Innateness as genetic adaptation

Lorenz redivivus (and revised)

Forthcoming in *Biology & Philosophy*

Nathan Cofnas

Darwin College, Cambridge

cofnas@cantab.net

Abstract

In 1965, Konrad Lorenz grounded the innate–acquired distinction in what he believed were the only two possible sources of information that can underlie adaptedness: phylogenetic and individual experience. Phylogenetic experience accumulates in the genome by the process of natural selection. Individual experience is acquired ontogenetically through interacting with the environment during the organism's lifetime. According to Lorenz, the adaptive information underlying innate traits is stored in the genome. Lorenz erred in arguing that genetic adaptation is the only means of accumulating information in phylogenetic (i.e., intergenerational) experience. Cultural adaptation also occurs over a phylogenetic time scale, and cultural tradition is a third source from which adaptive information can be extracted. This paper argues that genetic adaptation can be distinguished from individual *and cultural* adaptation in a species like *Homo sapiens*, in which even adaptations with a genetic component require cultural inputs and scaffolding to develop and be expressed. Examination of the way in which innateness is used in science suggests that scientists use the term, as Lorenz suggested, to designate genetic adaptations. The search for innate traits plays an essential role in generating hypotheses in ethology and psychology. In addition, designating a trait as innate establishes important facts that apply at the information-processing level of description.

Keywords: innateness, innate–acquired distinction, genetic information, adaptation, Konrad Lorenz

Introduction

An influential line of thinking in the philosophy of biology holds that "innateness" is a deeply problematic concept. Some philosophers contend that it has no legitimate role in science (e.g., Griffiths and Gray 1994; Griffiths 2002; Oyama 2000; see also Godfrey-Smith 2007; Lewens 2015: 171; Shea 2012a, 2012b). Others, that it can play a legitimate role when and if, through philosophical analysis, it can be clearly defined (see Mameli 2008). In contrast, scientists—except for a few philosophically minded ones (notably Bateson 1991; Tomasello 1999)—tend to regard innateness as a valid, useful concept with an obvious one-sentence definition, e.g., "it just means structure in advance of experience," as Haidt (2007) says. Even the philosophically informed Richerson and Boyd nonchalantly use the term a total of 28 times in *Not by Genes Alone* (2005). Statements like, "[m]uch of morality's innate" (Haidt 2007), "[i]nnate language-specific expectations may be held about words" (Bloom 2000: 77–78), "[i]nitial knowledge is innate" (Spelke 1994: 438), are commonplace. The term "innate" (or its derivatives) is used in 26 out of 29 issues of *Behavioral and Brain Sciences* published from 2010 to 2014, a total of 322 times.[1] Not only do scientists routinely use the word, seemingly substantive debates in biology and psychology concern whether certain phenotypes are innate or acquired. What are these scientists arguing about, if innateness is as confused a concept as some philosophers believe?

---

[1] In 2015 *BBS* stopped dividing volumes into issues. The count here does not include 25 instances in the bibliography. Besides two cases where Joshua Knobe refers to intuitions about innateness, all 322 references to innateness are made uncritically in the context of debating whether some traits are innate or not.

Consider the paradigm example of an innate behavior. Upon emerging from the egg, newborn chicks consume grain by pecking at it in a stereotyped manner. They perform this behavior before it is ever reinforced or observed in a conspecific. It seems that there is something special about this sort of behavior that distinguishes it from so-called “learned” behavior. But being “special” is a vague notion. More precise candidate definitions of innateness all seem to be subject to counterexamples or other objections (Mameli and Bateson 2006). Naïve definitions of innateness like “due to genes rather than environment” or “genetically coded” are untenable for obvious reasons. More subtle definitions face subtler but no less serious problems.

Criteria of a successful analysis of innateness

There are three ways a philosophical analysis of innateness could be successful. First, it could show that the concept is simply confused, that it picks out no meaningful biological property, and, despite the apparent role it plays in productive scientific debates, its influence in science has been nothing but pernicious. (This is essentially what developmental systems theorists take themselves to have demonstrated.) Second, it could show that when innateness is employed by scientists, it (at least often) does have a meaning, *and* this meaning accounts for the productive role that it plays in biology and psychology. The account need not vindicate every appeal to innateness in the history of science, but it should show that the concept plays a productive role in at least some substantive debates. Third, it could show that innateness refers to a variety of disparate concepts, some of which may be scientifically valuable in themselves, but which have a weak (or no) association with each other.

Regarding this third possibility, Mameli and Bateson (2006) list a number of properties associated with innateness—“i-properties”—including “reliably appearing in a particular stage of the life cycle,” being species typical, a Darwinian adaptation, developmentally environmentally canalized, and post-developmentally environmentally canalized. On Mameli and Bateson’s view, as well as Shea’s (2012a, 2012b), while each of these properties may have useful roles to play in biology and psychology, there is no innateness concept that conceptually or empirically makes a tight association among them. Since the term “innate” falsely implies that strong inferences among i-properties can be made, it should be abandoned in favor of specifying the i-properties relevant to a specific investigation.

Strong prima facie evidence that innateness is indeed meaningful is that the innate–acquired distinction underlies several major research programs, especially in psychology, which are widely regarded as being tremendously successful. Moral foundations theory, associated with Haidt, has probably become the standard psychological account of morality. Haidt’s thesis is that “morality [is] innate (as a set of evolved intuitions) and learned (as children learn to apply those intuitions within a particular culture)” (Haidt 2012: 31). It is difficult to believe that the explanatory power of moral foundations theory is an illusion, and that the whole thing is based on a conceptual mistake. Similarly, the celebrated findings concerning children’s “innate knowledge” of physics, associated especially with Spelke, seem very meaningful, but the whole research program is predicated on the innate–acquired distinction (Spelke 1994, 1998, 1999). Again, this is only prima facie evidence that innateness is meaningful—it could be that Haidt and Spelke’s theories would be even more successful if they were not founded on the innate–acquired distinction. But we should only adopt this conclusion if all attempts at making sense of innateness fail.

This paper will defend, with a modification, the account of innateness of a working scientist—Konrad Lorenz—though one who was made sensitive to the basic philosophical problems involved when his original, naïve account of innateness was exposed as conceptually confused by Daniel Lehrman. It will be argued that, where the innateness concept has been usefully employed in science, it has generally been used to refer to something like Lorenzian innateness.

Overview of Konrad Lorenz's later account of innateness

One candidate definition of innateness—proposed by Lorenz (1965)—has received short shrift in the literature. In debates about innateness (e.g., Ariew 1996), Lorenz is often remembered as the author of some 1930s-era papers that were demolished by Lehrman (1953). Lorenz (1937/1970) equated *innateness* with being "hereditarily fixed and predetermined" (268), as opposed to being "acquired by the individual in the course of its ontogeny" (269). His litmus test for whether a phenotype is innate or acquired was the "deprivation experiment": If a phenotype develops in the absence of environmental cues, it is innate; otherwise, it is acquired. Lehrman (1953) pointed out that no trait is "hereditarily fixed," and none develops independently of "experience and environment." Even a behavior that occurs from the moment of an animal's hatching/birth depends on experience in the egg or in utero. For example, according to Kuo (1932), pecking behavior in chicks requires very specific experience—*in the egg*—to appear. In the egg, the head of the developing chick is forced up and down by the beating of its heart. This movement forces fluid into the chick's throat, causing swallowing. This experience causes the chick to develop a stereotyped behavior pattern of moving its head and swallowing—i.e., pecking behavior. So, Lehrman argued, the deprivation experiment "by its very nature does not give a positive indication that behavior is ‘innate’ or indeed any information at all about what the process of development of the behavior really consisted of” (343).

