, not just causally on the environment, but specifically on reading adaptively-relevant correlational information that is specific to the environment into which the offspring are fostered.

If poverty of the stimulus arguments are couched in terms of the bare correlational information available in the environment, they tend to miss their mark. Whenever the development of a trait is causally dependent on an aspect of the environment, the environment will carry information—in the bare correlational sense—about the trait (Griffiths and Gray, 1994). But these are not examples of the mechanisms of development reading or consuming a representation. For example, primates do not 'read' information in dietary vitamin C in order to produce the outcome of normal skin structure (although dietary vitamin C vs. its absence does correlate with normal versus abnormal skin structure). Only when the way development reacts to a piece of correlational information in the environment is a matter of evolutionary design, with the function of producing a variable outcome depending upon the detected state of the environment, is it right to think of development as reading or consuming a representation in the environment. So the connection between the collection of *i*-properties and genetic representation is tighter than the connection between the *i*-properties and poverty of the stimulus arguments which understand information in just the correlational, causal sense (Khalidi, 2002, following Chomsky). An account which treats a trait as innate only if it would still emerge in an impoverished environment will end up wrongly excluding very many cases where genetic representations are relied upon, but where rich interaction with the environment is still needed for normal development. These range from non-biological supporting causes (temperature, light, nutrition), through the effects of stimulation on the nervous system (both passively, and through spontaneously

and exogenously-driven practice), to various practice effects that depend upon the physics of the limbs and the environment.

Pointing to semantic information allows a subtler treatment. Many cognitive traits that result from learning also depend upon pre-existing cognitive architecture, constraints or biases, in which case the trait develops partly in reliance on environmental information and partly in reliance on the inherited representations which are responsible for those constraints. Only with respect to the latter is an inference to other i-properties (universality, invariance, adaptation) inductively supported. Indeed, there may be cases where the cognitive trait would not develop at all in the context of impoverished stimuli, thus counting as non-innate on Khalidi's account; but where, when it does develop normally, some of the information in the mature structure has a non-environmental source. Face perception is a plausible example: in order to recognise faces we need a lot of experience of looking at faces, from which we extract information about regularities and invariants; but the basic neonatal bias to attend preferentially to faces may well depend on genetic representation, and be a developmentally-invariant, species-typical adaptation for developing face recognition abilities (Johnson *et al.*, 1991).

Poverty of the stimulus arguments test whether adaptively-relevant information has been extracted from the environment during development. Learning is the central example. Where the adaptive match between an aspect of a trait and the environment is due to learning, there is unlikely to be genetic representation of that same feature, or development would have been canalized to the adaptive outcome in any event. Conversely, genetic representation obviates the need for learning. These considerations underpin the success of the generalisation: genetically represented $\rightarrow$ not learnt¹³ (nor the result of a mechanism of adaptive plasticity); and also the converse: where a feature of a trait adaptively matches its environment but is not genetically represented, it is likely to be learnt or acquired via a mechanism of adaptive plasticity (again with many more exceptions in humans where non-genetic adaptations are much more important). That shows why there really is a connection between learning (i-property (viii)) and the other i-properties discussed above: universality, invariance and adaptation.

4.5 Other i-properties

As we saw above, genetic determination (i-property (ii)) cannot plausibly be construed as genetic determinism (too strong, contra the interactionist consensus), nor as there being some genetic cause (too weak—applies to every trait). There may be a good distinction to be drawn between active causes and background causes, or an account which could be given of the relative causal contributions of external and internal factors to the development of a trait. Intuitively, body plan depends more on internal causes than does dialect, say. Philosophers have attempted

¹³ Cp. innate $\rightarrow$ not learnt; equivalently learnt $\rightarrow$ not innate.

to give a more rigorous treatment these distinctions (e.g. Northcott, 2006), but although that project is not conclusively refuted (Godfrey-Smith, 2007), neither is it satisfactorily resolved. If it were, there could turn out to be a connection between being genetically represented and genetic factors making a greater causal contribution to development, although it is not clear why that should be.

