



King's Research Portal

Document Version

Early version, also known as pre-print

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

Citation for published version (APA):

Shea, N. (2014). Two Modes of Transgenerational Information Transmission. In K. Sterelny, R. Joyce, B. Calcott, & B. Fraser (Eds.), *Cooperation and Its Evolution* (pp. 289-312). (Life and mind: philosophical issues in biology and psychology). MIT Press. <http://mitpress.mit.edu/books/cooperation-and-its-evolution>

Citing this paper

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

General rights

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

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

Take down policy

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

Two Modes of Transgenerational Information Transmission

Nicholas Shea

Abstract

The explosion of scientific results about epigenetic and other parental effects appears bewilderingly diverse. An important distinction helps to bring order to the data. Firstly, parents can detect adaptively-relevant information and transmit it to their offspring who rely on it to set a plastic phenotype adaptively. Secondly, adaptively-relevant information may be generated by a process of selection on a reliably transmitted parental effect. The distinction is particularly valuable in revealing two quite different ways in which human cultural transmission may operate.

Contents

1. Introduction
2. Detection-based and selection-based information
 - Intergenerational Conflict*
3. Different channels of transgenerational information transmission
4. The distinction applied to hominin cultural transmission
 - Selection-based information: overimitation*
 - Detection-based information and group selection*
5. A theory of content applicable to transgenerational information transmission
 - Detection-based information*
 - Selection-based information*
6. Causal mechanisms and the information they carry
7. Conclusion

1 Introduction

Much of the organised complexity found in living things depends on communication. Many differences between forms of social organisation can be traced to differences in the way information is communicated between individuals. The complexity of an organism depends upon communication between cells, tissues and other component parts. These forms of horizontal communication rightly receive considerable attention. However, transgenerational communication is equally significant. Differences in the way that information is communicated down the generations also play an important role in explaining differences amongst living things. In particular, major innovations in transgenerational signalling have probably been especially important in recent hominin evolution.

It has long been recognised that teaching and learning transmits information between generations in humans, and that more limited forms of cultural inheritance are also found in other species. Now the burgeoning literature on epigenetic and other parental effects shows that there may be very many routes for communication between generations which go well beyond the bounds of cultural inheritance. We gain a useful perspective on these results by asking how information is being transmitted between generations in each case.

This paper argues for a distinction between two ways that adaptively-significant information may be transmitted between generations. The distinction is roughly between parental organisms blindly transmitting information and their finding out information for themselves. Both generate *correlational information* to which offspring can react adaptively. Correlational information is found whenever some entity's being in a particular state changes (usually raises) the probability that some other entity is in another particular state. Organisms can make use of correlational information because they can react to the information-carrier as they would if they had access directly to the thing with which it correlates. With detection-based transgenerational effects the parent detects a fact and transmits a correlate of it to its offspring. By contrast, selection-based effects depend upon parents reliably transmitting variants without doing any detection for themselves. Natural selection on such variants then generates the correlational information on which offspring organisms can rely.

We have two tasks: to draw the distinction and to justify that it is a matter of information transmission. The philosophical material on channels of semantic information is logically prior, but much more comprehensible if the distinction is first laid out and justified empirically. So section 2 will set out the distinction and section 3 will argue that to understand the evolutionary consequences of transgenerational effects it is important to categorise mechanisms in terms of how they form channels of information transmission. Section 4 completes the task of motivating the utility of the distinction by showing how it is a useful way of thinking about hominin cultural inheritance. Sections 5 and 6 turn to the second, more obviously philosophical task. Section 5 justifies the assumption that selection-based and detection-based effects are channels of information in a substantive sense that differs from other kinds of causal interaction by showing how they carry semantic information. Section 6 argues that appealing to semantic information offers a kind of explanatory purchase that is missing from a purely causal-mechanical take on transgenerational effects.

2 Detection-based and selection-based information

This section sets out the distinction between detection-based and selection-based effects and motivates it on the basis of empirical findings. An important class of our detection-based effects are studied in the empirical literature on transgenerational phenotypic plasticity. In such cases a range of phenotypic outcomes is open to the developing organism and which one it takes depends upon an effect from a parent. For example, the herb *Campanulastrum americanum* can follow one of two life history strategies. It can either germinate in the autumn and flower the following summer (annual), or it can germinate in the spring and flower in its second summer (biennial). Which strategy is adaptive depends upon the local environment in which the plant grows. Plants that grow in shady woodland understory have higher overall fitness if they adopt the biennial strategy. Those that grow in a light gap do better if they grow as annuals, making use of the extra light to reproduce more quickly. Interestingly, experiments have shown that the strategy adopted by a seedling is affected by whether its mother grew in understory or light gap (Galloway and Etterson 2007), and that this strategy is adaptive in the wild.

This form of transgenerational plasticity makes sense since a seedling is likely to be in the same kind of light environment as its mother, given limited seed dispersal. The choice between germinating in spring or autumn has to be made at an early stage, when it is presumably difficult or impossible for the germinating seed to detect directly which form of light environment it is in. Instead, the observed plasticity seems to be driven by some form of non-genetic maternal effect that correlates with the maternal light environment. Seedlings rely on this natural sign of light environment produce an adaptive developmental response. In this respect, the case is no different than within-generation examples of adaptive phenotypic plasticity, where the developing organism picks up on correlational information in its environment in order to fix on an adaptive developmental outcome; for example when the water flea *Daphnia pulex* reacts to chemical traces of predators by growing a protective outer shell.

Wherever there is causal interaction, there is information in the thin, correlational sense. However in these cases there is information in a stronger, semantic sense, as argued for in section 5. Out of all the ways that a parent organism causally affects its offspring, there is a principled reason for singling out these cases of transgenerational adaptive plasticity as involving information transmission in a substantive sense. The offspring acts on detection-based information transmitted by the parent.

In addition to detection-based effects, some epigenetic effects are stable over very many generations. For example, a single treatment of the Ciliate *Tetrahymena* with insulin leads to a significant increase in insulin binding in progeny for at least 664 generations (Csaba 2008; see Jablonka and Raz 2009 for further examples). Natural selection can act on stably-transmitted epigenetic effects just as it can on genes. Such cases should not be assimilated to the category of transgenerational plasticity. Instead, if such epigenetic effects carry information, it will be much more like the information carried by genes (Shea 2007a). Their effects may be context-sensitive, but these are cases where differences in phenotypic outcome are not due to epigenetic variation, but rather due to variation in the environment. In a standard heritability analysis these stably-transmitted epialleles will be indistinguishable from genes (Helanterä & Uller 2010).

