FINAL DRAFT

THE RESURRECTION OF INNATENESS

Maclaurin, J.

For Monist (2002)

Introduction

The idea that some biological characteristics are innate, while controversial is widespread in many academic disciplines. Neither philosophy nor science has outgrown the need to talk about traits, which, for a variety of reasons, appear to be inherent in biological populations. Philosophical claims of this nature are to be found in theories of moral sense, rational capacities, the way in which perception structures experience and so on. Scientific claims about innate traits are to be found in the study of animal behaviour and most famously in the relatively recent rise of nativism in cognitive science. In this tradition, Noam Chomsky (1965) and his heirs argue that much of our capacity to decipher verbal information is innate. David Marr (1982) defends a similar position with respect to the interpretation of visual information.

Disconcertingly, given our theoretical reliance on the idea, it turns out that 'Innateness', as a scientific notion, is the subject of considerable scepticism. Work in evolutionary and developmental biology has, for many, called into question the idea that there are such things as innate traits.

I The Trouble with Innateness

In *What is Innateness* (this issue), Paul Griffiths cogently sums up current concern over the notion of innateness. In short—different people mean different things by it. This should come as no surprise as the concept has been employed for many different theoretical purposes and, in the process, it has wended its way into 'folkbiology'. The origins of this biological scepticism regarding innateness appear in Daniel Lehrman's (1953) *Critique of Konrad Lorenz's Theory of Instinctive Behaviour*. The aim of Lehrman's paper was to demonstrate that the notion of innateness, as used in ethology (the study of animal behaviour founded by Lorenz), is confused because there are many ways in which the term can be interpreted. We can think of innate behaviours as:

- Traits that lack malleability;
- Traits that are characteristic of particular species;

- Traits that are evolutionary adaptations.
- Traits that are unlearned;
- Traits that develop in the absence of contact with conspecifics;
- Behaviours that develop fully formed in animals that have been prevented from practising them;

Many of these distinctions have their theoretical uses, but crucially they do not pick out the same groups of traits. Many birds have song patterns characteristic of their species, yet those behaviours do not develop unless the chicks are hatched by birds of the same species (Gould, 1991). Thus, that behaviour is innate in one of these senses, but not in another.

Lorenz's response to this criticism was to argue that while the identification of innate traits may be confusing, the notion of innateness itself is not confused. That is because it rests upon a single property that is true of all innate traits, namely that they are genetic (1950, p. 261). On its own, this claim is somewhat ambiguous. All traits have some genetic causes and none have only genetic causes. What then does it mean to say that 'innate traits are genetic traits?' Lorenz's answer was that we should read 'genetic trait' not as 'trait with a genetic cause' but as 'trait that is a product of genetic information' (1966, p. 37).

Finally, if we cast our gaze beyond the academy, we see, as Griffiths suggests, that colloquial use of the term often expresses one of a number of popularly believed biological falsehoods. Sometimes this is the biological determinism that sees innate traits as *genetic and thus immutable*. Sometimes it is the view that innate traits constitute the essential 'nature' that is common to all members of a species. Sometimes it is the view that innate traits are those characteristics that an individual is *supposed* to have. The vagaries of folk-biology are indeed fascinating. However, I suggest that they need not concern us in discussion of future scientific use of a term such as 'innateness'. Physicists do not after all allow folk-physics to colour their interpretations of terms such as 'force' and 'weight'. This said, even if we ignore folk-biology, there is still a good deal of scientific confusion stemming from the many different characterisations of innateness. Does this mean, as Griffiths suggests, that we should see 'innateness' as an antiquated theoretical construct that has had its day?

Perhaps the first thing to note is that the mere presence of confusion surrounding a theoretical term has not in the past constituted sufficient ground for abandoning its scientific use. Many important scientific terms have at some time been subject to a great deal of theoretical confusion. In the past, there have been many schools of thought as to how we ought best to characterise terms such as 'electricity', 'atom' and 'planet'. The mere fact of confusion attendant upon theoretical revision has not caused us to eliminate

those technical terms and it should not, on its own, cause us to eliminate 'innateness'.

If we would demonstrate that there is still a place for scientific (and philosophical) talk of innateness, we must show that 'innateness' still has a theoretical job to do and we must show that there is some clear way of characterising some useful property to which the term refers. We might further note that playing a valuable theoretical role appears to be more important than being subject to a single clear characterisation. Despite being the subject of heated debate for many decades, the term 'species' is in no danger of being abandoned. It is true that biologists working in different areas mean quite different things when they use the term (Claridge *et al* 1997). It is also true that there are cogent arguments for thinking that we will be unsuccessful in our efforts to find a single characterisation that will prove suitable to all purposes (Hull, 1997). Nonetheless, species, is such an important taxonomic rank that we will undoubtedly choose to put up with the confusion of retaining a widely used technical term that lacks a singular definition.

In this paper I will argue that 'innateness' does indeed have theoretical work to do. I will further argue that there is a clear way in which we might characterise 'innateness' that will capture a property relevant to the task at hand. In this respect I suggest that 'innateness' is somewhat better off than 'species'.

What is the Point of Talk About Innateness?

As noted above, many scientific and philosophical enterprises have made use of the notion of innateness. With this diversity of enterprises, comes a diversity of senses in which the term 'innateness' is employed. Furthermore, I agree with Griffiths that many of these enterprises now make use of newly discovered technical explanations of mechanisms that maintain particular traits in particular populations. In light of these developments Griffiths gives the following prescription for eliminating talk of innateness:

> If a trait is found in all healthy individuals or is pancultural, then say so. If it has an adaptive–historical explanation, then say that. If it is developmentally canalised with respect to some set of inputs or is generatively entrenched, then say that it is. If the best explanation of certain trait differences in a certain population is genetic, then call this a genetic difference. If you mean that the trait is present early in development, what could be simpler than to say so?
>
> Griffiths (this issue p.?)