In 1965, Lorenz published a response to Lehrman. He acknowledged Lehrman’s essential points—that all phenotypes are produced by an interaction of genes and environment, and, therefore, innate structures cannot be distinguished by their independence from the environment. But Lorenz argued that there is clearly something noteworthy about (as he put it, perhaps wryly) “what we formerly called innate” structures. He suggested that adaptive phenotypes reflect information about *environmental givens*. There are two sources from which organisms can acquire information about environmental givens: *individual* and *phylogenetic experience*. Innate phenotypes, he suggested, are those that reflect information about environmental givens acquired in the course of phylogenetic, rather than individual, experience. Phylogenetically acquired information is stored in the genome. The special relationship between genes and innate phenotypes is not that the former cause the latter, but that the adaptive information underlying innate phenotypes is stored in the genome.

Lorenz (1965) introduces his revised account of innateness by suggesting that there is something special about adaptive behavior that an organism performs in the absence of experience (in its lifetime—pre- or post-birth/hatching) that could potentially reveal the behavior’s adaptive properties. If a chick emerges from its egg and, without having any experience with food (pre- or post-hatching), without seeing another organism eating, walks over to an object in its environment that happens to be food (i.e., grain) and eats it with perfect form, this is a different kind of behavior than, for example, learning not to touch an electric fence.

What is the significance of the deprivation experiment? Several recent commentators have, citing Lehrman (1953), asserted that the “deprivation experiment” can reveal, at most, that a phenotype develops in the absence of selected features of the environment (e.g., Griffiths and Machery 2008: 404; Moore 2001: 266, n. 122). They observe that all phenotypes require *some* sort of environment in order to develop. Establishing that a phenotype develops in the absence of one particular environmental variable does not place it in a special category because, again, all phenotypes require some environment to develop. Since (these commentators argue) “innateness” requires arising “independently of the animal’s experience and environment” (Lehrman 1953: 341), the deprivation experiment cannot establish that a phenotype is “innate” because it is impossible to experimentally remove all experience and environment.

According to Lorenz (1965), the deprivation experiment can do more than merely reveal that a chosen environmental variable does not contribute to the development of a trait. *If* the environmental variables that are withheld from the organism are those that could potentially provide it with information about the adaptive value of engaging in a particular behavior, and the organism performs the behavior anyway, the organism seems to be revealing knowledge about its environment in the absence of experience that would be necessary to acquire this knowledge (in its lifetime). These sorts of behaviors, according to Lorenz, reflect information about environmental givens acquired in the course of *phylogenetic* experience: The genes correlated with these behaviors proliferated due to natural selection, so the information is “contained in the genome.”

Of course, a deprivation experiment can only potentially reveal that a behavior is based on phylogenetically acquired information if the environmental variables that are withheld are those that are relevant to providing information about the consequences of the behavior under consideration. If the environmental variables that are withheld are not relevant in this way, the experiment may reveal interesting facts about the ontogeny of the behavior, but it will not establish whether or not it reflects information gained in the course of phylogenetic experience. If one wishes to demonstrate that grain pecking in chicks is informed by phylogenetic experience, it will not help to conduct a deprivation experiment in which the chick's mother is removed from its environment before it hatches. Even if the mother's absence were, for some reason, to disturb the development of grain-pecking behavior, the mother's presence per se does not provide information relevant to recognizing that grain is a food or how to eat it.

Revising Lorenz

According to Lorenz, genetic change is the only possible way of adapting to the environment over phylogenetic (i.e., intergenerational)[2] time. This is a mistake. *Cultural evolution* is another way of adapting over this timescale (Bateson 1991: 20; Shea 2012b). Knowledge and traditions that increase the adaptive fit between organism and environment can be accumulated across generations and passed on culturally. Thus there are (at least) two sorts of phylogenetic adaptation—genetic and cultural. On Lorenz's view, a trait is innate if the information underlying it is inborn (*innatus*)—stored in the genome. Contra Lorenz, innateness cannot be identified with phylogenetic adaptation because cultural adaptation occurs over phylogenetic

---

[2] Haeckel coined the term "phylogeny" (Greek: "origin of the race") to refer to the process by which species diverge. Lorenz (1965), Eibl-Eibesfeldt (1989), Shea (2013), and some others have used the term to mean "intergenerational." Lorenz used "phylogenetic adaptation" to mean "genetic adaptation" on the assumption that the only sort of adaptation that can occur over intergenerational time is genetic.

time and cultural information is not stored in the genome.

To rescue, in a revised form, Lorenz’s account of innateness as genetic adaptation, it must be possible to determine whether the information underlying an adaptive phenotype is stored in the genome, is acquired in individual experience, *or is culturally transmitted*, and there must be some utility in marking genetic adaptations as having a special status. The following section seeks to explain what it means for organisms to extract information about the environment from their genome, individual experience, or culture. The next section argues that it is possible to determine the source of information underlying adaptedness in practice, even in a species (namely, ours) where adaptations with a genetic component often require cultural inputs and scaffolding to develop and be expressed. The penultimate section argues that epigenetic transmission does not undermine this account. The last argues that scientists generally use innateness to refer to genetic adaptations, and that this accounts for the value of the innateness concept in productive scientific debates.

Three sources of information about the environment

According to Lorenz (1977), both morphology and behavior, insofar as they are adapted to the environment, are an “image of that environment.” The exquisitely adapted fins and movement of a fish, for example, “reflect the hydrodynamic properties of water” (6). “Information [is] the root of all processes of adaptation,” he said, “using the word in its everyday sense to denote something that has a meaning and a purpose for whoever receives or possesses it” (23). The information underlying “innate” traits, of course, is supposed to be stored in the genome (Lorenz 1965).

Genetic adaptation

What exactly could it mean for traits to “reflect” the environment, and for the underlying information to be stored in the genome? Lorenz (1977: 22–23) explicitly states that he is referring to information in the semantic, rather than the information-theoretical (i.e., correlational), sense. This semantic informational content is determined by how it is used by an organism whose purpose is “biological survival.” Using the framework of teleosemantics, Shea (2007) spells out how the genome can carry information by *representing* the environment.