By contrast, the connection of genetic programming (i-property (vi)) to the other i-properties is vindicated. My claim is precisely that being genetically programmed, in the sense of being represented by the genome, is the underlying property which explains why the other i-properties cluster together in actual cases.

The connection with being present at birth (i-property (i)) also makes sense. Many traits are not present at birth, nor is the baby born with a capacity to develop the trait without environmental support, but where an aspect of a trait develops in reliance on a genetic representation, the information was indeed present at birth. The inference in the opposite direction is less secure since, as we have seen, the new-born baby has already undergone learning in utero, and may also have received signals about its likely environment directly from its mother in the form of maternal effects.

The connection with heritability (i-property (vii)) is less direct. Where there are different phenotypic variants which are each adaptive, then if those variants are genetically heritable it very likely that development of the trait depends upon reading a genetic representation (e.g. genetically-based adaptive ecotypes). In addition, natural selection on genes only operates if the phenotypes which they produce are heritable. But that only requires heritability at the time of selection, assessed relative to the range of genotypes in the population at the time of selection and the range of developmental environments that were encountered. Adaptations that have gone to fixation will no longer be heritable (except during episodes of stabilising selection). So a constrained kind of inference from heritability to genetic representation and developmental invariance is supported, but not the converse.

The extent of the connection between genetic representation and modularity (i-property (ix)) is more tenuous. As we will see in the next section, information which is implicit in the operation of an informationally-encapsulated psychological module need not derive from genetic representation, but may come from culturally inherited representations, or from learning.

Finally, genetic representation throws light on the intuitive idea canvassed at the outset: that innate traits are hard-wired and develop in a way that bypasses the processes of personal-level psychology (the will, etc.). When a genetic representation is consumed so as to lead to a particular developmental outcome, that is not a psychological-level process of reading or learning from a piece of information. From the perspective of the individual person, genetic representations form part of their constitution, as one of the set of factors that structure who they are and how they behave. Development in reliance on a genetic representation is not causally determinative of a particular outcome—the result may be malleable in the light of a person's intentions and decisions. Nevertheless, the way the genetic effect

produced an outcome at the time it was selected is unlikely to have proceeded via personal-level psychology.

For example, the facial expression of basic emotions ('affect programs') is widely shared across considerable variations in the kinds of environments in which people grow up. Plausibly, many of these dispositions have social-communicative evolutionary functions. But the development of these dispositions does not seem to depend upon individuals reasoning about how they should communicate their emotions and, once developed, the reason why we smile (rather than frown, say) when we're happy is not available at the personal level. Nevertheless, I can decide not to and, through a programme of the right kind of training, effectively give myself the appearance of a sulky old grump. So genetic representation gives us a sense in which innate traits, while not being causally determined, nevertheless normally develop by a route that does not depend on personal level mechanisms (or the will if such exists).

4.6 Application to Non-Adaptations

Consider pathological cases: genetic diseases and acquired pathology (e.g. aphasia caused by head trauma). My claim has been that, when applied to an adaptation, various i-properties are a rough guide to whether an aspect of a trait does or does not develop in reliance on a genetic representation. I have not discussed whether the i-properties might be connected or come apart when applied to features which are not adaptations (neutral or pathological traits). The existing literature on innateness suggests that they come apart in very many ways.

Nevertheless, people are inclined to treat genetic diseases as being innate and acquired pathology as being non-innate. Such inferences may not be vindicated by generalisations that really connect the i-properties, but they might still be explained by features of the concepts of innateness used by scientists or the folk. In particular, in both vernacular uses (Griffiths, Machery and Linquist, 2009; Linquist *et al.*, 2011) and scientific uses (Knobe and Samuels, in progress) fixity is a strong driver of innateness judgements. Indeed, in everyday uses fixity has become somewhat detached from the property of being present at birth, in talk of one's 'political DNA' and such like, for example: 'This music was so much a part of my upbringing that I feel very connected to it. It's totally part of my DNA'.¹⁴ I would suggest the following captures a common way that the concept is applied, at least in its everyday use:

- If
- (i) an organism O has trait T at some time $t' > t$; and
 - (ii) O would have had T at t' even if the environment with which O interacted between t and t' had been different in certain ways (specified by the context)

¹⁴ Annie Lennox, quoted in *Saga Magazine*, December 2010, p. 48.

then that is good evidence that T was innate to O at t; and the converse.