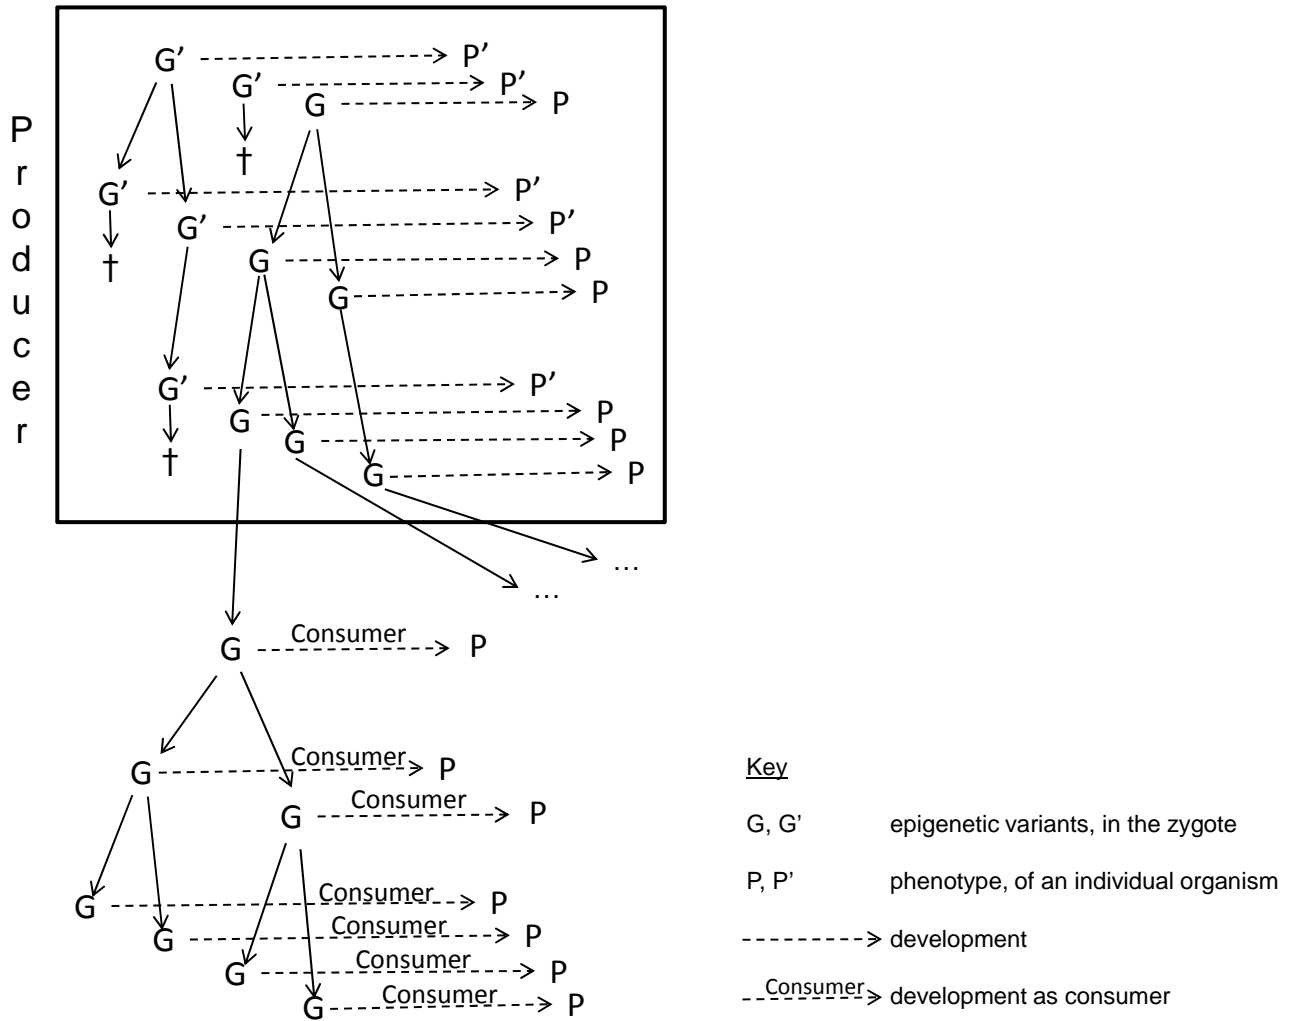


Figure 1 Epigenetic selection-based information

Information **production** takes place over many generations and involves selection. During this time, development (of phenotype P in response to epigenetic type G) is part of the process of representation production. At the outset, G arises by a random process (e.g. mutation) and leads to, and correlates with, phenotype P; at the end of representation production G correlates both with P and with some environmental factor(s) E in virtue of which G was selected (E not shown). **Consumption** of the selection-based information carried by G takes place in each subsequent generation, in every individual carrying G in the zygote and developing P as a result. (Content of G is: E is the case; develop P – see section 5.)

An initial reaction is that stably-transmitted epigenetic factors are not transmitting information at all. After all, they are unvarying, unlike the usual cases where a variable signal transmits information about a varying matter of fact. But that is a mistake. If an epigenetic factor is selected as a result of some phenotypic difference it makes, then the frequency of that factor in the population will increase. At the outset the epigenetic factor will correlate with the phenotypic result it produces. Assuming it arises at random, it will initially carry no information about the external environment. However, as a result of selection it will come to carry information about the external environment – it will correlate with whatever environmental parameter or parameters are conducive to the phenotypic effect with which it in turn correlates. The same process occurs when a genetic variant is selected. The process of selection leads genetic types to carry correlational information about features of the environment in which they were selected. To the extent that the environment has not changed in relevant respects, this information will still be useful to the developing organism.

With genes, it is not just fortuitous that the mechanisms of DNA expression and replication are acted on by natural selection. DNA has been designed by evolution to perform the function of transmitting selected phenotypes to future generations (Freeland and Hurst 1998, Haig and Hurst 1991). So the consumption of such DNA-environment correlations by the mechanisms of development, in giving rise to environmentally-appropriate phenotypes, is a matter of design (Shea 2007a). The same may be true of some mechanisms of epigenetic inheritance. They, too, may have been selected to perform the function of transmitting selected phenotypes. If so, such epigenetic factors will count as transmitting information in the semantic sense argued for in section 5 below. Unlike detection-based effects, the information they carry is due to selection. They are *selection-based effects* carrying *selection-based information*.

Sometimes an epigenetic variant will be selected because of an effect that it has within the organism, a physiological effect say, without altering the way the organism interacts with its external environment. In those cases the ‘environment’ about which the epigenetic variant carries information (as a result of selection) is the existing physiology of the organism. If the epigenetic effect is selected because of some effect that occurs within the organism, then there is some feature or features of the organism with which that effect interacts so as to increase fitness. The epigenetic factor comes to carry information about those features.

Selection-based effects can be context-sensitive. This is still information generated by selection, but information of a more complex sort. The selected epigenetic variant instructs development to produce a range of outcomes and carries the information that the environment is likely to be conducive to phenotypes within that range. Where each of the phenotypes induced in different environmental contexts is adaptive, we can say something more specific. Then the information is roughly: the environment is likely to be E1 or E2, in E1 produce phenotype P1 and in E2 produce P2. Selection on stably transmitted epigenetic factors could be the basis of such adaptive plasticity, which would be epigenetic accommodation of a plastic phenotype just as there can be genetic accommodation of a plastic phenotype (Gilbert and Epel 2009). A particular adaptive phenotype (P1) would then result from a combination of selection-based information about the likely range of environments and detection-based information about the specific environment (E1) within which a particular organism is developing. The latter could be detected by the developing organism itself or, as with *Campanulastrum*, could be detected by a parent and transmitted to the developing offspring.