Shea (2007) lays the groundwork for understanding genetic representation by introducing the idea of an “inheritance system.” An inheritance system is a developmental resource with the “meta-function” of producing heritable phenotypes. The most obvious (apparent) example is DNA. Certain DNA sequences—genes—have the teleofunction to give rise to the phenotypes whose selection caused their proliferation. *DNA itself*, along with its “associated developmental machinery,” exists in the form that it does because it is an effective mechanism for transmitting adaptive phenotypes (see also Lorenz 1965: 8). The particular genes under selection change over time, but the properties of DNA were fixed by selection because, by passing on particular genes from parents to offspring, they resulted in adaptive phenotypes reappearing from generation to generation. The function of being an inheritance system is “meta” because the system acquires the function indirectly as a result of direct selection on specific adaptive phenotypes.

Turning to the issue of genetic representation: On the standard naturalistic account, *representation* occurs when three elements of a system stand in a special relation to each other. A producer mechanism tracks variable environmental conditions by producing certain states. A consumer mechanism uses those states to respond appropriately to the variable conditions. The representational content of the produced state, or *signal*, corresponds to what must be the case for the consumer mechanism to perform its proper function(s) in the normal way (Millikan 1989b). The idea that genes serve a signaling role in this sense might seem like a nonstarter. While genes can trigger adaptive causal sequences, and be selected to have this effect, it is not clear how they can serve as signals concerning which of a variety of possible environmental states obtains (Papineau 2003: 121). A signal must have the potential to represent different states of the environment. The genes inherited by an organism remain the same no matter what happens.

Shea (2007), however, notices that if we extend the timescale over which we think about the relation between genes and environment, they do seem to serve a signaling role. As an inheritance system, DNA has the (meta-)function of preserving genes correlated with adaptive phenotypes—i.e., phenotypes that promote fitness in the selective environment. Therefore, any gene that is transmitted in the genome has the function of triggering the development of a phenotype that matches the environment. Consider the evolution of polar bear fur from brown to white. For simplicity, suppose that in the recent evolutionary history of bears there were two genetic variants correlated with differences in color: Gene-B was associated with brown fur, gene-W with white fur. The first bears to arrive in the Arctic all had B. When W arose as a mutation, its possessor had greater success in hunting because it was camouflaged against the snow. As a result, it had more offspring than the average brown bear, and its W-possessing descendants had the same advantage. Thus W proliferated and went to fixation. DNA fulfills its meta-function by preserving W, a gene that, relative to the alternative (B), is associated with a phenotype that promotes adaptedness in the environment.

Genes are what Millikan calls “pushmi-pullyu representations” (Shea 2007: 324, n. 9): Their representational content is simultaneously directive and descriptive/indicative. As Millikan (1995: 189) explains, a *directive* representation “has a proper function to guide the mechanisms that use it so that they produce its satisfaction condition.” The truth condition of a descriptive representation is the “condition to which it adapts its interpreters or users in the service of *their* proper functions.” Continuing with the example above, gene-W has the directive content to produce white fur, and the indicative content that the environment is white. The former can fail to be satisfied (if something happens in development that leads to nonwhite fur), while the latter can be false (if the environment becomes nonwhite). Qua descriptive representations, genes are signals.

There is a tricky question of who or what is supposed to be the *consumer* of genetic representations. Without identifying a consumer and its proper function(s), it is impossible to fix the indicative content of a signal, in this case a signal supposedly instantiated in genes. Godfrey-Smith (2007: 64) raises the issue pointedly: In order to make his account of representation work, “Shea is forced to argue for the reality of something that is described like a piece of biological machinery, but is in fact an abstraction.” It may be appealing to postulate the existence of consumers of genetic representations in order to vindicate our intuition that the genetic code represents something, but what is the “independent motivation” to think that these consumers actually exist? Shea (2013: 2–3) counters that the consumer in question “is not a neat single mechanism, but consists of all the messy, interactive and temporally extended processes of development.”

As noted above, Lorenz (1977: 22–23) contrasted the information-theoretical concept of information with a semantic concept that makes reference to the “purpose for whoever receives or possesses it”—the purpose for organisms being “biological survival” and, we could add, reproduction. The organism itself, having the purpose to survive and reproduce, may be the real “consumer” of semantic information in the genome. The “interactive and temporally extended processes of development” that Shea refers to do seem like an abstraction (as Godfrey-Smith alleges), since it is not clear how we would go about actually identifying specific “processes” that act as consumers. The organism itself is a concrete thing with a concrete proper function (to survive and reproduce) that is advanced by correctly using accurate genetic representations to produce biological devices that promote adaptive behavior in a given environment.

Exaptations, spandrels, and genetic representation

Gould and Vrba (1982) distinguish *exaptations*—structures designed by natural selection to serve a function and later coopted for a different “role”—from *adaptations*. To take perhaps their most famous example, feathers originally evolved for the purpose of insulation in dinosaurs, and came to be used for flight in (some of) the descendants of *Archaeopteryx*. Feathers are an exaptation, not an adaptation, for flight. On Gould and Vrba’s view, only adaptations have “functions,” whereas exaptations have only “effects.” That would seem to imply that the genes associated with feathers in (e.g.) macaws have representational content about the weather and the value of insulation during the Jurassic period, when some dinosaurs first evolved feathers, rather than about anything to do with (as Lorenz would say) the properties of the air as they relate to flying today.

About this particular example it should be noted that, since they first appeared, feathers have undergone a great deal of modification due to selection on genes. The feathers of penguins and bowerbirds, for example, were not merely exapted, but have been selected to serve new functions. The genes underlying differences in penguin and bowerbird feathers have indicative representational content about features of Antarctica and the ocean and Australasia and the air, respectively.

Even those genes that originally evolved in the Jurassic because they were associated with insulating feathers have been *preserved* by selection for different reasons in different bird species. They have been selected for different functions. As Millikan (1989a: 173) says, for natural selection to account for a trait, the selection may have occurred "'only yesterday.'" Indeed, she says, "the *main business* of natural selection" is stabilizing selection—protecting useful traits by weeding out harmful mutations. Taking this view, traits acquire new functions when selection pressures change (see also Godfrey-Smith 1994), and the underlying genes acquire new representational content at the same time.

But some adaptive traits come into existence without ever being selected for any reason. The sutures in the skulls of young birds and mammals are unfused. For birds—and presumably for the common ancestor of birds and mammals—the unfused sutures have (or had) no adaptive value. For mammals, unfused sutures make it possible to leave the birth canal uninjured (Gould and Vrba 1982). In this case, genes underlying unfused sutures were (probably) never selected in favor of alternatives that were associated with fused sutures, so we cannot say that those genes have indicative representational content about the birth canal. However, the genes that were selected for their role in producing the birth canal do have indicative content about *features of young mammalian skulls*, namely, that they can be compressed.

Individual experience

The foregoing gives an account of how the genome can contain information about the environment, namely, by representing it. Lorenz claimed that adaptive information can also be gained in “individual experience,” when the organism undergoes adaptive modification of its phenotype by extracting information from the environment in its lifetime. For this to occur, the organism needs some criterion for recognizing what counts as *adaptive* modification. For example, a goat that touches an electric fence gets shocked and does not touch the fence again. On Lorenz’s view, it has acquired information about its environment, reflected in the adaptive phenotype (avoiding the fence). In order to determine that the fence is harmful, the goat has to have a mechanism for identifying harmful objects in its environment. Certain types of nerve impulse—what we experience as “physical pain”—are correlated with tissue damage. The mechanism causes the goat to avoid whatever object contact with which immediately preceded the painful stimulus. The “pain” is used as a signal of “tissue damage.” That the *fence* causes tissue damage is information extracted from the environment in the goat’s individual experience.