Why, asks Griffiths, should we cling to the notion of innateness when we can now talk much more precisely? My response is to note that just because we can disaggregate a

general concept, such as 'innateness', does not mean that it is always appropriate to do so.

Radical new biotechnologies provide us with the opportunity to erase, enhance or modify characteristics of individuals and populations that until now seemed to be unchangeable consequences of phylogenetic history. Of course, such modifications can potentially be extremely deleterious. For this reason, it is very important that prior to embarking on such projects, we have a clear understanding of the processes which have caused the characteristics in question to become built-in to extant populations[1]. Griffith's putative justification for out elimination of theoretical use of 'innateness' is based on the idea that there is a wide variety of potential explanations for why traits become fixed within populations. However, crucially the elimination strategy discourages us from asking a very important question—why is it that trait *X* has become fixed within or built-into lineage *Y*? To frame that question clearly we require a general term to replace the rather clumsy talk of traits being 'fixed within or built-into' a lineage. I suggest that 'innate' is the obvious term to choose. Indeed there are many questions that would be aided by thinking of innateness in this way. For example, we might also want compare the speed and power of different mechanisms that produce innate traits. We might even want to ask why it is that some apparently possible mechanisms for producing innate traits do not in fact operate. Thus, I claim that it is in the context of this sort of general enquiry that disaggregating a notion such as 'innateness' does not satisfactorily serve our explanatory purposes. What we need is an 'umbrella term' that picks out a group comprising of many different phenomena in just the same way that 'electricity' now picks out a wide variety of electrical phenomena.

So my proposal is that we reinvent 'innateness' as an umbrella term picking out all traits that are, by whatever mechanism, built-in to populations. The task of providing some good characterisation of just what this property is, will be pursued at length in the remainder of this paper. Of course, if 'innateness' is to be an umbrella term, we will have to ensure that our characterisation of it is suitably agnostic on many of the issues mentioned so far. It must accept traits that are built into populations by a variety of different mechanisms (eg, generative entrenchment, heterozygote superiority etc.) and it must accept traits that are built into populations in a variety of different ways. This might include traits that are not common to all the members of a population (you might think of Huntingtons as built-in to human populations even though it is rare) or traits that are malleable in the course of development (you might think of language as built-in to human populations even though it is not developmentally fixed).

---

[1] We will often be interested in traits that are common in populations, however, the type of analysis I develop will apply just as well to groups smaller than whole populations or to larger taxonomic entities such as species or higher taxa.

In section three of this paper, I shall canvas a number of suggestions that seek to reinvent the term innateness, avoiding the ambiguities mentioned in the previous section. Briefly, these solutions fall into two groups. One camp, including Stephen Stich (1975) and Elliot Sober (1998) argue that we should interpret claims about innateness as claims about dispositions. Thus, for them, '*X* is innate in *Y*' means approximately that '*Y* has a disposition to develop *X*'. The other camp, including André Ariew (1999) and Bill Wimsatt (1999) argue that talk of innateness could be made respectable once more if we were to identify that property with a specific biological process. The view that I will develop in this paper stands midway between those two camps. I think innateness can be characterised so as to refer to some set of biological processes, all of which cause traits to become 'built-in' to evolving lineages. It should also give us some indication of what it is that such processes have in common and thus of how we might profitably compare them. Thus, I see Stich's and Sober's accounts as too 'thin'. Conversely, given that I think we can characterise a set of processes which produce innate traits, I see Wimsatt's and Ariew's single process accounts as too 'thick'. Arguments in support of my 'medium width' account of innateness are supplied in section three. The intervening pages give a detailed argument to the effect that there is indeed a principled way of characterising a set of biological processes, all of which produce 'built-in' traits.

My suggested solution to the problem of characterising this new general notion of innateness is a descendant of the claim by Konrad Lorenz that all innate traits are the products of genetic information. In what follows, I shall argue that innate traits are products of information provided by developmental resources[2] that are maintained in biological populations by a variety of mechanisms. It is the persistence of these resources that maintains innate traits in populations over many generations. Unlike Lorenz, I shall not insist that the developmental resources in question should be genetic. Just as there is a variety of developmental resources that underpin innate traits, so there is a variety of biological mechanisms acting to stabilise evolving populations, maintaining the resources in question.

Followers of this literature will know that in the last thirty years much ink has been spilt in an effort to show Lorenz and his followers that theirs is not in fact a promising solution to the problem (see for example Bateson (1976), Griffiths and Gray (1994), and Oyama (1985)). Indeed, some have suggested that the whole idea of a genetic trait is conceptually suspect (see for example Lewontin (1985, 1991)). As my characterisation of 'innateness' is similar to Lorenz's, some of those arguments will have to be addressed in this paper.

---

[2] A developmental resource is any factor that influences the development of a biological individual. These include genetic, nutritional, cytological, embryological and behavioral factors.

The most important hurdle facing my characterisation of innateness is that it is theoretically difficult to single out particular developmental resources as sources of information for particular traits. In the next section, I set out the problem and propose a solution. I shall do so by initially evaluating the Lorenzian idea that innate traits are products of genetic information. It will allow me to set out the way in which a set of developmental resources can be thought of as a source of information for a biological trait and to evaluate some of the more radical claims about the theoretical importance (and indeed the theoretical plausibility) of the notion of a genetic trait. It will also allow me to explain why I choose not to equate innate traits with genetic traits. I shall leave behind talk of genetic traits and focus more closely upon innateness in section three.

The Nature of Genetic Information

There are two ways to characterise information. One is to treat it as a semantic property of intentional systems. I can tell you what bits of information on my computer *are about* because they were put there by an intentional agent. So I can say which files contain pictures of whom and which contain ideas about what. That will not work as a characterisation of genetic information because genomes are configured, not by intentional agents, but by evolution.

A more promising idea is explored by Sterelny, Smith and Dickerson (1996). They suggest that a portion of the genome carries semantic information about some trait just if it has evolved for the purpose of producing that trait. So, this notion of information underpins the idea that the genome can be a source of developmental information about a trait, while environmental factors necessary for the successful development of the trait may nonetheless not be sources of developmental information about that trait.