Application of this inferential frame to a pathological condition caused by an environmental insult will count it as non-innate. Application to a disease where it is a genetic difference from the normal population which accounts for why an individual has a dysfunctional outcome will, against most contexts of comparison, count the disease as innate. In short, innateness judgements about non-adaptations probably result from application of means of identification that gain their efficacy from the fact that the i-properties tend to cluster together when applied to adaptations.

It is worth remarking that this brings out the contrast between the purely causal notion of genetic determination and our precise sense in which an aspect of a trait may develop in reliance on genetic representation. A phenotypic difference may be caused by a genetic difference (against some background context held fixed), as in the case of a genetic disease—in neutral cases such traits may go to fixation through drift (especially in small populations)—but without selection, there is no genetic representation, and nor is there an explanans at the outcome of development (adaptive match) to which genetic representation would be the explanandum.

5. Where the Inferences Go Wrong

There are several further reasons why inferences between the i-properties go awry, especially when applied to humans, despite the role of genetic representation in generating a genuine cluster. The first is that DNA is not the only source of inherited representations. We noted above that there may be other systems of inheritance that have the function of transmitting phenotypes down the generations. Chromatin marking is one candidate. Learning by imitation in humans is another.

Epigenetic inheritance systems like chromatin marking that are associated with the genome and pass on its state of gene activity or expression stably down many generations will be a source of inherited representation that acts much like genetic representation. When an adaptive trait has evolved on the basis of epigenetic inheritance the information on the basis of which development arrives at an outcome that adaptively matches the environment was present at birth. That information is read in the course of development, so the trait can still develop in the absence of any environmental correlate of the adaptively-relevant matter of fact (poverty of the stimulus). The outcome is likely to be universal or species-typical, and developmentally invariant or canalized against variations in the developmental environment, within the range of environments in which it evolved. Selection of the trait will have depended on there being a heritable epigenetic difference during the course of selection, and if there is still relevant variation in the population the trait will look just like a genetically-heritable trait in a population analysis (Badyaev and Uller, 2009). Only the inference to ‘programmed in the genes’ (i-property (vi)) is undermined. If epigenetic inheritance systems turn out to be an important source of adaptations (Jablonka and Lamb, 2005), then in many cases inferences

amongst the i-properties will be licensed not by genetic representation, but by other sources of inherited representation that are present at birth.

With humans the problem becomes even more acute. Learning by 'over-imitation' in humans may be an adaptation for transmitting behavioural phenotypes down the generations (Shea, 2009). Where a behavioural adaptation evolves by natural selection on variants that are passed on by over-imitation the correlational information that the particular behaviour is fitness enhancing (e.g. that a certain food preparation practice has long-term benefits) is generated by a process of selection. There may be nothing in the environment that would allow an individual to learn for themselves about the value of the behaviour (e.g. because feedback is too stochastic or too delayed). Nevertheless, by faithfully imitating the behaviour of their parents, children could come to develop an adaptive phenotype. So not all inherited representation is present at birth, as Lorenz (1965) claimed.

Outcome-insensitive over-imitation is one example of the way human development may rely on *inherited* representation—where correlations are not detected in individual learning but through a process of selection over many generations—that is read from the developmental *environment*. Aspects of a trait that develop in that way will tend to have the i-properties: they will arise in development invariantly across the range of environmental variation present when they were selected, they may well become universal or typical of the species, and they develop in the absence of any information or feedback that tells the organism that the resulting behaviour is adaptively-relevant. Therefore, development in reliance on environmentally-mediated inherited representation has quite different characteristics from development that involves learning from the environment or other kinds of adaptive plasticity based on reading or consuming informational cues in the environment. That is a further respect in which our account of inherited representation represents a major break from the later Lorenz (1965). The deprivation experiment is a particularly poor test for inherited representation when a substantial amount of inherited representational content is not present at birth, but is transmitted through the environment.