With the basic distinction between detection-based and selection-based information in hand, we can revisit *Campanulastrum*. Why does its form of developmental plasticity rely on a maternal effect rather than a genetic cue? Modelling work by Leimar et al. (2006) suggests that that question should be answered by examining the quality of the information available through those different routes. Genes would be unable to build up correlational information at the available timescale. Paternal DNA is dispersed by the wind and so is drawn from a much larger area than the spatial scale on which the variation between light gap and understory is found, making paternal information about the light environment unreliable. And lineages of offspring do not remain in the same light environment as their mothers for sufficiently many generations for natural selection to act and build up maternally-inherited information about being in a light gap, say (either in genes, or in long-term maternal epigenetic effects). In short, the structure of the problem suggests that detection-based information from the mother will be the most reliable source of correlational information about the local light environment, if a mechanism for conveying such signals is available – as it seems to be. The literature on the evolution of individual and social learning also focuses on the circumstances in which detection-based information (of various sorts) is superior to selection-based information (Boyd and Richerson 1985, Feldman et al. 1996, Sterelny 2003, pp. 162-171).

In addition to selection-based information that produces particular adaptive phenotypes, there are other cases where epigenetic mechanisms serve to produce undirected variation in offspring phenotypes. Typically environmental stress can trigger epigenetically-caused variability in offspring. For example, oxidative stress or heat shock can directly alter epigenetic marks in the germ line, leading to random variations in offspring phenotypes. Such variability can be adaptive for the parent, and while there may be cases where it is an adaptation (a bet-hedging response to a change in environment), in many cases the epigenetic variability is just a direct causal effect of the environmental stress and not the result of selection. Such cases would not then count as cases of the parent transmitting adaptively-relevant semantic information to the offspring.

Intergenerational Conflict

Detection-based epigenetic information transmission raises the same issues about the possibility of honest signalling as are canvassed in the existing literature on animal signalling. If the interests of sender and receiver diverge, what prevents the sender from sending dishonest signals? But if the signals are unreliable, then receivers would not have a reason to react to the signal with an informationally-relevant response. How then could a signalling system ever evolve? This problem is overcome if the signalling occurs in a co-operative context, so that the interests of the sender and receiver are aligned. For example, non-reproductive worker bees all have the same reproductive interest in their shared hive. These obstacles are of course overcome in many cases of animal signalling that are not fully co-operative, as explained by factors such as costly signalling of fitness, or production of signals that are otherwise hard to fake, or where dishonest signalling imposes large developmental costs. Should we expect similar obstacles to the evolution of transgenerational detection-based information transmission?

The fitness interests of parents and offspring are usually closely-aligned in most respects. That provides the basis for co-operative signalling about many aspects of the

environment. However, they do come apart in non-signalling contexts, for example where the parent has limited resources to distribute amongst its progeny. Parental fitness will often be best served by distributing resources equally amongst several progeny, whereas each one of those individuals has a stronger fitness interest in getting an unequally large share of the resources.

Something similar may occur with respect to offspring of offspring. Suppose that, in resource-poor environments, it is adaptive for a mother to limit the number of offspring produced in turn by each of her progeny. When resources are in short supply it might be in the mother's best interests for each daughter to have on average m offspring, whereas each daughter's individual fitness interest would be to have $n > m$ offspring. The mother could send a signal to its offspring that correlates with whether the environment is likely to be resource-poor or resource-rich. But would the offspring react to the signal by producing the behaviour that is adaptive for the parent? Modelling work strongly suggests that they would not, unless the mother has a means to take control of the relevant aspect of her offspring's development and canalize it towards her preferred result [Uller and Pen, in submission]. It is easy to see why. Offspring will evolve to react to correlational information in a way that is adaptive for themselves (cp. the non-informational treatment in Hinde et al. 2010). The result is that, as with regular animal signalling, transgenerational detection-based epigenetic information transmission will not evolve in non-cooperative contexts unless there is some special mechanism in place to deal with the conflict between how the receiver can best act in the light of a piece of correlational information, on the one hand, and how the sender would like the receiver to act on that information, on the other.¹

3 Different channels of transgenerational information transmission

Selection-based effects and detection-based effects are two ways of transmitting information between generations. This way of categorising cross-cuts standard ways of dividing up epigenetic effects based on the types of mechanism involved. In this section we argue that to understand the different channels by which information flows between generations, and their evolutionary consequences, it is important to consider the way information is generated and transmitted as well as the mechanisms that are involved.

The transgenerational information flow is fundamental to the evolution of complexity. The process of differential survival and reproduction – selection – generates information about which phenotypic variants are better suited to the environment in which they exist. Life can evolve only if that information is somehow transmitted to subsequent generations. DNA and its associated machinery of transcription, translation and replication is a wonderful solution to that problem. But it is becoming clear that there are several other mechanisms of heredity (Jablonka & Lamb 1995, 2005). Jablonka and Lamb argue that information should be used as the common denominator by which different systems of heredity are compared (2007, p. 382).

¹ Where the mother takes control of an aspect of the offspring's development in her own interests, that would not be a case of semantic information transmission of the type considered here. Those cases are better treated as manipulations (Dawkins and Krebs 1978).

One of the themes of this volume is that communication and signalling is an evolutionary challenge to all, not only to cognitively sophisticated creatures like humans. Even the simplest forms of multicellularity depend upon efficient signalling between cells to solve co-ordination problems. A parallel point can be made about transgenerational information transmission. It used to be thought that genetic inheritance was the only mechanism of transgenerational information transmission in all but the most cognitively sophisticated creatures. It is now becoming clear that there are many more mechanisms of transgenerational information transmission than just DNA, on the one hand, and human culture, on the other. Following Jablonka & Lamb's suggestion, an appreciation of the different ways in which information is transmitted across the generations can help us to understand aspects of various branches of the tree of life.

As we saw in the last section, influences that are unified mechanistically under the label 'parental effects' can generate and transmit information in different ways. Reliably transmitted epigenetic factors can generate selection-based information that works like genes, acting as a channel of inherited information generated through a process of selection. Epigenetic factors probably operate on a different timescale to genes, due to different mutation rates and levels of transmission fidelity. For example, some forms of antibiotic resistance in *E. coli* works by epigenetically-based selection (Adam et al. 2008), suggesting that some epigenetic mechanisms are well-suited to rapid evolution of adaptation.

When detection-based information is transmitted from parents to offspring, parental effects operate at the timescale of a single lifetime, with the parent detecting some fitness-relevant environmental parameter and the offspring relying on a signal from the parent to fix a developmental parameter for life. (Of course, for such effects to be adaptive, there must also have been natural selection on the intergenerational disposition; but the epigenetic signals themselves vary generation-by-generation.) So epigenetic mechanisms bring in information channels that operate in parallel to genes over at least two different timescales.

The animal signalling literature fills in the story at the shortest timescale, where organisms signal to one another about transient matters of fact, with receivers producing not a fixed response that lasts a lifetime, but a transient response that is appropriate to the particular fact signalled (e.g. fleeing in response to the information that there is a predator overhead now). Of course, signals at this timescale can be sent between generations as well. Parents signal significant pieces of transient information to their offspring. Given their highly overlapping reproductive interests, there are many cases of co-operative communication of transient matters of fact between generations.

4 The distinction applied to hominin cultural transmission

Section 2 distinguished between two different ways that information is transmitted between generations by epigenetic factors. In the last section we argued that it is useful for all kinds of inheritance mechanisms, genetic and non-genetic, to analyse the way information is generated and transmitted between generations. In this section we will follow that advice in examining hominin cultural transmission.

Although it is now clear that there are many non-genetic channels of transgenerational information transmission throughout the natural world, it remains true that humans are

particularly well-endowed with means for sending information between generations. Humans have far richer cultural traditions than any other species, showing that transgenerational transmission of phenotypes has an especially central role for us. In this section we argue that the distinction we drew above between selection-based and detection-based information is particularly germane to understanding human cultural transmission.