In the current context however, this spelling-out of a semantic notion of genetic information has two drawbacks. Lorenz suggested that we use the notion of genetic information as a means of characterising innate traits. But, consider a group of traits well-known from studies of human visual perception. Human beings are subject to a variety of visual illusions. That is to say that all humans misinterpret certain types of visual data. These perceptual failings are not adaptations. They are simply by-products of the evolutionary process—trade-offs between cost and efficiency. If this is true then, there are no genetic, cytological, behavioural factors… that have evolved for the purpose of producing these traits. Thus, on the semantic account of information discussed here, there is no set of developmental resources (genetic or otherwise) that is a source of information about the development of these traits. Therefore, a notion of innateness that is based on this characterisation of information will not see these traits as innate. That seems

undesirable, given that these traits appear to be built-in to our species in just the same way as our perceptual adaptations. A second problem for the semantic approach is that it gives us no means by which to quantify the extent to which different developmental resources are informationally responsible for the development of particular traits. I shall say more on this later.

A very different account of information is spelled out in the mathematical information theory developed by Shannon and Weaver (1949). According to that theory:

- Information has flowed from *A* to *B* if the number of possible states of affairs at *A* decreases given some state of affairs at *B*. Thus, receiving the signal at *B*, reduces our uncertainty about states of affairs at *A*. Thought of in this way, tremors can signal volcanic eruptions and sniffles can signal colds.
- Signals travel from sources, via channels, to receivers.
- A channel is noisy if it contributes information to the receiver that doesn't come from the desired source. Static on a radio conveys information about something. Perhaps it's an approaching storm, perhaps a fault in the radio's wiring. It doesn't though convey the information we want, ie, the information being transmitted by our preferred station.
- A signal is equivocal when it transmits only partial information about states of affairs occurring at a source. When you manage to take down half a number plate; when you hear a distant but unidentified rumbling; when your car just doesn't go as well as it used to…. In all these cases, you are in receipt of equivocal information.

It is this characterisation of developmental information that I shall employ in my characterisation of 'innateness'. For reasons that will become clear later, it is important to note that this type of informational analysis is not profitably thought of as a causal analysis[3]. My television set conveys information (albeit equivocally) about what is on my neighbour's television set. Despite this, my neighbour's television set is not a cause of the information displayed on my television set. Indeed, I can know all there is to know about the informational relationship that holds between our two sets without knowing anything about the wider set of causal relationships which synchronise the two systems.

In this sense, informational analysis is not the same as causal analysis. This is not to say that informational relationships are somehow non-physical, or that they are not underpinned by ordinary causal relationships. My claim is simply that I can have the full informational picture of the system comprised of the neighbouring televisions, without having the full causal picture of that system. If this paper is right, there is value to be had

[3] Note: some authors label this type of information 'causal information' in an effort to distinguish it from semantic information (see for example Sterelny and Griffiths (1999, p. 101))

in an informational analysis of the relationship between developmental resources and phenotype. This said, let us now return to the question of whether this style of informational analysis would have provided Lorenz with a useful characterisation of innateness.

At first glance, it appears promising. Innate traits, Lorenz could have said, result from genetic signals. The development of such traits requires the presence of certain environmental factors, but these can be thought of as mere channel conditions. The genome provides the signals; any information from the environment is just noise.

The parity principle

In *Developmental Systems and Evolutionary Explanation,* Paul Griffiths and Russell Gray (1994) ask, 'Why see it that way'. They point out that mathematical information theory was not designed to (and is not capable of) individuating the sources of particular pieces of information. That is to say, it has no principled means of distinguishing signal from noise. By its lights, what we count as signal and what we count as noise is a fact about us, not about the systems we analyse.

So, when I refer to atmospheric conditions interfering with my television set, it would be just as correct (according to the theory) to refer to tonight's television news bulletin as interfering with my "atmospheric conditions detector". To say that some information is signal is to say that it's relevant, and no more. Thus, if it is possible to see the environment as a channel across which genetic information flows, then mathematical information theory will allow us to invert this analysis and so see the genome as a channel across which environmental information flows. This is the parity principle and it tells us that if we treat the genome as a source of developmental information, then we have to accord the same status to environmental factors required in development.

So, for Griffiths and Gray and for Susan Oyama (Oyama 1985), the problem with analysing biological development in terms of information is that it works too well. There are many and varied sources of developmental information and no way to pick out which are the most important. A consequence of this is that they reject the dichotomy that sees all traits as either wholly genetic or wholly environmental. They think it makes no sense to talk in these terms.

A much stronger conclusion, and I think a much more dubious one, is the rejection of the idea that we can ever say of a trait that it is more genetic than some other trait. Griffiths, Gray and Oyama all use this conclusion to licence their rejection of claims that innate traits are genetic traits. Their justification for this conclusion (explained in detail below) is based on the fact that 'genetic traits' are in fact the products of interactions

between a vast array of genetic and non-genetic developmental resources. Indeed, as we shall discover, their argument appears to spell trouble for any notion of innateness (such as my own) based on the idea that innate traits are the products of information provided by particular developmental resources which persist from generation to generation.

II Are some traits more genetic than others?

The parity principle has apparently cost us explanatory power. It seems to leave us no way to assess the relative importance of any multiple sources of information that together produce a particular phenomenon. This may be inevitable but it has some disconcerting consequences.

After a talk by Susan Oyama at the 1999 ISPHSSB[4] meeting, David Hull commented that the problem with the Griffiths/Gray/Oyama prohibition on talk about traits being genetic, is that we tend to lose the idea of variance. We've lost the idea that some traits differ in the extent to which they are buffered against environmental change (or for that matter against differences in genotype).

The problem is most apparent when we try to make sense of the medical use of the term 'genetic'. The medical community thinks of genetic disorders as heritable medical conditions whose presence is signalled by a characteristic mutation, or by a mitochondrial or chromosomal abnormality. The genetic abnormality is not present in individuals without the condition and thus, some claim that the condition is *caused* by the abnormality in question.