Lorenz argued that development consists in the interaction between heredity and the environment, Lehrman that it consists in the interaction between an organism and its environment (Lehrman, 1953, p. 345). The full picture is that development consists in the interaction between all three: the organism, its heredity and the environment.

Although there are some behavioural traditions in other species, the role of inherited representational content that is not present at birth is particularly important for humans. Humans have lived in social groups or tribes between which there are large behavioural differences, for example between hunter-gatherers and pastoralists and agriculturalists. Anthropology has given us a rich catalogue of these large differences in ways of life, skills, tools, food preparation practices, rituals and mystical beliefs. There is a strong tendency to think of these differences as being due to differences that are present at birth, just as differences between different species or breeds of animals and plants are. The finding that, for most of these

characteristics, humans turn into the kind of people they grow up with, and not the kind of people their parents were (when the two diverge), was a revolutionary discovery which everyday thinking has not yet fully assimilated. It led to a crisis in the world of anthropology, precipitating a split between biological and cultural anthropology (Kronfeldner, 2009).

Several theorists have argued that humans have crossed a Darwinian threshold in the extent to which cultural processes are responsible for producing adaptive phenotypes (Maynard Smith and Szathmáry, 1995; Sterelny, 2006). If so, the developmental environment, in particular the social environment, is an especially important source of inherited representation in human development. We can distinguish between internal causes and environmental causes in development, although the development of most behavioural phenotypes will depend upon both. We can also distinguish between correlational information that an *individual* could detect and learn from, and information that is generated by a process of selection. Those two distinctions do not align at all well in humans: much inherited representation is located, causally, in the developmental environment. It is perfectly cogent to ask whether some trait is biologically or culturally inherited, but it is a mistake to think this lines up with the distinction between development being caused by internal or external factors. Whether an aspect of a trait develops in reliance on biological inherited representation (genetic, epigenetic) or cultural inherited representation (imitation, etc.), the cause of its development is likely to be a rich interplay between internal and external factors.

It follows that the use of the innateness concept—with an associated set of conceptions that includes all or most of the *i*-properties—is particularly problematic when applied to human behavioural traits. In such cases we have little inductive or abductive reason, much less a guarantee, that the inferences involving the concept of innateness will take us to true conclusions.

So when evolutionary psychologists identify a behavioural or psychological capacity that looks to be adapted to features of human lifeways, and claim that it is innate, we should be slow to conclude that it has evolved through genetic selection and develops in reliance on a genetic representation. Culturally inherited representation is likely to be important, as may be processes of individual learning. The development of such adaptive capacities is likely to be much less informationally encapsulated than nativist theorists suppose, even if the online operation of the mature capacity has become relatively informationally encapsulated. For example, development of the human capacity to keep track of the mental states of others may depend heavily on niche construction and persisting changes to the learning environment (downstream epistemic engineering), so that it arises reliably across the wide range of environments in which people grow up, but not in a way that is independent of information from the environment (Sterelny, 2003). Human moral psychology, and its links with the emotions, is another likely example.

Putative universal features of human grammar offer a contrast case. If there are such universals, and if they are adaptations (for example, as a solution to the problem of serial communication), then they may well be due to gene-based

selection. The opposing view is that language universals are due to learnability constraints, with cultural selection on different languages leading to the selection of those that humans can learn most easily. As with many other psychological capacities (e.g. intuitive physics, multi-step reasoning), the question of the relative contribution of biological inherited representation, cultural inherited representation and environmental information is yet to be settled empirically.