It is tempting to think that human cultural learning all falls in the category of detection-based information. Individuals learn about some aspect of their environment (how to exploit some resource, how to perform some skill, how to turn some raw material into a useful tool), and then pass on that knowledge to conspecifics, including those in the next generation. If those individuals in turn pass on their knowledge, then this pattern of teaching and learning will produce cultural traditions. But the utility of the behavioural phenotypes is usually thought to be detectable by the individuals who learn them. Unlike the information carried by genes, or by long-term epigenetic effects, the generation of useful information does not itself depend upon a history of selection. The standard idea is that, at some point, an individual learns for itself something useful. That information can then be passed on and preserved by cultural inheritance, so that it does not need to be re-learned anew by the individuals who come after. When the first individual uncovers a useful fact or hits on a clever trick, perhaps by trial and error learning, that individual's behavioural phenotype immediately carries correlational information about a fitness-relevant aspect of the environment. A process of individual learning has generated that information.

However, not all cultural traditions work like that. Suppose offspring were to copy the behavioural patterns of their parents blindly, without ever assessing the utility of those forms of behaviour. Then behavioural phenotypes would follow reproductive lineages. If there were fitness differences between individuals with different behavioural phenotypes, then those with higher relative fitness would be selected. Even if novel behavioural phenotypes were introduced at random, such a process would give rise to selection, and so generate correlational information about the selective environment. In short, given reasonably high-fidelity vertical or oblique transmission, selection on behavioural phenotypes would work, informationally, just like genetic transmission and long-term epigenetic effects. That is, behavioural phenotypes would carry selection-based information down the generations.

The distinction between these two broad ways that cultural learning may operate, which are often elided, emerges clearly when we take Jablonka and Lamb's advice and use information as a common denominator. If individuals generate exploitable correlational information through individual learning then the process of cultural inheritance is Lamarckian, with useful characteristics acquired in individual ontogeny then passed down the generations. The role played by individual learning also makes the transmission process inherently noisy, as individuals in subsequent generations go on learning for themselves. If the exploitable correlational information depends upon a history of selection, then the process loses its directed character. A potential for higher-fidelity transmission is balanced by a reduction in the speed with which the system can respond to selection pressures. Nevertheless, even behavioural traditions that depend on selection – those which carry selection-based information – can likely be selected on a much more rapid timescale than gene-based selection. (Compare the different timescales over which selection generates information for genetic vs. epigenetic inherited information, as we saw above.)

Selection-based information: overimitation

There is evidence that some forms of human imitation-based learning do operate so as to carry selection-based information. The baseline is that both humans (Meltzoff 1988) and other primates (Voelkl and Huber 2000) are capable of copying the means by which a demonstrated action is performed, and of doing so rationally, appreciating the difference between means to an intended goal and irrelevant actions (Horner and Whiten 2005, Buttleman et al. 2007, Wood and Hauser 2008, Tennie et al. 2009, cf. Povinelli 2000).

Curiously, humans have a disposition to imitate even actions which seem to them to be irrelevant to achieving the demonstrated goal. When chimpanzees are shown a series of actions performed on an opaque box they copy them all so as to obtain a reward, as do children. A transparent version of the same apparatus makes it obvious that some of the demonstrated steps are irrelevant to achieving the reward. Chimps omit unnecessary steps but children copy them (Horner and Whiten 2005, Lyons et al. 2007). There seems to be a developmental trajectory in which infants (12 and 18 months old) are poor at copying specific actions and are more focused on outcomes (Nielsen 2006), 3 year olds are capable of copying specific actions (Williamson et al. 2010) but less inclined to engage in overimitation, and 5 year olds imitate obviously irrelevant actions in a wide range of conditions (McGuigan et al. 2007, McGuigan and Whiten 2009, Whiten et al. 2009, but cf. Nielsen and Tomaselli 2010). Overimitation appears to be even more pronounced in adults (McGuigan, Makinson and Whiten 2010).

5 year olds can tell the difference between the necessary and apparently unnecessary actions, but seem to think that the latter have an unspecified purpose (Kenward et al. 2010, cp. Hernik & Csibra 2009). Their tendency nevertheless to copy these steps may just be the most rational hypothesis about the underlying causal structure (Buchsbaum et al. 2010), but the developmental results suggest that the tendency may be more pervasive than causal rationality would dictate.

It is counterintuitive that humans should engage in such ‘overimitation’ when other apes do not. Humans seem to be copying in the less intelligent way. Overimitation makes more sense when we realise that it may be the basis on which selection-based information is transmitted. The tendency to imitate actions that are registered as intentional but causally irrelevant to the outcome may be an adaptation to increase the fidelity with which behavioural phenotypes are transmitted down the generations. Arguments for and against that hypothesis are detailed in Shea (2009) and only summarised here.

Since there is a fitness cost of transmitting useful phenotypes to non-kin, our hypothesis that overimitation may be the basis of a selection-based inheritance system is only plausible if vertical transmission is somehow favoured. If it is, selection on phenotypic variants would have the power to accumulate behavioural adaptations. Evidence that human children selectively imitate in ‘natural pedagogical contexts’ (Gergely & Csibra 2006, Csibra and Gergely 2009) plausibly encourages vertical transmission, lending support to the inheritance system hypothesis.

However, the evidence is by no means one way. The importance of imitation may instead be to allow the transmission of behaviours whose utility has been discovered through individual learning – detection-based information. Studies of diffusion chains reach different conclusions about the reliability of child-to-child copying of irrelevant actions (McGuigan and Graham 2009; Whiten and Flynn 2010), although they do seem to show fidelity to a particular

method of achieving an outcome (Flynn 2008). Social transmission of a particular method of performing an action is not distinctive of hominid cultural transmission (Hobaiter and Byrne 2010), so some of these results count against the claim that there is a distinctively hominin form of imitation-based transmission of selection-based information. Instead they suggest that individual children are still behaving on the basis of an assessment of the consequences of their actions. Considered as an inheritance system, such modifications would reduce the fidelity of long-term transmission. On the other hand these studies do not replicate the chains of parent-child dyads that are essential for vertical inheritance. For now the importance of overimitation for transmitting selection-based information remains empirically open.

Detection-based information and group selection

Detection-based behavioural effects raise their own theoretical problems. Once individual learning is allowed free reign, the transmission process becomes noisy, making it hard to see how new behaviours, starting at low frequency, would be selected and then maintained in a population. Boyd and Richerson (1985, 2005) have shown how various kinds of learning biases can help to overcome these problems, for example for learners to observe several different techniques and copy the best.

The problem of honest signalling from the animal communication literature also raises concerns. If useful skills and information spreads horizontally to non-kin, it will often reduce the relative fitness of the sender (*vis a vis* those competitors). Even if spreading the information is good for both senders and receivers (in the everyday sense of ‘good for’), senders would have an incentive not to transmit useful information if keeping it secret is more detrimental to others than it is to themselves. For example, even if the sender would be healthier, on average, if everyone in the group were to adopt his clever hygiene practice, he will increase his relative fitness by keeping it to himself. There will be more disease going around for everyone, but he has a relative advantage of a behaviour that protects him a bit more than others.