The problem with this way of characterising genetic disorders is that mutations are not the only cause of the symptoms of genetic disorders. Like all biological traits, the symptoms of genetic disorders are the products of a vast array of interacting causes. Furthermore, we now know that the symptoms of a variety of genetic disorders can be ameliorated by changing the environment of the individual with the mutation in question. PKU (phenylketonuria) is a metabolic disorder caused by the absence of the enzyme phenylalanine hydroxylase, preventing the body from converting phenylalanine into tyrosine. Untreated, PKU causes a variety of symptoms, most notably mental retardation. The discovery that the symptoms are the result of the failure of a metabolic pathway means that the disorder can now be treated by putting the patient on a diet low in phenylalanine.

What this tells us is that the symptoms of PKU are best thought of as the result of interaction occurring between a variety of genetic and non-genetic developmental

---

[4] The 1999 meeting of the International Society for the Philosophy, History and Social Studies of Biology in Oaxaca (Mexico).

resources. Indeed, this is true of all genetic disorders, although of course many are not similarly responsive to therapeutic intervention. Thus, if we think of the notion of a genetic trait as conceptually suspect, then it seems we ought to think of the notion of a genetic disorder as suspect for the same reasons. These disorders are caused by genetic and non-genetic factors. If the genome can be said to carry information about the symptoms of these disorders then so can environments in which the disease would develop. So, while we can of course give a description of the discovery that PKU can be treated by changing the patients nutritional environment (e.g. in terms of malleability), we cannot describe it in the familiar terms that many would choose. That is to say, if we agree with Griffiths and Gray that traits do not differ in *the extent to which* they are genetic, then we cannot take ourselves to have learned that PKU is not a wholly genetic disorder or that it is less genetic than we thought it was.

In what remains of this paper I will argue that the problem here is not the parity principle, but in the conclusions that we draw from it. I will suggest that the explanatory power, to which Hull referred, is still within our grasp if we pay more attention to the nature of developmental information.

More genetic traits are less equivocal traits

I agree with the parity principle. Both the genome and the environment provide information to the developing individual, in just the same way. But I disagree with the idea that it makes no sense to speak of some traits as being more genetic than others. I think this problem can be solved by recruiting the information theory notion of equivocation. In so doing, we will come one step closer to making sense of the notion of innateness.

Jaundice (the yellowing of body fluids and tissues due to excess biliruben) is a symptom of a number of medical conditions including some types of anaemia, cirrhosis, hepatitis and RH factor reaction in newborns. Fever is similarly a symptom of disease but it's much more equivocal. Many medical conditions result in fever. Thus, one type of symptom is less equivocal than the other. It achieves a greater reduction of uncertainty. There are many such examples. We are happy to have scientists talk of signs of an impending earthquake, signs of an impending economic depression etc. We are also happy that some of these signs will be more equivocal than others.

If symptoms of disease can vary in the extent to which they are equivocal, why not think that developmental resources can differ in the extent to which they are equivocal as sources of information about phenotypic outcomes? With this in mind, I claim that trait *A* is more genetic than trait *B* if the genome is a less equivocal source of information

concerning the development of *A* than the development of *B*. So for example your genome is a very equivocal source of information about the language you will speak. It is a much less equivocal source of information about the number of vertebrae you will have.

So, what I am proposing is that *with respect to individual traits* and particular sets of developmental resources we can ask 'how equivocal is this set of resources as a source of information about the development of the trait or traits in question?' Given this, it now appears that we can accept the parity principle and still think that some traits are more genetic than others. At least we would be able to if we had some definite proposal as to how it is that we compare the extent to which different developmental resources are equivocal as sources of phenotypic information. *Prima facie*, a suitable solution seems to be provided by the analysis of variance. However, that approach is potentially problematic for reasons first set out by Richard Lewontin in the mid-seventies.

Lewontin on the Analysis of Variance

In 'The Analysis of Variance and the Analysis of Causes' (1974), Lewontin provides a critique of the analysis of variance much utilised by researchers in human genetics. To understand the problems with such analyses, we need first to understand their purpose.

Analysing the relative importance of two causes that contribute to produce a single phenotypic result is only possible if the causes in question are truly independent. Lewontin gives the following analogy:

> …if two men lay bricks to build a wall, we may quite fairly measure their contributions by counting the number laid by each; but if one mixes the mortar and the other lays the bricks, it would be absurd to measure their relative quantitative contributions by measuring the volumes of bricks and of mortar.
>
> (1974, p. 402)

He goes on to say that it is just as absurd to ascribe so many inches of a man's height to his genes and so many to his environment. This fact about causal analysis seems to frustrate the project of determining the relative importance of genes versus environment in the development of phenotype. In response, scientists have turned from causal analysis to the analysis of variance. This purports to tell us what proportion of some trait's deviation from the population mean is due to deviation of the individual's environment from the mean environment on the one hand or deviation of the individual's genetic value from the mean genetic value on the other.

Having acquainted us with the purpose of analysis of variance, Lewontin points to two important problems. One is that researchers mistakenly interpret variance data as allowing us to apportion *causal* responsibility. This is simply failing to understand the point of the bricklayer example. However, it need not concern us here, as our project is informational rather than causal analysis.

A second problem is that researchers occasionally lose sight of the fact that analysis of variance is, by its nature, a local analysis. It tells us about the effects of changes in environment on a particular population with a given genetic make-up. Or it tells us about the phenotypic consequences of different genotypes in a given environment. It does not tell us about global relationships that hold between genetic and other developmental resources that collaborate to produce particular phenotypic outcomes. These global relationships are expressed in graphs of norms of reaction. Figure 1 (taken from Lewontin 1974) shows a variety of norms of reaction. Crucially, what it shows is that if we simply hold the environment constant and investigate the effects of different genotypes on phenotypic outcomes, the results will be very dependent on which environment we choose. The converse is also true.