6. Conclusion

Philosophers and others have offered various analyses and theoretical explications of the concept of innateness. The leading candidates are developmental invariance or canalization; developing in response to an impoverished stimulus; not being dependent on learning; and being an evolutionary adaptation. None of these treatments lines up with an intuitive classification of traits into the innate and non-innate. Each fails even on some central cases. However, we can understand the attraction of each proposal when we see that they are roughly and non-accidentally aligned. The reason is that each tends to follow when an aspect of an adaptive trait develops in reliance on a genetic representation.

Genetic information had been passed over in the innateness literature for lack of a satisfactory account of that property. This article shows that one merit of a recent but by no means uncontroversial account of (semantic) genetic information (Shea, 2007a) is that it can explain why the *i*-properties tend to cluster together. Armed with a good theory of genetic representation, it becomes clear that there are non-accidental connections between the *i*-properties and the property of developing in reliance on a genetic representation. It also reveals why the concept of innateness is defective, by showing that inferences amongst the *i*-properties are likely to go awry, especially with respect to human psychology and behaviour.

One reaction is that the innateness concept should be salvaged and defined in terms of genetic representation. That would be a mistake, for two reasons. First it would obscure the way the various *i*-properties dissociate even for non-human traits. Working directly with the *i*-properties themselves makes it more apparent where the inferences will fail. Using the concept of innateness makes bad inferences more likely. Secondly, inferences amongst the *i*-properties are particularly likely to fail when applied to human traits, especially to human psychological traits and behaviour. Deploying a concept of innateness in this field, when the concept derives its inductive support elsewhere, has been the source of a series of mistakes in reasoning about human cognition and behaviour. We will understand human psychology much better when we jettison the concept of innateness in favour of sharper distinctions made in terms of the *i*-properties directly.

*Faculty of Philosophy and Somerville College,
University of Oxford*

References

- Ariew, A. 1999: Innateness is canalization: in defense of a developmental account of innateness. In V. G. Hardcastle (ed.), *Where Biology Meets Psychology: Philosophical Essays*. Cambridge, MA: MIT Press.
- Badyaev, A. V. and Uller, T. 2009: Parental effects in ecology and evolution: mechanisms, processes and implications. *Philosophical Transactions of the Royal Society B*, 364, 1169–77.
- Bateson, P. 1991: Are there principles of behavioural development? In P. Bateson (ed.), *The Development and Integration of Behaviour*. Cambridge: Cambridge University Press.
- Bergstrom, C. T. and M. Rosvall 2011: The transmission sense of information. *Biology & Philosophy*, 26, 195–200.
- Boyd, R. 1991: Realism, antifoundationalism, and the enthusiasm for natural kinds. *Philosophical Studies*, 61, 127–48.
- Bromham, L. 2008: *Reading the Story in DNA*. Oxford: Oxford University Press.
- Cowie, F. 1999: *What's Within? Nativism Reconsidered*. Oxford: Oxford University Press.
- Cowie, F. 2009: Why isn't Stich an eliminativist? In D. Murphy and M. Bishop (eds), *Stephen Stich and His Critics*. Oxford: Blackwell.
- Dawkins, M. S. 1995: *Unravelling Animal Behaviour*, 2nd edn. Harlow: Longman.
- Gilbert, S. F. and Epel, D. 2009: *Ecological Developmental Biology*. Sunderland, MA: Sinauer Associates.
- Godfrey-Smith, P. 2007: Innateness and genetic information. In P. Carruthers and S. Laurence and S. Stich (eds), *The Innate Mind: Foundations and the Future*. Oxford: Oxford University Press, 55–68.
- Griffiths, P. E. 2001: Genetic information: a metaphor in search of a theory. *Philosophy of Science*, 68, 394–412.
- Griffiths, P. E. 2002: What is innateness? *Monist*, 85, 70–85.
- Griffiths, P. E. 2004: Instinct in the '50s: the British reception of Konrad Lorenz's theory of instinctive behavior. *Biology and Philosophy*, 19, 609–31.
- Griffiths, P. E. 2009: The distinction between innate and acquired characteristics. In E. Zalta (ed.), *Stanford Encyclopedia of Philosophy*, online at <http://plato.stanford.edu>
- Griffiths, P. E. and Gray, R. D. 1994: Developmental systems and evolutionary explanations. *Journal of Philosophy*, 91, 277–304.
- Griffiths, P. E. and Gray, R. D. 2005: Discussion: three ways to misunderstand developmental systems theory. *Biology and Philosophy*, 20, 417–25.
- Griffiths, P. E. and Machery, E. 2008: Innateness, canalization, and 'biologizing the mind'. *Philosophical Psychology*, 21, 397–414.
- Griffiths, P. E., Machery, E. and Linquist, S. 2009: The vernacular concept of innateness. *Mind & Language*, 24, 605–30.