So there is a puzzle about how there can be the right kinds of co-operative contexts for detection-based information transmission to evolve. A restriction to vertical transmission would help here, too; but unlike in the case of blind imitation, once agents are reasoning for themselves about how to behave it is unclear that a disposition to copy only the behaviours of kin forebears would be evolutionarily stable. A population of agents who copied detection-based information only from their parents would be invaded by individuals with a disposition to copy useful detection-based information from non-kin as well.

The problem of why information should be shared with non-kin can be overcome if there is strong group selection. Then the relative fitness advantage of keeping information secret is outweighed by the relative fitness advantage of being in a group of agents disposed to transmit detection-based information. Several of the mechanisms pointed to by Boyd and Richerson for improving the maintenance of adaptive behaviours in social groups work best given group selection. Although there is widespread scepticism about the importance of group selection in other species, it may have been particularly important in human evolution. The increasing richness of human culture in our recent evolutionary past may have gone hand in hand with the operation of a powerful form of group selection between small tribes.

We saw above that vertical or oblique overimitation of behavioural phenotypes by kin could have formed a channel of selection-based information in hominid evolutionary history. In

this section we have argued that this is not the exclusive channel of cultural inheritance. Socially-mediated learning may also transmit detection-based information about behavioural phenotypes, if the right kind of social structure is in place to underpin group selection.

So it seems likely that recent hominin evolution has involved both detection-based and selection-based information transmission. Their respective contributions are still unclear. The answer has implications for the timescale on which cultural changes can respond to selection pressures, and the kinds of mechanisms that may be in place to support the transmission of information down the generations. Asking Jablonka and Lamb's question about how information is generated and transmitted in cultural inheritance produces the insight that phenomena that are standardly grouped together as being based on social learning might actually operate in ways that have significantly different evolutionary consequences. Much work has implicitly assumed that social learning transmits detection-based information. Our focus on selection-based culturally-mediated information is a useful counterweight. It may make sense of empirical phenomena like the human tendency towards overimitation. Periods of apparent stasis in evolutionary history of hominin tools or behavioural practices may also be due to their being based on the relatively high-fidelity transmission of selection-based information.² By contrast, where we find behavioural phenotypes and cultural practices that seem to vary widely between populations but spread rapidly within a population, it is more likely that detection-based information has mediated their transmission.

5 A theory of content applicable to transgenerational information transmission

The sections above have set out the argument that a distinction between selection-based and detection-based information transmission can do useful explanatory work. It is now time to pay out on the promissory note – to show that it is legitimate to treat these as channels of information transmission in a substantive sense, over and above the mere existence of causal correlations. In this section we argue that both cases involve the transmission of semantic information. That invites the question: 'but what does semantic information add to a full causal understanding of these interactions?' To answer that we must first set out the framework that gives rise to semantic information. That is the task of this section, before turning in the final section to show how this semantic information plays an explanatory role over and above the explanations afforded by underlying causal structure.

If the causal intercourse between parents and offspring described above was nothing more than chains of causal interactions, then it would remain true that we could describe these interactions using the tools of information theory, but there would be no justification for thinking that there was any way of distinguishing between channels of information transmission and other kinds of causal interactions. All the myriad ways in which parents affect offspring carry information in the thin causal sense, whether or not they are adaptive or even heritable, as do causal interactions between an offspring and all other aspects of its environment. Our claim

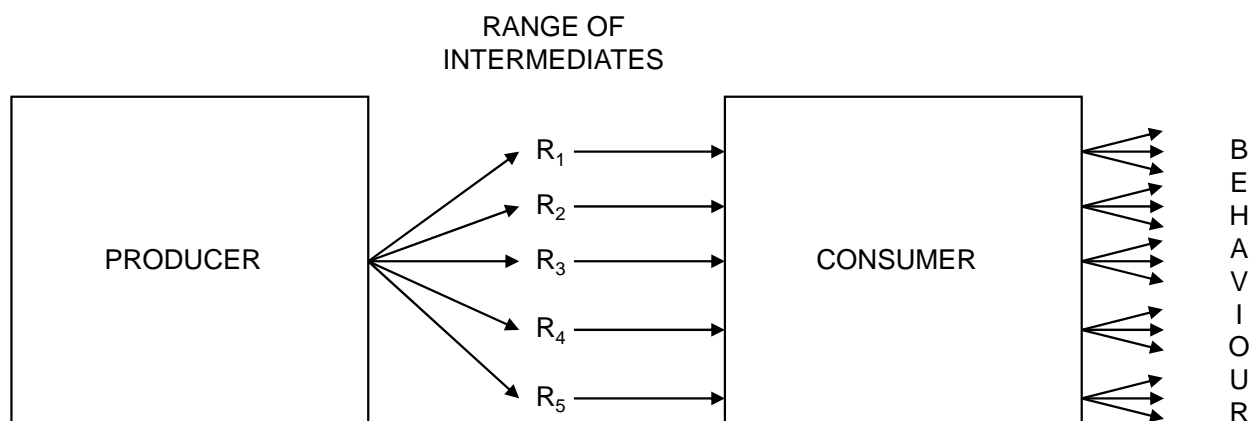
² Haldane (1956, p. 9) suggested that the stability of the Acheulean biface could have been due to its being transmitted by selection-based information. He had in mind a genetically transmitted instinct, but culturally-transmitted selection-based information could have a similar effect.

that it is useful to distinguish between two different ways that information is generated and transmitted in inheritance would be correspondingly thin were this not information in some more substantial sense.

The stronger, substantial sense is furnished by infotel semantics. Infotel semantics is a philosophical framework that characterises systems that are adaptations for information transmission. It shows how some information-carrying elements in the system are representations with semantic content. Having semantic content is a matter of there being correctness conditions or satisfaction conditions attached to each carrier of information. This is significant for philosophical purposes because it differentiates representations from mere correlations. For our purposes the importance of the framework is to delineate those cases where there is an adaptation to transmit information in certain ways, which gives substance to the claim that these are channels of information transmission. So we argue for the claim that detection-based effects and selection-based effects are based on information channels by showing how they involve representations with semantic contents.

Detection-based information

The framework of infotel semantics is well suited to capturing the semantic content of low-level varieties of representation, like those found in animal signalling and in inheritance systems (Shea 2007b). It is a development of the teleosemantics pioneered by Millikan (1984) and Papineau (1987), but with an additional requirement that representations should correlate with the facts that they represent (i.e. that they should also carry information in the ‘thin’ sense of information theory).³ The framework applies when there is some producer system that produces a range of different natural signs and a consumer system that responds to those natural signs, behaving in a characteristic way or ways with respect to each.



³ Standard teleosemantics bases content purely on the way that a consumer system is disposed to react to a putative representation and so is indifferent to whether representations transmit information of the correlational sort.

An intermediate carries correlational information if its occurrence affects the probability that some further state of affairs obtains. Correlational information is cheap, so the natural signs sent by a producer system will typically carry correlational information about very many states of affairs. An intermediate has content, however, only if the consumer system reacts to it with a form of behaviour that is appropriate given the obtaining of one of the states of affairs with which it correlates. In the cases we are interested in, the idea of the consumer system producing an output that is appropriate given the information carried by an intermediate should be understood in evolutionary terms. Behaving in that way is one of the evolutionary functions of the consumer system; for that function to lead to survival and reproduction of the consumer in an evolutionarily normal way a specific condition has to obtain; and the intermediate carries correlational information about that condition. In short, the content of the intermediate is that condition the production of the intermediate correlated with during a history of selection and the obtaining of which explains why the specific forms of behaviour prompted by that intermediate led to the survival and reproduction of the consumer system.