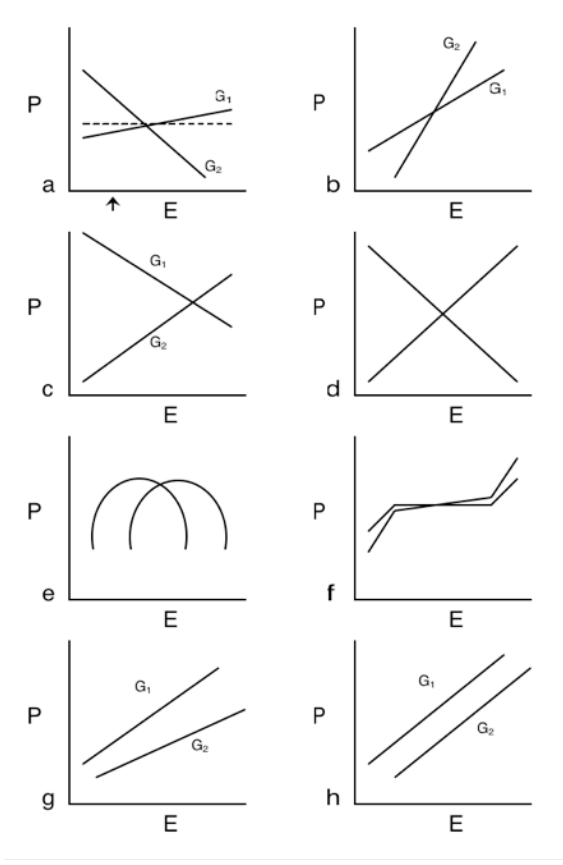

Fig. 1 Examples of different forms of reaction norms (after Lewontin (1974)). In each case the phenotype (P) is plotted as a function of environment (E) for different genotypes (G_1,

G_2).

So does the move from analysis of variance to norms of reaction prevent us from drawing conclusions about the extent to which genotype (or alternatively environment) is a source of information about phenotypic outcomes? It is a very important question, for if Lorenz's notion of innateness fails here, then so does my own. After all, my claim is just that there is some set of developmental resources that is highly correlated with innate phenotypes (Lorenz claims that there is some set of purely genetic resources that is highly correlated with innate phenotypes). As it happens, the answer to this crucial question is—it depends on which graph we are talking about.

Graph h in figure one illustrates perfect additivity between genetic and environmental causes of development. The phenotypic consequences of differences in genotype are the same in all environments. Likewise, the phenotypic consequences of differences in environment are the same for both genotypes. So in this case there is a clear and global correlation between genetic factors, environmental factors and phenotypic outcome. Thus, it seems uncontroversial to claim that it is telling us about the extent to which genetic or environmental factors are a source of information about phenotypic outcomes. After all non-semantic information just is correlation.

But herein lies the problem. Perfect additivity is the exception, not the rule. Indeed, this is precisely Lewontin's point. All the other graphs in figure 1 represent a lack of additivity between genetic and environmental factors that contribute to developed outcome. Technically speaking, we cannot partition the variance for the traits in question into a genetic and an environmental component. Put in the language of Lewontin's example: in most cases, genes and environment are not independent causes of phenotype. They are interacting causes of phenotype in the way that the bricklayer and the mortar mixer were interacting causes of the wall.

I said above that Griffiths and Gray reject even the claim that some traits are more genetic than others. Now we have the explanation. Almost all interaction during biological development is non-additive to some degree or other. That is why they think that we cannot say that skeletal structure is *largely due to* our genetic make-up, while our taste in literature is *largely due to* environmental factors.

As regards innateness, the same point has been put by Elliot Sober. He says that the contributions of genes and environment are 'not made in a common currency' and thus they cannot be compared (1994). But Sober's claim is expressly a claim about apportioning *causal* responsibility. So the question now becomes—Does lack of additivity prevent assessment of the extent to which different factors are equivocal *as sources of information* about developed outcomes? I claim that it does not.

The first thing to note here is that, from the point of view of informational analysis,

additivity is not an "all or nothing" property. Graph h depicts a perfectly additive relationship. Graph g shows an *approximately* additive relationship between a set of genetic and environmental factors. This, Lewontin acknowledges (1974, p. 408). So, knowing this, can we still read off facts from graph g about the extent to which phenotypic outcome is correlated with either genetic make-up or environmental factors?

The answer to this question is 'yes' but there is a qualification. We can see this by inspecting graphs g and h. In graph h, the extent to which genotype is a source of information about phenotype is represented by the distance between the lines G_1 and G_2. The extent to which environment is a source of information about phenotype is represented by the slope of the lines G_1 and G_2. In Graph g, the lack of additivity limits the precision with which we can read off facts about the correlation between developmental inputs and phenotypic outputs. In effect, lack of additivity introduces an element of error into our analysis. Thus in graph g, the most precise claim we can make is that the extent to which genotype is a source of information about phenotype is represented by the average distance between the lines G_1 and G_2 ± half the maximum distance between the lines. Correspondingly, the extent to which environment is a source of information about phenotype is represented by the average slope of the lines G_1 and G_2 ± half the maximum difference in slope.

We should note that this simple analysis of graph g is one that would not work on other data sets. Talking about the average slope makes sense for graphs such as g, h, the centre portion of f, and perhaps even b (although the "error" associated with b would be large). However, we would learn little about development by trying to calculate the average slope of the lines in graph e whose curved lines of graph e draw our attention to an important issue. The mathematics of determining the extent to which a data set falls away from additivity is, in reality, much more complex than merely taking an average and calculating an error. However, philosophically we need not be concerned at the complexity of the mathematical details. All that matters in the current context, is that there is a mathematical means of spelling out what we mean when we say that a relationship is somewhat additive, very nearly additive etc. The crucial conclusion to draw here is that lack of additivity does not preclude us from investigating the extent to which sets of developmental resources are correlated with particular phenotypic outcomes. What it does do is to limit the precision with which we can investigate these relationships. To be sure, there are data sets (for example that represented in graph d) that are entirely lacking in additivity. In these cases we cannot single out genes (or indeed any other set of developmental resources included in the data set) as having important informational input in the development of the trait in question. My aim here is just to show that we can sometimes apportion informational responsibility even in cases where the relationship between the contributing factors is not perfectly additive.