Biological mechanisms with the proper function of bringing about a certain relation between the organism and an element of the environment have what Millikan (1984) calls “relational proper functions.” The elements of the environment toward which they have the function of bringing about a relationship are their “adaptors.” For a mechanism to successfully perform its relational proper function (in the normal way) it must correctly represent the environment—specifically, its adaptor (Cofnas 2016: 512–513). Mechanisms with relational proper functions can themselves be established by genetic adaptation, by cultural adaptation, or through individual experience. In the case of the goat avoiding the electric fence, the mechanism in its nervous system that causes it to represent and avoid harmful objects in the environment (namely, objects that cause tissue damage) is presumably a genetic adaptation. But representation of the adaptor *must* occur in the lifetime of the individual—this information must be gained in individual experience, and it is stored as a representation.

Cultural adaptation

Lewens (2015: 58) notes that sometimes cultural evolutionists use the term "cultural information" to refer to the "collections of intentional states" that are transmitted and adaptively modified from one generation to the next, sometimes to refer to the "potentially far broader range of resources that enable this" transmission to occur. He writes:

> In this mode, the search for cultural information is the search for those factors that explain the organism's dispositions to behave appropriately. Loosely, it is the search for factors that help to answer Shea's question, 'How does the organism know what to do?'

But not all culturally acquired intentional states (or the resources that enable their transmission) are adaptive in a biological sense—that is, they do not "explain the organism's dispositions to behave appropriately." Most cultural evolutionists define culture and cultural information in terms of the means by which it is transmitted and without reference to its adaptive value. Sperber (1996: 1) characterizes culture as "made up, first and foremost, of...contagious ideas. It is made up also of all the productions (writings, artworks, tools, etc.) the presence of which in the shared environment of a human group permits the propagation of ideas." Richerson and Boyd (2005: 5) say that by "culture" they mean "information capable of affecting individuals' behavior that they acquire from other members of their species through teaching, imitation, and other forms of social transmission." "By *information* [they] mean any kind of mental state, conscious or not, that is acquired or modified by social learning and affects behavior." Jablonka and Lamb (2005: 160) "see culture as *a system of socially transmitted patterns of behavior, preferences, and products of animal activities that characterize a group of social animals*."

On this characterization of culture, the copycat effect that occurs after a celebrity commits suicide (Henrich 2016: 49–50) is an example of "cultural transmission." The suicide generates ideas in other people that potentially affect their behavior. It is, in Richerson and Boyd's words, "information capable of affecting individuals' behavior that they acquire from other" people.

Copycat suicide may be cultural transmission, but it is not cultural *adaptation*. Cultural adaptation relies on a different concept of information than is usually employed in studies of cultural evolution. For the present purposes, it is crucial not to conflate these phenomena. The mechanisms underlying cultural transmission in general—whether they themselves are genetic or cultural adaptations (see Heyes 2012a)—presumably have the proper function to accomplish cultural *adaptation*, but they frequently fail to perform this function, especially under modern conditions.

On Shea's account, genes represent only because they serve a signaling function in an inheritance system (i.e., a mechanism with the meta-function to produce heritable phenotypes). "Otherwise the representational framework would not apply. So only genuine inheritance systems represent in the way genes do" (Shea 2007: 327). For him, whether cultural information has the same kind of representational capacity as genes hinges on whether humans possess an inheritance system with the meta-function to transmit heritable cultural variants. He sees evidence that our dispositions for receiving cultural variants are designed by selection to operate in a way analogous to DNA (Shea 2012b). Humans have a tendency to overimitate when learning. In contrast to chimpanzees, we tend to copy all elements of an action rather than just those elements that appear causally necessary to achieve the desired end. Shea argues that we have evolved dispositions that tend to make cultural transmission more vertical than horizontal or oblique—a necessary condition for linking the frequency of cultural variants with their associated fitness. Thus, he says, under the right conditions, behavioral dispositions in a population that were transmitted by overimitation "will carry information about past environments" just like genes (Shea 2012b: 2240).

It may be true that being an inheritance system is necessary to represent "*in the way genes do*," but there are other ways in which culturally transmitted information can represent the environment. The *adaptive* information in a cultural tradition can be generated in two fundamental ways. First, adaptive phenotypes can be based on information acquired in individual experience, and these phenotypes (at least behavioral ones) can potentially be transmitted by imitation, emulation, scaffolding the learning environment of the next generation, and so on. As a result of cumulative cultural transmission, individuals can, using insight, build on the discoveries of previous generations to produce behaviors and suites of behaviors that no one individual could ever have devised from scratch in their lifetime (e.g., how to control fire or do calculus). (Some nonhuman animals might have cultural traditions of this sort, although their capacity for cumulative cultural transmission appears to be severely limited.)

The second way in which adaptive phenotypes arise in a cultural tradition is by natural selection acting on cultural variants that vary more or less randomly (Eibl-Eibesfeldt 1989: 11–12, 16; Richerson and Boyd 2005: 4, 13–14, 68–79; Rogers and Ehrlich 2008; Shea 2012b:

2240). Insofar as children tend to blindly imitate their parents, cultural variants can serve a function analogous to genes, being pushmi-pullyu representations with the person as consumer. When a cultural tradition has accumulated a body of naturally selected information, selection can then favor learning dispositions that prompt people to acquire these variants horizontally or obliquely. The disposition to adopt the beliefs and practices of the majority, which may itself be a genetic adaptation (Richerson and Boyd 2005), is one way in which people can access the body of naturally selected cultural information in their environment without inheriting it vertically. If so, the disposition to conform is rooted in a genetic representation of the environment with the indicative content that the most common cultural variants in a group are probably associated with the greatest fitness.

Psychologists point out that people accomplish adaptive ends by adopting beliefs that, in a variety of contexts, are “systematically false” (Pinker 2005: 18). It is important to distinguish two senses of “truth” and “falsity.” A belief in a true proposition may be *false* qua signal of the environment used by a consumer mechanism with the proper function to bring about an adaptive relationship between organism and environment. Consider two different propositions, which are culturally transmitted as beliefs:

(1) God hates birth control and commands us to be fruitful and multiply.

(2) Using birth control to prevent pregnancy will elicit no punishment in either this or any other life.

Proposition (1) is presumably false while (2) is true. Nevertheless, belief in (1) confers a strong selective advantage under current conditions. Evidence shows that, since religious beliefs are often passed from parents to children (Richerson and Boyd 2005: 76–77), belief in (1) is proliferating (Norris and Inglehart 2011). Given natural selection for belief in (1), the belief reflects information about the environment in the same way as a genetic variant associated with an adaptive phenotype. The belief is a pushmi-pullyu representation with the directive content to have children and the indicative content that conditions will support having many children (i.e., resources are available to raise many children who will themselves be able to survive and reproduce).