Less abstractly, consider how this applies to systems of animal communication. Suppose female frogs preferentially mate with males that produce low pitch calls. In the animal communication literature, researchers substantiate the idea that communication is occurring by showing:

- (i) 'that call frequency is well correlated with caller size'
- (ii) 'that females show a behavioural preference for calls of lower frequency'
- (iii) 'that size is what matters to the receivers' [e.g.] 'that female frogs benefit from mating with larger males but not from mating with older ones'

(Searcy & Nowicki 2005, p. 3)

According to infotel semantics the content of a signal is fixed by a set of conditions that match the factors listed above: correlation on the input side and consumption for content on the output side – consumption for content being a matter of the condition under which the behavioural response to the signal will be evolutionarily beneficial.

The infotel framework carries over straightforwardly to cases of detection-based information. Consider the epigenetic signal of light environment sent by maternal *Campanulastrum*. It correlates with whether the mother plant was in an understory or light gap environment, and hence whether the progeny are likely to be in understory or a light gap. The developing seedling responds to that signal by germinating in the spring or in the autumn, with the respective outputs showing relative fitness benefits if spring germination occurs in understory and if autumn germination occurs in a light gap. So we have correlation on the input side (maternal signal of light gap) and consumption for content on the output side (offspring behaviour that is adaptive on condition that it is in a light gap), with the input correlation matching the output condition. In addition to ascribing indicative contents – *that male frog is large* or *the local environment is sunny* – infotel semantics also ascribes imperative contents to these kinds of signals, corresponding to the functions that they serve. For example, the epigenetic signal in *Campanulastrum* carries the imperative content *grow as an annual* or *grow as a biennial*.

The cultural cases may also include adaptations for information transmission. Where individuals are learning for themselves then they are using a learning mechanisms which is an adaptation for picking up on correlational information. Where that correlational information


happens to have been generated by their parents then it will fall into our category of transgenerational detection-based information. In some cases the connection between parent and offspring may be tighter, for example if there have been adaptations for teaching and learning (a ‘pedagogy complex’, see Sterelny forthcoming).

For exposition, it is simplest to think in terms of simple correlations between natural signs and the external affairs they signify: some epigenetic effect is produced when and only when the maternal plant is in the understory and a different epigenetic effect is produced when and only when the mother is in a light gap. The framework does not depend on this simplification. The correlational information carried by a sign emitted by a sender will often consist in a joint probability distribution between a range of properties the sign can have and a range of possible external states of affairs. These correlations can be captured by Shannon information theory, but also by other means, such as by the Kullback-Leibler divergence (Cover and Thomas 2006) which measures how informative the range of states of one system (with their probabilities) are about the range and probability distribution of states of another system. My umbrella term ‘correlational information’ is intended to be neutral as between the various ways of capturing such correlations.

The treatment in terms of a single correctness condition or satisfaction condition is also a dispensable idealisation, albeit one that is innocuous in very many applications. In some cases there may be no single condition the obtaining of which explains the fitness advantage of the behaviour produced in response to some natural sign. Just as a particular sign is informative about the probabilities of a range of states of affairs, the behavioural output produced in response to the sign may have aided survival and reproduction in a normal way in a range of different conditions, each of which are amongst the states of affairs with which the sign correlates. We should give the content of such signals not in terms of a single correctness condition, but in terms of a set of conditions and a degree of probability or credence which the system attaches to each.

Selection-based information

Application of the infotel framework to selection-based information involves an additional subtlety. In these cases correlational information is generated through a process of selection on reliably-transmitted epigenetic variants over many generations. Once natural selection has generated correlational information, that information is reliably transmitted down the generations in the germ line of chains of individuals. The consumer is the individual developing organism, and the transmission chain also consists of a series of individuals. However, the producer in the infotel model is realised by the process of selection over many generations.



Infotel semantics delivers contents when there is a producer system designed to send a variety of different signals and a consumer system designed to react to the various signals with different outputs. With reliably-transmitted epigenetic variants there are genuine producers and consumers designed to send and respond to such transgenerational effects only if it is an evolutionary function of that family of epigenetic variants that they should be reliably transmitted down the generations. It is not enough that they lead to heritable effects so that selection can in fact build up correlational information. We want to claim that there is a producer system designed to send natural signs which correlate with some environmental state of affairs. If the producer-consumer system has the evolutionary function of transmitting heritable phenotypes, then the generation of correlational information in a period of selection is a particular instance of the performance of that function, so will constitute the operation of a real producer system in generating correlational information.

It is reasonably clear that DNA has the relevant function (Shea 2007a) – a meta-function, in addition to its myriad object-level functions of producing particular selected phenotypes. The system of DNA transcription, translation and replication has evolved in order effectively to fulfil the function of producing heritable phenotypes. It is less clear whether any non-genetic mechanisms have evolved in order to produce heritable phenotypes (rather than just happening to be a basis on which some heritable phenotypes arise). The evidence discussed above that some features of the human imitation appear to be adaptations for the high-fidelity transmission of behavioural phenotypes is an indication that human blind over-imitation, at least, may also meet the criterion.

Given the metafunction, the rest of the framework applies directly. A range of different messages are produced by episodes of selection in different historical epochs. The consumer (the developing individual) responds to each with a particular behaviour or set of behaviours. That behaviour is adaptive in the environmental conditions in which it was selected; and the result of selection is that the signal sent between generations carries correlational information about that environmental condition.

In short, infotel semantics delivers a constrained, substantive sense in which our transgenerational selection-based effects and detection-based effects are based on channels of information transmission.

6 Causal mechanisms and the information they carry

The last section vindicated the claim that selection-based and detection-based transgenerational effects are based on channels for transmitting semantic information. This section highlights the kind of explanatory payoff that can be achieved by focusing on the semantic information that is transmitted in these interactions. The point applies equally to detection-based and selection-based information, so to keep things simple we'll focus on detection-based information.

With any kind of explanation that appeals to representational contents there is always a question about the explanatory work that is done by adverting to the fact that various states bear contents. What does it add to a causal explanation in terms of non-contentful properties of those states? Returning to the example of the mate choice of female frogs, it looks as if researchers had a complete causal story once they discovered that females respond to low pitch calls with mating behaviour. The causal sequences are:

- (a) male emits lower pitch call → female hears lower pitch call → female mates with male
- (b) male emits higher pitch call → female hears higher pitch call → female doesn't mate with male

Even if it's true that, according to a defensible theory of content, the low pitched call has the representational content *my body size is large*, what additional explanatory purchase does that provide?

The answer is that the content-based story addresses a further explanandum: why is the system set up to have that causal structure? Why do females choose mates based on the pitch of male songs? Of course the answer is that pitch correlates with body size and, in the evolutionary past, females who mated with larger males had a higher fitness. So we can understand why the signalling system is set up as it is when we see the female responding to the male call in virtue of its content. We also get a corresponding explanation of a particular episode of behaviour. Why did *that* female engage in mating behaviour with *that* male? Because she was responding to a representation of his body size. We get *an* explanation of her behaviour by citing the causal regularities (a) and (b) above and noting that this bit of female behaviour exemplified causal regularity (a). But we get a deeper explanation when we advert to the representational content of the call.