What of Lewontin's other worry? Analysis of variance is sometimes interpreted as providing information about global relationships which hold between developmental resources and phenotypic outcomes. But norm of reaction graphs appear to show us that such relationships are often local. How can we say that a genome (or indeed any set of developmental resources) carries information about some phenotypic outcome, when the extent to which it is correlated with changes in that outcome differs markedly in different environments? Again, this is crucial to the claim that I want to make about innateness. I am suggesting that innate traits are products of information from special groups of developmental resources, but no set of developmental resources can produce the same innate trait in absolutely any environment whatsoever and that fact appears to go against the idea that these resources are really sources of information about the traits in question. The solution to this problem is that informational analysis *is* local analysis.

The property of being a source of information

To see why, we have to look more closely at what we claim when we say that some system is a source of information. When I say that my radio is a source of information about distant systems I am not saying that states of my radio will tell me about distant systems in all possible worlds. There are many worlds in which such a device would not function. Conversely, I am not merely making a claim about the actual world. Tomorrow I may wear a black jumper or a red one. I certainly don't think that will affect whether or not my radio continues to work. Finally, I'm not making a claim about most possible worlds. In most possible worlds my radio doesn't exist. What I am doing is limiting my claims to a particular set of possible worlds that I take to be relevant.

The same is true when I assert some particular medical significance to yellowing of the whites of the eyes. Jaundice signals a particular set of medical conditions, but it doesn't of course do so in all possible worlds. Similarly, signs of an impending earthquake provide us with information but only in a restricted set of possible worlds. All these cases are instances of one system providing us with information about another, but crucially that claim does not imply that the correlations involved hold globally (or if you like, across all possible worlds). This is just to say that, philosophically speaking, the property of being a source of information is a dispositional property[5], just like properties such as being dangerous or being fragile. It is a property whose effects are felt only in some particular

[5] Note: this style of analysing dispositional properties, while common, is not universally accepted (see for example Mellor (2000)). However, fortuitously, the subtleties of this debate are not relevant to the current issue. What matters here is only that we accept that being a source of information is a claim about a particular set of relevant possible worlds.

set of relevant possible worlds.

Thus, when I assert that system *A* is a source of information about system *B*, what I am suggesting is that *A* and *B* are correlated, not in all possible worlds, but in some local set of relevant worlds.

To come back to the biological case, norm of reaction graphs tell us that correlations between sets of developmental resources and phenotypic outcomes are likely to be local. If we accept that being a source of information is a dispositional property then we can still maintain that these relationships are informational. Crucially, in the biological case, we will have to be able to spell out which environments will be included in the set of relevant local environments.

Before I go any further it is important to clear up a potential ambiguity, viz, what is meant by 'local' in the above explanation? The term 'local' here simply means 'in some specifiable set of developmental circumstances'. So, I am not suggesting here that 'local' means 'geographically local' or that it means inhabiting extremely similar environments. So, I think we can say that universal grammar is innate in our species even though humans occupy an extremely varied set of environments and the developmental histories of some individuals (for example those who suffer severe brain trauma) include factors which thwart the development of the trait in question. Of course, that claim rests squarely on the ability of the researchers in question to spell out why it might be that certain bizarre environments (or perhaps even genotypes) do not count as relevant for the purposes of the investigation in question.

I should also note at this point that I have, up to now, written of traits being built into populations. Actually, this is a simplification. We will often be interested in traits that are common in populations, however, the type of analysis I have discussed is not limited to any particular group of organisms as the unit of study. It will apply just as well to groups smaller than whole populations or to larger taxonomic entities such as species or higher taxa. All that matters from my perspective is that the group under study be relevantly similar with respect the developmental resources that give rise to the traits under study. I will say more about this issue in the next section.

In summary—I agree that developmental resources are often not independent in a way that would allow us to determine the extent to which they are *causal* contributors to phenotypic outcomes. However, I argue that this does not preclude us from determining the extent to which they are *informational* contributors to phenotypic outcomes. This is because lack of additivity can, at least in some cases, be accounted for as error and because the local nature of relationships that hold between sets of developmental resources and phenotypic outcomes is consistent with the idea that informational analysis is inherently local. I conclude that this allows us to evaluate claims about the informational importance of particular sets of developmental resources to the production

of particular phenotypic traits. That in turn allows us to investigate the extent to which traits are or are not genetic.

Does this account of developmental information give any hope to nativists who claim that there is some set of innate traits or to the heirs of Lorenz who claim that we can cash out 'innate trait' in terms of genetic information?

III The Resurrection of Innateness

Was Lorenz right?

Lorenz argued that imprinting in ducklings is innate because the genome is the source of the information that produces it. In light of the principle of parity, this is clearly a lop-sided picture. There are many non-genetic factors that will covary with imprinting behaviour. Ducklings will not imprint if they are subject to experimental conditions under which they are so deprived of perceptual input that there is nothing for them to imprint on. Despite this I think we can give sense to Lorenz's claims that imprinting is genetic and thus innate.

As regards individuals reared under bizarre experimental conditions, I think we are entitled to argue that such cases are not relevant in assessing whether or not some trait is innate. Informational analysis is local. To claim that imprinting is affected only by differences in genetic make-up is not to claim that this would be true in all possible environments. Rather it is to claim that it is true of some set of relevant environments. Thus, an ethologist could argue that their claim is directed at the informational status of the genomes of ducklings reared in the wild. Perhaps opponents of innateness would respond that this is an *ad hoc* delimitation of the set of relevant environments. Many developmental biologists have after all worked on the assumption that any environment found to be capable of influencing the development of a trait, shows that that trait is not innate[6].