Distinguishing genetic, individual, and cultural adaptation

The previous section argued that there is a meaningful theoretical distinction between the genome, individual experience, and culture as sources of information. In the real-life messy process of ontogeny, a single phenotype often reflects information acquired from different sources, or it may develop in stages, each stage guided by information from a different source(s). Because of gene–culture coevolution, the genes underlying some human genetic adaptations reflect information about the *cultural* environment. This section address the question of how to distinguish genetic, individual, and cultural adaptation in practice.

How do individuals extract naturally selected cultural information?

Naturally selected cultural variants reflect information about the environment just as genes do.

To use this information, you need to extract it somehow. Whether you receive it vertically, horizontally, or obliquely, you need procedures for identifying what to copy from which potential models. The heuristics that we use to guide the acquisition of cultural variants can themselves be acquired through genetic adaptation, individual experience, or cultural inheritance. Adaptive behavior may be informed by integrating genetically, individually, and culturally acquired information in complicated ways. Though difficult in practice, it is always conceptually possible to distinguish the source(s) of information underlying adaptive phenotypes.

Consider the human disposition to "overimitate" mentioned earlier. When copying goal-directed behaviors, human children from around the age of 5, in contrast to chimpanzees, often imitate all elements of the model's behavior, including those which are obviously causally irrelevant to obtaining the goal (Horner and Whiten 2005; Nielsen and Tomaselli 2010). Surprisingly, after witnessing a model extract an object from a box using a tool, adults are *more* prone to copy obviously causally irrelevant actions than 3- or 5-year-old children (McGuigan et al. 2011). Children are over-imitators. Adults are "super-copiers" (McGuigan et al. 2011), though their tendency to overimitate is sensitive to characteristics of the model (McGuigan 2012).

Some commentators propose that the tendency to overimitate could be an adaptation that allows us to acquire culturally accumulated information, the adaptive value of which is not evident to us (e.g., Henrich 2016: 108–109; Nielsen and Tomaselli 2010; Shea 2012b: 2239–2240). Shea (2012b: 2240) suggests that the tendency may have evolved genetically *or* culturally. Heyes (2012b) argues forcefully against the genetic-adaptation hypothesis, providing evidence that both the ability and disposition to overimitate arise as a result of intensive training from infancy. The resulting tendency for overimitation is rooted in associative learning, she says—it is not “innately human” (200). Whether the tendency to overimitate is a genetic adaptation or is culturally acquired is not important for the present discussion. What is important is that neither scenario raises problems for the distinction between genetically, individually, and culturally acquired information.

Consider first the possibility that overimitation is a genetic adaptation. Because the benefits of inherited behavior are often opaque to the individual, we are often better off fitness-wise relying on selection-tested traditions than on our own rational insight (Eibl-Eibesfeldt 1989: 11–12; Henrich 2016: chapter 7). Suppose that, among early hominins, some genetic variants were associated with a tendency to overimitate, while others were associated with a tendency to imitate only those actions that seem causally relevant to the goal. So, for example, over-imitators would tend to attach feathers to the back of the arrows they made, if this was the practice of the models from whom they learned arrow making. Non-over-imitators would make un-fletched arrows if the purpose of feathers was not evident to them. In the case of arrow making and a thousand other examples, the over-imitator would be more successful. If the genes underlying overimitation proliferated, the phenotype reflects information about the environment, namely, that all elements of (certain) models’ goal-directed behavior are important, even if their causal relevance is not apparent.

Assuming the story given in the preceding paragraph is correct, the genome stores information about an aspect of the environment (namely, culture). The phenotype—the tendency to overimitate—is innate. That fact that the phenotype reflects information about culture rather than some other aspect of the environment does not pose any special conceptual difficulty. The tendency to overimitate will not manifest itself unless a person is exposed to a cultural environment, but, as noted, innate traits even in non-cultural animals require the right environmental conditions to develop and be expressed. There is no obvious conceptual problem arising from the claim that we are genetically adapted to a cultural environment, thus that information about culture is stored in the genome, and the genetically adapted trait is innate.

Second, consider the possibility that overimitation is acquired culturally in the way Heyes (2012b) suggests. Adults frequently mimic infants and respond positively when infants mimic them. This establishes associations in developing children between their own movements and the corresponding movements in other people. Mimicry training constitutes a surprisingly large proportion of adult–infant contact. In face-to-face interactions between mothers and their infants, mother and infant mimic each other about once a minute. Usually (79% of the time) these episodes are initiated by the mother. Experience like this throughout infancy and early childhood accounts for our ability to imitate and our tendency to *over*imitate.

On Heyes's view, the disposition to overimitate is passed down vertically (parents to children) and obliquely (non-parental caregivers to children). Suppose this culturally inherited trait has been naturally selected: In our evolutionary history, those whose parents and caregivers mimicked them in infancy, and rewarded them for mimicking adults, grew up to be overimitators. As overimitators, when learning new skills they copied actions whose purposes they did not understand, such as attaching feathers to the back of arrows. Overimitators were more successful as individuals and groups because they could better extract adaptive information from their cultural traditions. Overimitators had more children, and their children acquired the same disposition because of the experience given to them by their caregivers.

Such a scenario does not pose any great conceptual challenge. A trait that helps us to tap into culturally preserved information (overimitation) is itself culturally acquired. The trait carries the same information about the environment whether that information is stored in the genome or is passed on culturally—namely, when copying certain models to achieve a goal, seemingly causally irrelevant actions actually serve a purpose.

Niche construction

The physical and biotic environment in which organisms develop, acquire skills, and act may be altered in a variety of ways by themselves and by other members (or generations) of their species. Often, organisms will not develop normally or act successfully unless the environment has been altered in certain ways. In such cases, what is the source of information underlying adaptive phenotypes?

Start with a simple case, where a species is genetically adapted to manipulate the environment in a specific way, and to develop and make its living in such a manipulated environment. Beavers are genetically adapted to build dams. They have an instinct (an innate behavior) to cover anything that sounds like running water with sticks and mud. This trait reflects information about their environment: Something emitting the sound of running water is a river and covering rivers with sticks and mud will create an environment that is beneficial (to beavers). Beavers have other innate traits that will only manifest in a dam environment. These other innate traits reflect information about the dam environment, which is itself produced as a consequence of other innate traits that reflect information about *how to build dams*. The information about how to build dams and also how to navigate them is stored in beaver genomes.

On Sterelny's (2012) account, among humans (and our hominin ancestors) complex skills are transferred from generation to generation by a process of apprenticeship. Hominin skills are too demanding to impart by means of local or stimulus enhancement or even mere teaching. The apprentice learning by which we acquire technical skills is "learning by doing in an environment seeded with informational resources," including materials "in both their raw and processed forms," examples of products the learner is trying to make in various stages of completion, as well as guidance by experts who not only give explicit instructions but also shape the learning environment in a way that conduces to mastering the relevant skills step by step (35–36). According to Sterelny, this process also "depends on individual cognitive adaptations for social learning" (34).