Since we have offered a naturalistic reduction of content in these simple systems, content is nothing over and above a particular conjunction of producer, consumer, correlational information and evolutionary function. But that conjunction is a real property that is important in nature. Explanations in terms of the conjunction identify real patterns that generalise across a wide range of cases. The reduced property (semantic information) is ontologically real and explanatorily useful.

The same dialectic applies to detection-based transgenerational effects. The experimental literature on epigenetic effects focuses on demonstrating the existence of such effects and documenting their robustness. So we observe that progeny of *Campanulastrum americanum* that grow in light gaps tend, in their turn, to grow as annuals (Galloway and Etterson 2007). That causal knowledge offers us some explanatory purchase. Why did *that* plant grow as an annual? Because it grew from a seed from a plant in a light gap. However, the informational treatment offers additional explanatory purchase. It fits with the other half of Galloway and Etterson's research on *Campanulastrum*, which investigates the adaptive significance of that strategy. The causal story is supplemented when we understand that the annual strategy is an adaptive response to an epigenetic signal that correlates with a fitness-relevant fact about the environment.

The point can be made most cleanly by reference to abstract models of signalling systems. Brian Skyrms (2010) has adapted the sender-receiver games discussed by Lewis (1969) and modelled them in terms of evolutionary game theory. The sender has private access to some fitness-relevant facts about the environment and the receiver can perform some actions whose fitness may depend upon the environment (see table below). Senders can send a range of messages which the receiver can react to. A sender strategy is a rule for sending messages, which can be conditional on the environment. A receiver strategy is a rule for action, which can be conditional on the message. Senders and receivers jointly receive the payoffs from an

interaction. Agents are endowed with strategies consisting of a pair of sender and receiver strategies and then the system is allowed to evolve using replicator dynamics.

Example of a sender strategy	
<i>Environment</i>	<i>Signal</i>
E1	S1
E2	S2

Example of a receiver strategy	
<i>Signal</i>	<i>Action</i>
S1	A1
S2	A2

Payoffs for both sender and receiver
A1: 1 in E1, 0 in E2
A2: 0 in E1, 1 in E2

A *signalling system* is a complementary set of strategies where players always achieve a payoff of 1. For example, we can see that the combination of the sender and receiver strategies illustrated above constitutes a signalling system:

In E1, sender sends signal S1, receiver acts on S1 to do action A1. Payoff: 1.

In E2, sender sends signal S2, receiver acts on S2 to do action A2. Payoff: 1.

Skyrms' modelling work shows that, under standard assumptions, the signalling systems are attractors whose basins of attraction cover almost all the possible initial distributions of strategies. There are many other combinations of strategies that are less adaptive; for example the sender can send the same signal irrespective of the environment, or the receiver can perform the same action irrespective of the signal.

The important result from the modelling work is that embedding sender-receiver strategies in a selective regime where there are payoffs for actions, with payoffs conditional on the state of the environment, a certain kind of communication system is likely to evolve. Not only will the strategies evolve to maximise the transmission of information, in the sense that the sender will correlate its signal to the state of the environment and the receiver will correlate its action to the state of the signal. But they will also correlate in a particular way. Correlations are being efficiently transferred in the pure anti-signalling case, too, where the sender correlates its signals to the state of the environment and the receiver acts on them conditionally to perform the action that is not rewarded in those circumstances:

Anti-signalling

In E1, sender sends signal S1, receiver does action A2. Payoff: 0.

In E2, sender sends signal S2, receiver does action A1. Payoff: 0.

From the point of view of information theory, for example of the Shannon variety, there is maximum information flow in both cases. Representational content marks out the case where there is transmission of correlation in the service of maximising payoffs. The payoff structure allows us to see why the evolutionary dynamics tend to result in pairs of sender-receiver strategies that both (i) transmit correlational information, and (ii) do so in ways that help the pair.

In the kinds of cases considered here, representational properties arise where there is a certain internal structure – transmission of correlation – that is embedded in the context of success conditions for its outputs. Therefore, when we explain some behaviour in terms of the content of a representation which caused it we get a different, and in some ways deeper explanation than we would have if we were just to appeal to causal regularities about the operation of the system. Content depends upon the mechanism in question being embedded in

a structure of payoffs, fitness, success or some such. Correspondingly, we can explain the success of particular pieces of behaviour by the fact that a system is responding to some true content.

7 Conclusion

Within the broad category of adaptive transgenerational effects an important distinction should be drawn between selection-based and detection-based information. With detection-based effects an individual in one generation detects some fitness-relevant parameter and conveys correlational information about it to individuals in the next generation, who react with a fitness-relevant output (phenotype or behaviour). With selection-based effects the fitness-relevant information is instead generated by a process of natural selection. The distinction is particularly illuminating when applied to information transmission in human cultural traditions.

Acknowledgements

The author would like to thank audiences at the Australian National University (RSSS), Oxford Brookes University and Bristol University; and Brett Calcott, Ben Fraser, Richard Joyce, Kim Sterelny and Tobias Uller for helpful comments. This research was supported by the Oxford University Press John Fell Research Fund, the Oxford Martin School, Somerville College (Mary Somerville Junior Research Fellowship) and the Wellcome Trust (grant 086041 to the Oxford Centre for Neuroethics).

References

- Adam, M., Murali B., Glenn, N. & Potter, S.S.. 2008. Epigenetic inheritance based evolution of antibiotic resistance in bacteria. *BMC Evol. Biol.* **8**:52.
- Boyd, R. and Richerson, P.J. 1985: *Culture and the Evolutionary Process*. Chicago: University of Chicago Press.
- Boyd, R. and Richerson, P.J. 2005. *The origin and evolution of cultures*. Oxford University Press, Oxford.
- Buchsbaum, D., Gopnik, A., and Griffiths, T. L. (2010). Learning from actions and their consequences: Inferring causal variables from continuous sequences of human action. *Proceedings of the 32nd Annual Conference of the Cognitive Science Society*.
- Buttelmann, D. Carpenter, M., Call, J., Tomasello, M. 2007. Enculturated chimpanzees imitate rationally. *Dev. Sci.*, 10, F31–F38.
- Cover, T. M., & Thomas, J. A. (2006). *Elements of Information Theory* (2nd ed.). Hoboken, NJ: John Wiley & Sons.
- Csaba G. 2008. Hormonal imprinting: phylogeny, ontogeny, diseases and possible role in present-day human evolution. *Cell Biochemistry and Function* 26(1):1–10.
- Csibra, G., & Gergely, G. (2009). Natural pedagogy. *Trends in Cognitive Sciences*, 13, 148-153.
- Dawkins, R., and Krebs, J. R. (1978). Animal signals: information or manipulation. In J. R. Krebs and N. B. Davies (eds.), *Behavioural Ecology: An Evolutionary Approach*, Oxford: Blackwell.