- Jablonka, E. 2002: Information: its interpretation, its inheritance, and its sharing. *Philosophy of Science*, 69, 578–605.
- Jablonka, E. and Lamb, M. J. 2005: *Evolution in Four Dimensions: Genetic, Epigenetic, Behavioral, and Symbolic Variation in the History of Life*. Cambridge, MA: MIT Press.
- Jablonka, E. and Lamb, M. J. 2007: Précis of *Evolution in Four Dimensions*. *Behavioral & Brain Sciences*, 30, 353–92.
- Jablonka, E. and Raz, G. 2009: Transgenerational epigenetic inheritance: prevalence, mechanisms, and implications for the study of heredity and evolution. *The Quarterly Review of Biology*, 84, 131–76.
- Johnson, M. H., Dziurawiec, S., Ellis, H. and Morton, J. 1991: Newborns' preferential tracking of face-like stimuli and its subsequent decline. *Cognition*, 40, 1–19.
- Khalidi, M. A. 2002: Nature and nurture in cognition. *British Journal for the Philosophy of Science*, 53, 251–72.
- Knobe, J. and Samuels, R. in progress: Thinking like a scientist: innateness as a case study.
- Kronfeldner, M. E. 2009: 'If there is nothing beyond the organic ...': heredity and culture at the boundaries of anthropology in the work of Alfred L. Kroeber. *NTM Zeitschrift für Geschichte der Wissenschaften, Technik und Medizin*, 17(2), 107–33.
- Lehrman, D. S. 1953: A critique of Konrad Lorenz's theory of instinctive behavior. *The Quarterly Review of Biology*, 28, 337–363. Excerpts reprinted in Oyama *et al.*, 2001.
- Lehrman, D. S. 1970: Semantic and conceptual issues in the nature-nurture problem. In L. Aronson, E. Tobach. D. S. Lehrman and J. Rosenblatt (eds), *Development and Evolution of Behavior*. San Francisco, CA: Freeman.
- Leimar, O., Hammerstein, P. and Van Dooren, T. J. M. 2006: A new perspective on developmental plasticity and the principles of adaptive morph determination. *The American Naturalist*, 167(3), 367–76.
- Lewontin, R. 1974: The analysis of variance and the analysis of causes. *American Journal of Human Genetics*, 26, 400–11.
- Lickliter, R. and Berry, T. D. 1990: The phylogeny fallacy: developmental psychology's misapplication of evolutionary theory. *Developmental Review*, 10, 348–64.
- Lind, M. I. and Johansson, F. 2007: The degree of adaptive phenotypic plasticity is correlated with the spatial environmental heterogeneity experienced by island populations of *Rana temporaria*. *Journal of Evolutionary Biology*, 20, 1288–97.
- Linguist, S., Machery, E., Griffiths, P. E. and Stotz, K. 2011: Exploring the folkbiological conception of human nature. *Philosophical Transactions of the Royal Society B*, 366, 444–53.
- Lorenz, K. 1937 [1985]: Über den Begriff der Instinkthandlung. *Folia Biotheoretica*, 2, 17–50. Reprinted in C. J. Mellor and D. Gove, trans.; G. M. Burghardt, ed., *Foundations of Comparative Ethology*. New York: Van Nostrand Reinhold, pp. 405–32.