How can we disentangle the sources of information that we draw on during apprentice learning, the success of which depends on niche construction, individual and various kinds of social learning, and adaptations that are the product of gene–culture coevolution? The practice of arranging the learning environment of apprentices can be the result of individual insight passed from one generation of teachers to another, or it could be a naturally selected cultural practice that proliferated because it aided the transfer of adaptive skills from generation to generation, and it could even have a genetic component.

Some animals appear to have genetic adaptations for structuring the *learning* environment of their young. Several carnivores including domestic cats (*Felis catus*), cheetahs (*Acinonyx jubatus*), otters (*Lontra canadensis*), and sparrowhawks (*Accipiter nisus*) catch and release live prey in the presence of their young, allowing them to practice hunting (Byrne 1995: 140–141). Experimental evidence with domestic cats shows that this experience aids the development of hunting skills (Caro 1980). Human adults seem predisposed to speak to babies—and, as Byrne points out, pet animals and sometimes even house-plants—in "baby talk": "'What's that you've got there? A nice, red ball, is it? Isn't that pretty? Wouldn't it be fun to drop it?'" (Byrne 1995: 37). Baby talk contains little novel information—what information it does contain is usually obvious—but it employs a lot of vocabulary and complex syntax, and is spoken in an "exaggerated" way that makes the syntactic structure clearer and phonemes more differentiated.

In cases like the aforementioned, it is still possible to distinguish "innate" elements of the behavior that develop from those that are acquired. The fact that human language acquisition may be aided by baby talk does not undermine poverty of the stimulus arguments for an innate language faculty, if indeed the experience provided by such talk, along with the other data that we are exposed to, does not provide as much information as is manifest in the linguistic competence that we ultimately attain. In Chomsky's (2000: 2) theory, "universal grammar" refers to our "innate biological endowment" that provides information about language that goes beyond what can be extracted from the environment.[3] Language competence would not develop at all in the absence of certain linguistic experience: Some of the information necessary to acquire language is not innate. From adults' perspective, the tendency to speak to children in baby talk is (possibly) a genetic adaptation.

Many reliably developing human capacities are scaffolded by the interventions of teachers/models. The crucial interventions may be naturally selected cultural practices, genetic

---

[3] Chomsky has recently suggested that the key human-specific structures underlying language competence are exaptations rather than adaptations for language (Hauser et al. 2002). (On whether exaptations are "innate" in the sense defended in this paper see the subsection "Exaptations, spandrels, and genetic representation," above.) Leaving aside the empirical tenability of this claim, many, of not the majority, of linguists think that poverty of the stimulus arguments are a useful tool to help identify possible genetic adaptations for language.

adaptations, or the result of individual insight (about how to impart skills) that is passed from one generation of teachers to another. The capacities might not develop unless information from all these sources is integrated, and the environment is engineered in the right way. Still, it is conceptually possible to classify information as being drawn from one of the three sources—genetic, individual, and cultural.

Epigenetics

Shea does not think that DNA and cultural learning by overimitation are the only inheritance systems. He suggests that there are also epigenetic inheritance systems. This subsection briefly argues that epigenetic inheritance is not a fourth source of information, and does not undermine the innate–acquired distinction.

Citing Jablonka and Lamb (2005), Shea (2012b: 2238) claims that “[s]ome epigenetic effects may be transmitted with high fidelity down many generations, in which case they can accumulate information by selection on epigenetic variants.” While this is theoretically possible, there is no clear evidence that any actual adaptation evolved by the accumulation of epimutations (Haig 2007). Haig (2007: 421–422) suggests that the “low-fidelity of epigenetic inheritance” largely precludes such adaptation. Adaptive epigenetic variants would tend to be lost too quickly for complex adaptations to be based upon them. As an empirical proposition it is highly questionable that there are epigenetic inheritance systems that have the meta-function of accumulating information about the environment in a way that is analogous to DNA.

Even if epigenetic inheritance *were* sufficiently high fidelity to accumulate information about the environment in a way analogous to DNA sequences, this would necessitate only a minor revision to the account of innateness defended in this paper. We could simply say that innate traits are phylogenetic adaptations where the information is stored in the genome (i.e., DNA sequences) *or epigenetic states*. Of course we would care whether the information was stored in one or the other, but the innateness concept would not be radically affected.

An epigenetic phenomenon that cannot be so straightforwardly accommodated is "intergenerational switches." This is when environmental conditions trigger epigenetic changes, which can be inherited by the subsequent generation, and which lead the descendants to develop adaptations for the triggering conditions. Is the concept of Lorenzian innateness undermined by this phenomenon? Haig (2007: 421) observes that the ability to switch epigenetic states is itself a property of the genome, and so can evolve by conventional natural selection acting on genetic mutations. Epigenetic switches with adaptive effects—e.g., that cause an animal to grow a thicker coat if its parents experienced a harsh winter—will be naturally selected, while those with maladaptive effects will tend to be eliminated. The switches themselves are genetic adaptations for responding to environmental conditions encountered in the individual experience of previous generations.

Haig thinks of heritable epigenetic switches as "an expression of transgenerational *ontogeny*." We typically think of ontogeny as beginning with fertilization, but in the case of epigenetic switches we could think of it as beginning sooner—before meiosis. The ontogeny of methylation patterns begins in the previous generation or two. If the timeframe of ontogeny is pushed back in this way, epigenetic switches are (as noted) genetic adaptations, while the environmental trigger is *experienced* (pre-meiosis).

How innateness is used in actual scientific discourse—and to what end

After Tomasello (1999) rejects the concept of innateness as based on pre-Darwinian thinking, he acknowledges that "the search for innate aspects of human cognition has...led to some very important insights" (50). The example he gives is the discovery that, contra Piaget, it is not necessary for infants to manually manipulate objects in order to acquire understanding of objects in space, since this understanding develops before such experience (citing work by Baillargeon and Spelke). "But," he says,

> this discovery should not stop the process of inquiry—we should not simply say that X is innate and so our job is done—but rather it should lead us to ask other questions, for example, the role of visual experience by itself in the absence of manual manipulations in the development of a concept of object. (50–51)

This echoes Lehrman's (1953: 344) concern that designating traits as innate "adds nothing to an understanding of the developmental process involved," and "discourages examination of the embryological processes leading" to the trait. If Lehrman's statement is meant to be an empirical proposition—that seeking to identify "innate" traits will discourage research on ontogeny—it seems to have been falsified. Many scientists seek to classify traits as innate or acquired, and the study of development is a flourishing subfield of biology. In the history of biology and psychology, there may indeed have been investigators who saw the discovery of "innate" traits as the end of research. The majority do not seem to have fallen into this trap.

But innateness will not be vindicated as a valuable concept if all that can be said for it is that it does not cause too much confusion. What is *gained* by identifying a trait as innate—as a genetic adaptation—besides the fact that a particular environmental variable is not necessary for its ontogeny? In Shea's (2012b: 2236) view, innateness claims are made "on the basis of, and used to make inferences to, the various i-properties" (with "genetic coding" being just one of the i-properties rather than the essence of innateness, as argued here). A scientifically useful innateness concept would, therefore, explain why reliable inferences of this sort can be made. The problem, he says, is that the i-properties are mostly uncorrelated among human cognitive traits, so innateness claims cannot be reliably made in regard to "traits that have arisen in recent hominin evolution." This final section argues that innateness claims are for the most part made to designate genetic adaptations, not to make inferences among i-properties. Their primary usefulness is in hypothesis generation and describing traits at the information-processing level of description.