I think such arguments fail. In other sciences, we leave it up to the members of the discipline concerned to cash out their own dispositional properties. In part, this consists of their telling us what they take to be the appropriate set of relevant possible environments. An interest in the developmental trajectories of animals in, and only in, their normal environment is not arbitrary. Nor for that matter is an interest in the developmental trajectories of animals in, and only in, zoos. So, ethologists, if they choose, could use "innateness" as a theoretical term describing the trajectories of individuals in such clearly

---

[6] See for example Gottlieb (1971) and Bateson and Vauclair (1975).

delimited groups of organisms.

Am I suggesting in all this that Lorenz may yet turn out to be right in his characterisation of innate traits as genetic traits? None of what I have said is intended to return us to the conceptual environment of the 1950s in which ethologists felt entitled to presume that traits which display the hallmarks of an innate trait (mentioned at the beginning of this paper) are genetic traits? Such confidence turns out to be misplaced for two reasons.

In part, Lorenz's program failed because he misunderstood the nature of information. He took information to be a monadic property, something that could inhere in the genome. Developmental information is not like this. Rather, it is a very complex set of relations that holds between a vast array of developmental resources and phenotypic outcomes.

Secondly, it was well documented long before the 1950s that some traits that we would now see as carrying all the hallmarks of innateness are nonetheless subject to clear influence from environmental factors. Kinder (1927) demonstrated that nest-building activity in rats is inversely correlated with temperature despite in other respects appearing to be innate. We now know that apparently innate traits, whose development requires stereotypical environmental input, are in fact common. Given this it appears pigheaded to cling to the notion that only genetic resources can be necessary developmental precursors to innate traits. Innate traits are expressions of phylogenetically inherited information, but not all such information is genetic

The consequences for nativism

I began this paper by setting out an objection to the view that we can divide traits into those that are innate and those that are acquired. Sometimes natural selection acts so as to ensure that certain skills are reliably learned in successive generations. Therefore, if we think of innate traits as products of our phylogenetic inheritance, then these traits are both innate and acquired. Given this, how can we make sense of, for example, Chomsky's claim that universal grammar is built into me in a way that English is not?

The Lorenzian answer to this question would have been to say that universal grammar is genetic. My answer is to argue that universal grammar is the product of a set of developmental resources which is maintained in human populations by some mechanism or mechanisms. But herein lies a problem. While this may well be true of universal grammar, it is also true of natural languages such as English. The set of resources underpinning the speaking of English is of course cultural. However, I have just argued that there seems no good argument to justify the claim that only genetic resources

can be sources of developmental information for innate traits.

This leads to two major differences between my account of innateness and those that I would have it replace. On my account, acquired traits can be innate. Breaking my thumb at the age of 8 is not innate as there is no mechanism in my population that regularly produces such outcomes in my lineage. However, there are mechanisms that maintain all sorts of behavioural characteristics in human populations. So on my story, highly stereotypical, inherited cultural phenomena (such as religious beliefs) could turn out to be innate.

The second major difference is that I claim that traits vary in the extent to which they are innate. None of the traits I have just listed are as innate in my as having a four chambered heart. The mechanisms that maintain that trait in human populations are much less equivocal sources of developmental information than the cultural mechanisms which maintain continuity in human behaviour.

So, for these reasons I am forced to recast the claim made by nativists. On my account, what Chomsky has discovered is that universal grammar is much more innate than the speaking of particular natural languages and it is a product of developmental information provided by a different set of developmental resources than those which underpin particular natural languages.

We may further hypothesise that the set of developmental resources in question is wholly, or in large part, genetic. But that is a further question and it is an empirical question. All I require for the existence of an innate trait is that there exists within the population some mechanism or process that maintains the developmental resources which very reliably produce the trait in question.

Some Alternative Suggestions

Obviously, the questions I raise in this paper have been tackled many times before and one group of proposed solutions is, in some respects, similar to my own. So, in closing I want to very briefly outline some alternate solutions to the problem and to say why I take them to be less satisfactory than the solution I propose.

Bill Wimsatt (1999) has recently argued that innateness is caused by generative entrenchment. This is the evolutionary process by which initially unimportant phenotypic characteristics become locked into lineages, as characteristics crucial for survival evolve so as to be developmentally dependent upon them. I see no reason for thinking that all innate traits will be the product of this process. André Ariew has recently proposed an account which seems similar to that of Wimsatt. His account is based on C. H.

Waddington's notion of developmental canalisation. Waddington famously described the process of biological development as being analogous to a ball rolling down a slope that is folded into hills and valleys. Ariew claims that:

> To reflect the significance of development, on the canalisation account innateness is treated as a matter of degree: the greater the environmental range against which a developmental pathway is buffered, the more canalised is the developmental pathway. The steepness of an entrenched pathway (or "chreode") in Waddington's representation of development represents the degree to which a pathway is canalised.
>
> (1999 p.129)

This appearance may be illusory. Waddington's account provides us with an apt and useful metaphor for the way in which a multitude of constraining factors influence biological development. That said, it is still a metaphor. Nobody thinks there are actually balls rolling down slopes in the course of any given episode of biological development. If the metaphor is cashed out, presumably as a specific developmental mechanism, then my response is the same as to Wimsatt: there is no reason to think that all innate traits are a product of a single mechanism.

Generative entrenchment (and perhaps canalisation) clearly is a mechanism that would maintain appropriate developmental information within a population from generation to generation. Thus, I can agree with Wimsatt at least in thinking that many innate traits will be the products of this process.