- Feldman, M. W., Aoki, K., & Kumm, J. (1996). Individual versus social learning: Evolutionary analysis in a fluctuating environment. *Anthropological Science*, 104, 209-231.
- Flynn, E. (2008). Investigating children as cultural magnets: Do young children transmit redundant information along diffusion chains? *Philosophical Transactions of the Royal Society of London, Series B: Biological Sciences*, 363, 3541-3551.
- Freeland, S. J., and Hurst, L. D. (1998). The genetic code is one in a million. *Journal of Molecular Evolution*, 47, pp. 238-248.
- Gergely, G., Csibra, G. 2006. Sylvia's recipe: The role of imitation and pedagogy in the transmission of cultural knowledge In S. Levenson & N. Enfield (Eds.) *Roots of Human Sociality: Culture, Cognition, and Human Interaction*, Oxford: Berg Publishers
- Gilbert, S.F. & Epel, D. 2009. *Ecological developmental biology. Integrating epigenetics, medicine and evolution*. Sinauer Associates Inc, Sunderland, MA.
- Haig, D. and Hurst, L. D. (1991). A quantitative measure of error minimization in the genetic code. *Journal of Molecular Evolution* 33, pp. 412-417.
- Haldane J.B.S. (1956). The argument from animals to men: an examination of its validity for anthropology. *The Journal of the Royal Anthropological Institute of Great Britain and Ireland*, 86(2): 1-14.
- Helanterä, H. and Uller, T. 2010. The Price equation and extended inheritance. *Philosophy and Theory in Biology* 2: e101.
- Hernik, M. and Csibra, G. (2009). Functional understanding facilitates learning about tools in human children. *Current Opinion in Neurobiology*, 19:34–38.
- Hinde, C. A., Johnstone, R. A. and Kilner, R. M. (2010). Parent-offspring conflict and coadaptation. *Science*, 327, 1373-6.
- Hobaiter, C., and Byrne, R. W. (2010). Able-Bodied Wild Chimpanzees Imitate a Motor Procedure Used by a Disabled Individual to Overcome Handicap. *PLoS ONE* 5(8): e11959. doi:10.1371/journal.pone.0011959
- Horner, V. K. & Whiten, A. 2005 Causal knowledge and imitation/emulation switching in chimpanzees (*Pan troglodytes*) and children. *Anim. Cogn.* 8, 164-181.
- Jablonka, E., & Lamb, M. J. (2005). *Evolution in Four Dimensions: Genetic, Epigenetic, Behavioral, and Symbolic Variation in the History of Life*. Cambridge, MA: MIT Press.
- Jablonka, E., & Lamb, M. J. (2007). Précis of Evolution in Four Dimensions. *Behavioral & Brain Sciences*, 30, 353–392.
- Jablonka, E., and Raz, G. (2009). Transgenerational epigenetic inheritance: prevalence, mechanisms, and implications for the study of heredity and evolution. *The Quarterly Review of Biology*, 84(2), 131-176.
- Kenward, B., Karlsson, M., and Persson, J. (2010). Over-imitation is better explained by norm learning than by distorted causal learning. *Proceedings of the Royal Society B*, published online 13 October 2010. doi: 10.1098/rspb.2010.1399
- Leimar, O., Hammerstein, P., & Van Dooren, T. J. M. (2006). A New Perspective on Developmental Plasticity and the Principles of Adaptive Morph Determination. *The American Naturalist*, 167(3), 367-376.
- Lewis, D. (1969) *Convention: A Philosophical Study*. (Oxford, Basil Blackwell)
- Lyons, D. E., Young, A. G. & Keil, F. C. 2007 The hidden structure of overimitation. *Proc. Nat. Acad. Sci. USA* 104, 19751-19756.

- McGuigan, N., and Graham, M. (2009). Cultural transmission of irrelevant tool actions in diffusion chains of 3- and 5-yearold children. *European Journal of Developmental Psychology*. doi:10.1080/17405620902858125
- McGuigan, N., Makinson, J., Whiten, A. (2010). From over-imitation to super-copying: Adults imitate causally irrelevant aspects of tool use with higher fidelity than young children. *Br J Psychol* [epub]
- McGuigan, N., Whiten, A., Flynn, E. & Horner, V. 2007 Imitation of causally-opaque versus causally transparent tool use by 3- and 5-year-old children. *Cognitive Development*, 22, 353-364.
- McGuigan, N. and Whiten, A. (2009). Emulation and ‘overemulation’ in the social learning of causally opaque versus causally transparent tool use by 23-and 30-month-olds. *J. Exp. Child Psychol.* 104, 367–381.
- Meltzoff, A. 1988. Infant imitation after a 1-week delay: Long-term memory for novel acts and multiple stimuli. *Developmental Psychology* 24:470–76.
- Millikan, R.G. 1984: *Language, Thought and Other Biological Categories*. Cambridge, MA: MIT Press.
- Nielsen, M. (2006). Copying actions and copying outcomes: Social learning through the second year. *Developmental Psychology*, 42, 555–565.
- Nielsen, M. & Tomaselli, K. (2010). Overimitation in Kalahari Bushman children and the origins of human cultural cognition. *Psychol. Sci.* 21, 729–736.
- Papineau, D. 1987: *Reality and Representation*. Oxford: Blackwell.
- Povinelli, D. 2000. *Folk physics for apes: the chimpanzee's theory of how the world works*. Oxford / New York: OUP.
- Searcy, W. A., & Nowicki, S. (2005). *The Evolution of Animal Communication*. Princeton: Princeton University Press.
- Shea, N. (2007a). Representation in the genome, and in other inheritance systems. *Biology and Philosophy*, 22, 313-331.
- Shea, N. (2007b). Consumers Need Information: supplementing teleosemantics with an input condition. *Philosophy and Phenomenological Research*, 75(2), 404-435.
- Shea, N. (2009). Imitation as an inheritance system. *Philosophical Transactions of the Royal Society B*, 364, 2429-2443.
- Shea, N. (forthcoming a). Developmental Systems Theory Formulated as a Claim About Inherited Representations. *Philosophy of Science*.
- Shea, N. (forthcoming b). Inherited Representations are Read in Development. *British Journal for the Philosophy of Science*.
- Shea, N., Pen, I. and Uller, T. (in submission). Three epigenetic information channels and their different roles in evolution.
- Skyrms, B. (2010) *Signals* (Oxford / New York, O.U.P.).
- Sterelny, K. (2003) *Thought in a Hostile World* (Oxford / Malden, MA, Blackwell).
- Sterelny, K. (forthcoming). *The Evolved Apprentice* (London / Boston MA, MIT Press).
- Tennie, C., Call, J. & Tomasello, M. 2009. Ratcheting up the ratchet effect: On the evolution of cumulative culture. *Phil. Trans. R. Soc. B*.
- Uller, T. and Pen, I. (in submission). Evolution of Maternal Effects under Conflicts, Costs and Constraints.

- Voelkl, B. & Huber, L. 2000 True imitation in marmosets. *Animal Behaviour* 60:195–202.
- Whiten, A., and Flynn, E. (2010). The Transmission and Evolution of Experimental Microcultures in Groups of Young Children. *Developmental Psychology*, 46(6), 1694-1709.
- Whiten, A., McGuigan, N., Marshall-Pescini, S., & Hopper, L. (2009). Emulation, imitation, overimitation and the scope of culture for child and chimpanzee. *Philosophical Transactions of the Royal Society of London, Series B: Biological Sciences*, 364, 2417-2428.
- Williamson, R. A., Meltzoff, A. N., and Markman, E. M. (2008). Prior Experiences and Perceived Efficacy Influence 3-Year-Olds' Imitation. *Developmental Psychology*, 44(1): 275–285.
- Wood, J. N. & Hauser, M. D. 2008 Action comprehension in non-human primates: motor simulation or inferential reasoning? *Trends in Cognitive Sciences* 12, 461-46