Genetic disease: not a counterexample

Conceiving innateness as genetically based adaptation means that genetic diseases are not innate. Is this a problem? Many philosophers take it for granted that, assuming innateness means anything, genetic diseases are innate (e.g., Mameli 2008; Mameli and Bateson 2006; Stich 1975). But this does not seem to correspond with scientific practice. A (2016) search on the website of one of the leading journals devoted to genetics and disease—the *Journal of Medical Genetics*—finds 60 articles that have used the word "innate" since that publication was founded in 1964. If we discount 13 articles where "innate" appears only in the bibliography, one where it is used in a description of historical understandings of evolution, and two obituaries, that leaves 44. In 33 of these articles the term is used in connection with the immune system. “Innate” and “acquired immunity” are technical terms that refer to pathogen-specific receptors that develop before, or in response to, exposure to specific pathogens. In this usage innateness means genetic adaptation. In almost half of the remaining 11 papers, “innate” is used to refer to genetically based differences among people. In five cases it is used to describe some sort of failure, the cause of which is internal as opposed to external to the body, e.g., “the cause of the subfertility was obstructive rather than being due to an innate failure of germ-cell production.” In only one of these five papers is an innate failure tied to genes: “the condition is...more likely to be the result of innate errors of development, probably of genetic origin.” This is the only passage in the history of the *Journal of Medical Genetics* that comes close to calling a genetic disease innate. It could also be interpreted as saying that an innate process of development can be disturbed by genetic mutations, i.e., normal development is innate but it can be disrupted if the underlying genes mutate.

These data suggest that by far the most common usage of innateness in medical genetics is to refer to genetic adaptations, usually the genetically adapted aspects of the immune system. Very rarely it refers to genetically based differences among people. Genetic diseases may or may not be innate according to the “folk” understanding of that concept. Despite the fact that genetic diseases have been presented as paradigm cases of innateness in some philosophical discussions, there is very little, if any, support for this from scientific discourse.

The usefulness of the innateness concept

Lorenz (1965) reports that his own discoveries in ethology were first formulated as hypotheses about innateness. As discussed above, Tomasello himself notes that the search for “innate” cognitive traits has led to important discoveries. In fact, the innateness concept plays a crucial role in suggesting hypotheses to test, and this accounts for an important part of its usefulness in productive scientific work.

In an exchange in *Developmental Science*, Smith (1999) argues that nativism is a confused research program, while Baillargeon (1999) argues that it is a potentially fruitful one. Spelke (1999: 145) points out that Smith and Baillargeon reach opposite conclusions after presenting essentially the same arguments: Both cite a series of experiments on the development of infants’ knowledge about objects, then report new experiments in which they observe and manipulate infants’ developing knowledge. Although their research programs are “closely parallel,” Baillargeon sees her research as testing nativist hypotheses, while Smith sees her research as contributing to a “mechanistic” explanation of development to which the innate–acquired or innate–learned distinction contributes nothing.

The case of Smith vs. Baillargeon illustrates that research on psychology can be carried out in more or less the same way regardless of whether or not the researcher uses the word “innate.” The important question is whether psychologists like Smith and Tomasello who reject the word have also jettisoned the *concept* of genetic adaptation. There is reason to think not (cf. Lorenz 1965: 4).

Consider some research described in the course of Smith’s (1999) anti-nativist argument. Around the age of 2, children begin to manifest the shape bias in object naming. That is, when they learn that a name applies to an object, they assume the same name applies to other objects of similar shape (as opposed to size, color, material, or some other property). Smith gives compelling evidence that this bias arises from the fact that the first nouns that children tend to learn categorize objects by shape. In her laboratory, she caused children to develop the bias earlier than usual through training. She concludes from this that the “belief” about the significance of shape in categorizing objects

> is neither a primitive nor one that is made out of more primitive beliefs. Instead, the knowledge that objects are named by shape is made out of something different, the association of one perceptible cue with another. Thus, it would appear that children’s smart learning of object names begins with neither the specialness of shape nor the specialness of nouns but with more ordinary psychology. (138)

Although Smith rejects the word innate, it seems she has given an answer to Lorenz’s question of what is the source of information underlying adaptedness—the genome or individual experience (or culture)? Her conclusion is that information about the role of shape in object classification is acquired in individual experience, relying on a mechanism for associative learning that is a genetic adaptation. She does not imply that we should refrain from investigating the development of the capacity for associative learning, but she implicitly framed her experiment as an investigation of the source of information underlying adaptation.

Like Lorenz, Spelke (1998: 192) sees investigations about innateness in human cognition as concerning “the sources of human knowledge.” Specifically addressing issues raised by Oyama, she says that “the nativist-empiricist dialogue is not about the interaction of genes and their environment, but about whether knowledge of things in the external world develops on the basis of encounters with those things.” The fact that an organism interacts adaptively with its environment means that it has somehow acquired information about the environment. Educated guesses can be made about what information is likely to be contained in the genome, what extracted from the environment (in individual or cultural experience), and those guesses provide the basis for a great deal of experimentation in psychology.

Take another illustrative example. Six- to 18-month-old infants tend to grab at objects within their reach, exploring and often putting them in their mouths. However, they will generally not grab at *plants*, and, if they do, they will not put them in their mouths unless they witnessed an adult eating the same kind of plant. The experiments that uncovered this fact were designed by Wertz and Wynn (2014) on the assumption that young people have a genetic adaptation to extract culturally acquired information about local flora—to avoid plants that are dangerous, and consume those that are nutritious.

Even if thinking in terms of genetic adaptation (i.e., innateness), whether implicitly or explicitly, is crucial for designing studies like those of Wertz and Wynn, Smith, Baillargeon, and Spelke, is there any value in labelling a trait "innate" *after* it has already been discovered? If not, searching for genetic adaptations might be a useful research strategy, but once innate traits are discovered there would be no more utility to labelling them as such.

The value in labeling a trait as innate is to convey scientifically important facts about it that apply at the information-processing level of description. The study of moral reasoning, language acquisition, and object recognition is not conducted at the level of physiological causes (cf. Spelke et al. 1992: 629). This is partly due to the incompleteness of our current science, and, perhaps more important, to the fact that understanding these phenomena at a higher level of description is an end in itself (even though identifying traits as "innate" is just the first step toward investigating their physical basis). For example, the principles that guide language acquisition and use are not instantiated in the brain as “principles.” Even if we knew the whole physical process by which auditory stimuli impinge on developing brains, and how the capacity for language generation is stored and verbal behavior performed, the source of information being employed at this level would be a separate question, to be answered by reference to innateness.

Acknowledgements

Thanks to Patrick Bateson, Cecelia Heyes, Tim Lewens, David Papineau, Neven Sesardić, Kim Sterelny, and two reviewers for *Biology & Philosophy* for helpful comments on earlier drafts of this paper.