Stephen Stich gives an account of innateness in terms of dispositions:

> A person has a disease innately at time *t*, if and only if, from the beginning of his life to *t* it has been true of him that if he is or were of the appropriate age (or at the appropriate stage of life) then he has or in the normal course of events would have the disease's symptoms
>
> (1975, p. 6)

Elliot Sober adds in the idea of a relevant range of environments:

> [A] phenotypic trait is innate for a given genotype if and only if that phenotype will emerge in all of a range of developmental environments
>
> (1998)

Both Sober and Stich's accounts come to the idea that an innate trait is one that we have a disposition to express. I agree, but I think this doesn't go far enough. The problem with dispositional accounts of innateness is that they fail to mark the relationship between developmental resources and the innate traits to which they contribute. Thus, we buy a philosophical description at the expense of forgoing a scientific explanation of the nature of innateness and innate traits. Some will baulk at this and claim that I also fail to provide a scientific explanation of innateness. Surely, innateness would only be fully explained if I could point to a particular scientific process by which traits become locked into populations. Pace Wimsatt, I think we ought not to fix on some particular mechanism and claim that it is the only metric by which a characteristic, fixed in a population can be said to be innate.

I think it would be an unnecessarily onerous requirement to insist that science could only describe universal grammar as innate if it could be shown to be the result of, say, generative entrenchment. What if it turned out not to be a product of that process? There is still much to be learned about Specific Language Impairment, but at least some sufferers appear to be completely normal except for a severe deficit in their ability to apply innate rules of grammar (van der Lely et al 1998). This implies that universal grammar is not generatively entrenched. If that is true, should we therefore coin a new term 'innate*'? It seems to me that we are much better off to think of innateness as multiply realisable.

IV Conclusion

The notion of innateness has fallen into disrepute. This is due in large measure to the tendency of both philosophers and scientists to see innateness as a claim about the relative causal contributions of genes and environment in the development of certain traits. As these causes are not independent, the claim is metaphysically suspect. Given that we now accept that phylogenetic inheritance need not be genetic inheritance, the claim is biologically suspect. In this paper, I have argued that we can respect the metaphysics by making innateness an informational rather than a causal property. I have further argued that we can respect the biology if we reject the requirement that innate traits be products of genetic information. Instead we need only see them as products of information from some set of reliably recurring developmental resources. It remains to be seen whether philosophers and scientists will accept innateness as a matter of degree.

REFERENCES

A. Ariew 1999 "Innateness is Canalization: In Defense of a Developmental Account of Innateness", In, *Where Biology Meets Psychology*, pp. 117-138. Cambridge Mass.: MIT Press.

P. P. G. Bateson 1976 "Specificity and the Origins of Behaviour", *Advances in the Study of Behaviour*, 6, 1 - 20.

P. P. G. Bateson and J. Vauclair 1975 "Prior Exposure to Light and Pecking Accuracy in Chicks." *Behaviour*, 52, 196 - 201.

N. Chomsky 1965 *Aspects of the Theory of Syntax,* Cambridge Mass.: MIT Press.

M. F. Claridge, J. A. Darwah and M. R. Wilson 1997 "Practical approaches to species concepts for living organisms", In M. F. Claridge, J. A. Darwah and M. R. Wilson (eds.), *Species: The Units of Biodiversity*, pp. 1 - 16. London: Chapman and Hall.

G. Gottieb 1971 *The Development of Species Identification in Birds,* Chicago: University of Chicago Press.

J. Gould and P. Marler 1991 "Learning by Instinct", In D. Mock (ed.), *Behaviour and Evolution of Birds*, pp. 4-19. San Francisco: Freeman & Co.

P. Griffiths and R. Gray 1994 "Developmental Systems and Evolutionary Explanation", *Journal of Philosphy*, XCI, 277 - 304.

D. Hull 1997 "The ideal species concept — and why we can't get it", In M. F. Claridge, J. A. Darwah and M. R. Wilson (eds.), *Species: The Units of Biodiversity*, pp. 357 - 380. London: Chapman and Hall.

E. F. Kinder 1927 "A study of the nest building activity of the albino rat", *Journal of Experimental Zoology*, 47, 117-161.

D. S. Lehrman 1953 "A Critique of Konrad Lorenz's Theory of Instinctive Behaviour", *Quarterly Review of Biology*, 28, 337 - 363.

R. C. Lewontin 1974 "The Analysis of Variance and the Analysis of Causes", *American Journal of Human Genetics*, 26, 400 - 411.

R. C. Lewontin 1985 "The Organism as Subject and Object of Evolution", In R. C. Lewontin and R. Levins (eds.), *The Dialectical Biologist*. Cambridge: Harvard University Press.

R. C. Lewontin 1991 *The Doctrine of DNA: Biology as Ideology,* New York: Harper.

K. Lorenz 1950 "The Comparative Method in Studying Innate Behaviour Patterns", *Symposia of the Society for Experimental Biology*, 261 - 268.

K. Lorenz 1966 *Evolution and the Modification of Behaviour,* London: Methuen.

D. Marr 1982 *Vision,* New York: W. H. Freeman.

D. H. Mellor 2000 "The semantics and ontology of dispositions", *Mind*, 109, 757-780.

S. Oyama 1985 *The Ontogeny of Information,* New York: Cambridge University Press.

C. Shannon and W. Weaver 1949 *The Mathematical Theory of Communication,* Urbana: University of Illinois Press.

E. Sober 1994 "Apportioning Causal Responsbility", In E. Sober (ed.), *From a Biological Point of View*. New York: Cambridge University Press.

E. Sober 1998 "Innate Knowledge", In, *Routledge Encyclopedia of Philosophy*, pp. 794-797Routledge.

K. Sterelny and P. E. Griffiths 1999 *Sex and death : an introduction to philosophy of biology,* Chicago, Ill.: University of Chicago Press.

K. Sterelny, K. Smith and M. Dickison 1996 "The Extended Replicator", *Biology and Philosophy*, 11, 377 - 403.

S. P. Stich 1975 *Innate ideas,* Berkeley: University of California Press.

H. K. J. van der Lely, S. Rosen and A. McClelland 1998 "Evidence for a grammar-specific deficit in children", *Current Biology*, 8, 1253-1258.

W. C. Wimsatt 1999 "Generativity, Entrenchment, Evolution, and Innateness: Philosophy, Evolutionary Biology, and Conceptual Foundations of Science", In, *Where Biology Meets Psychology*, pp. 139-179. Cambridge Mass.: MIT Press.