- Lorenz, K. 1965: *Evolution and Modification of Behavior*. Chicago, IL: University of Chicago Press.
- Mameli, M. 2008: On innateness: the clutter hypothesis and the cluster hypothesis. *Journal of Philosophy*, 55, 719–36.
- Mameli, M. and Bateson, P. 2006: Innateness and the sciences. *Biology and Philosophy*, 22, 155–88.
- Mameli, M. and Bateson, P. 2011: An evaluation of the concept of innateness. *Philosophical Transactions of the Royal Society B*, 366, 436–43.
- Maynard Smith, J. 2000: The concept of information in biology. *Philosophy of Science*, 67, 177–94.
- Maynard Smith, J. and Szathmáry, E. 1995: *The Major Transitions in Evolution*. Oxford: Freeman.
- Millikan, R. G. 1984: *Language, Thought and Other Biological Categories*. Cambridge, MA: MIT Press.
- Moss, L. 2003: *What Genes Can't Do*. Cambridge, MA: MIT Press.
- Nijhout, H. F. 2001: The ontogeny of phenotypes. In Oyama *et al.*, 2001, pp. 129–140.
- Northcott, R. 2006: Causal efficacy and the analysis of variance. *Biology and Philosophy*, 21, 253–76.
- Pen, I., Uller, T., Feldmeyer, B., Harts, A., While, G. M. and Wapstra, E. 2010: Climate-driven population divergence in sex-determining systems. *Nature*, 468, 436–39.
- Samuels, R. 2002: Nativism in cognitive science. *Mind & Language*, 17, 233–65.
- Sarkar, S. 2003: Genes encode information for phenotypic traits. In C. Hitchcock (ed.), *Contemporary Debates in Philosophy of Science*. Oxford: Blackwell, pp. 259–72.
- Shea, N. 2007a: Representation in the genome in other inheritance systems. *Biology and Philosophy*, 22, 313–31.
- Shea, N. 2007b: Consumers need information: supplementing teleosemantics with an input condition. *Philosophy and Phenomenological Research*, 75, 404–35.
- Shea, N. 2009: Imitation as an inheritance system. *Philosophical Transactions of the Royal Society B*, 364, 2429–43.
- Shea, N. 2011: What's transmitted? Inherited information. *Biology and Philosophy*, 26, 183–89.
- Shea, N. forthcoming a: Inherited representations are read in development, *British Journal for the Philosophy of Science*, doi: 10.1093/bjps/axr050.
- Shea, N. forthcoming b: Two modes of transgenerational information transmission, in B. Calcott, B. Fraser, R. Joyce, and K. Sterelny (eds.) *Signaling, Commitment, and Emotion*. Cambridge, MA: MIT Press.
- Shea, N., Pen, I. and Uller, T. 2011: Three epigenetic information channels and their different roles in evolution. *Journal of Evolutionary Biology*, 24, 1178–87.

- Stegmann, U. E. 2005: Genetic information as instructional content. *Philosophy of Science*, 72, 425–43.
- Sterelny, K. 2003: *Thought in a Hostile World*. Oxford: Blackwell.
- Sterelny, K. 2006: The evolution and evolvability of culture. *Mind & Language*, 21, 137–65.
- Sterelny, K. and Griffiths, P. E. 1999: *Sex and Death*. Chicago, IL: University of Chicago Press.
- Sterelny, K., Smith, K. C. and Dickson, M. 1996: The extended replicator. *Biology and Philosophy*, 11, 377–403.
- Stich, S. P. 1975: Introduction: the idea of innateness. In S. P. Stich (ed.), *Innate Ideas*. Berkeley, CA: University of California Press.
- Tinbergen, N. 1942: An objectivist study of the innate behaviour of animals. *Bibliotheca Biotheoretica*, 1, 39–98.