References

Ariew A (1996) Innateness and canalization. Philos Sci 63:S19–S27

Baillargeon R (1999) Young infants’ expectations about hidden objects: a reply to three challenges. Dev Sci 2:115–163

Bateson P (1991) Are there principles of behavioural development? In: Bateson P (ed) The development and integration of behaviour: essays in honour of Robert Hinde. Cambridge University Press, Cambridge, pp 19–39

Bloom P (2000) How children learn the meanings of words. MIT Press, Cambridge

Byrne RW (1995) The thinking ape: the evolutionary origins of intelligence. Oxford University Press, Oxford

Caro TM (1980) Predatory behaviour in domestic cat mothers. Behaviour 74:128–147

Chomsky N (2000) New horizons in the study of language and mind. Cambridge University Press, Cambridge

Cofnas N (2016) A teleofunctional account of evolutionary mismatch. Biol Philos 31:507–525

Eibl-Eibesfeldt I (1989) Human ethology. Aldine de Gruyter, New York

Godfrey-Smith P (1994) A modern history theory of functions. Noûs 28:344–362

Godfrey-Smith P (2007) Innateness and genetic information. In: Carruthers P, Laurence S, Stich SP (eds) The innate mind, vol. 3: foundations and the future. Oxford University Press, Oxford, pp 55–68

Gould SJ, Vrba ES (1982) Exaptation—a missing term in the science of form. Paleobiology 8:4–15

Griffiths PE (2002) What is innateness? Monist 85:70–85

Griffiths PE, Gray RD (1994) Developmental systems and evolutionary explanation. J Philos 91:277–304

Griffiths PE, Machery E (2008) Innateness, canalization, and ‘biologicizing the mind’. Philos Psychol 21:397–414

Haidt J (2007) Enlightenment 2.0 requires morality 2.0. Beyond Belief: Enlightenment 2.0. http://thesciencenetwork.org/programs/beyond-belief-enlightenment-2-0/jonathan-haidt

Haidt J (2012) The righteous mind: why good people are divided by politics and religion. Pantheon Books, New York

Haig D (2007) Weismann rules! OK? Epigenetics and the Lamarckian temptation. Biol Philos 22:415–428

Hauser MD, Chomsky N, Fitch WT (2002) The faculty of language: what is it, who has it, and how did it evolve? Science 298:1569–1579

Henrich J (2016) The secret of our success: how culture is driving human evolution, domesticating our species, and making us smarter. Princeton University Press, Princeton, NJ

Heyes C (2012a) Grist and mills: on the cultural origins of cultural learning. Philos Trans R Soc B 367:2181–2191

Heyes C (2012b) What’s social about social learning? J Comp Psychol 126:193–202

Horner V, Whiten A (2005) Causal knowledge and imitation/emulation switching in chimpanzees (*Pan troglodytes*) and children (*Homo sapiens*). Anim Cogn 8:164–181

Jablonka E, Lamb MJ (2005) Evolution in four dimensions: genetic, epigenetic, behavioral, and symbolic variation in the history of life. MIT Press, Cambridge

Kuo ZY (1932) Ontogeny of embryonic behavior in Aves. IV. The influence of embryonic movements upon the behavior after hatching. J Comp Psychol 14:109–122

Lehrman DS (1953) A critique of Konrad Lorenz’s theory of instinctive behavior. Q Rev Biol 28:337–363

Lewens T (2015) Cultural evolution: conceptual challenges. Oxford University Press, Oxford

Lorenz K (1937/1970) The establishment of the instinct concept. In: Studies in animal and human behaviour, vol 1. Harvard University Press, Cambridge, pp 259–315

Lorenz K (1965) Evolution and modification of behavior. University of Chicago Press, Chicago

Lorenz K (1977) Behind the mirror: a search for a natural history of human knowledge. Methuen, London

Mameli M (2008) On innateness: the clutter hypothesis and the cluster hypothesis. J Philos 105:719–736

Mameli M, Bateson P (2006) Innateness and the sciences. Biol Philos 21:155–188

McGuigan N (2012) The role of transmission biases in the cultural diffusion of irrelevant actions. J Comp Psychol 126:150–160

McGuigan N, Makinson J, Whiten A (2011) From over-imitation to super-copying: adults imitate causally irrelevant aspects of tool use with higher fidelity than young children. Br J Psychol 102:1–18

Millikan RG (1984) Language, thought, and other biological categories: new foundations for realism. MIT Press, Cambridge

Millikan RG (1989a) An ambiguity in the notion “function.” Biol Philos 4:172–176

Millikan RG (1989b) Biosemantics. J Philos 86:281–297

Millikan RG (1995) Pushmi-pullyu representations. Philos Perspect 9:185–200

Moore DS (2001) The dependent gene: the fallacy of “nature vs. nurture.” Freeman, New York

Nielsen M, Tomaselli K (2010) Overimitation in Kalahari Bushman children and the origins of human cultural cognition. Psychol Sci 21:729–736

Norris P, Inglehart R (2011) Sacred and secular: religion and politics worldwide, 2nd edn. Cambridge University Press, New York

Oyama S (2000) Evolution’s eye: a systems view of the biology–culture divide. Duke University Press, Durham, NC

Papineau D (2003) Is representation rife? Ratio 16:107–123

Pinker S (2005) So how *does* the mind work? Mind Lang 20:1–24

Richerson PJ, Boyd R (2005) Not by genes alone: how culture transformed human evolution. University of Chicago Press, Chicago

Rogers DS, Ehrlich PR (2008) Natural selection and cultural rates of change. Proc Natl Acad Sci USA 105:3416–3420

Shea N (2007) Representation in the genome and in other inheritance systems. Biol Philos 22:313–331

Shea N (2012a) Genetic representation explains the cluster of innateness-related properties. Mind Lang 27:466–493

Shea N (2012b) New thinking, innateness and inherited representation. Philos Trans R Soc B 367:2234–2244

Shea N (2013) Inherited representations are read in development. Br J Philos Sci 64:1–31

Smith LB (1999) Do infants possess innate knowledge structures? The con side. Dev Sci 2:133–144

Spelke ES (1994) Initial knowledge: six suggestions. Cognition 50:431–445

Spelke ES (1998) Nativism, empiricism, and the origins of knowledge. Infant Behav Dev 21:181–200

Spelke ES (1999) Innateness, learning and the development of object representation. Dev Sci 2:145–148

Spelke ES, Breinlinger K, Macomber J, Jacobson K (1992) Origins of knowledge. Psychol Rev 99:605–632

Sperber D (1996) Explaining culture: a naturalistic approach. Blackwell, Oxford

Sterelny K (2012) The evolved apprentice: how evolution made humans unique. MIT Press, Cambridge

Stich SP (1975) Introduction: the idea of innateness. In: Stich SP (ed) Innate ideas. University of California Press, Berkeley, pp 1–22

Tomasello M (1999) The cultural origins of human cognition. Harvard University Press, Cambridge

Wertz AE, Wynn K (2014) Selective social learning of plant edibility in 6- and 18-month-old infants. Psychol Sci 25:874–882