

Draft of paper to appear in Matthen, M., and Stephens C., eds., Handbook of the Philosophy of Biology, Volume 3 of the Elsevier Handbook of the Philosophy of Science under the general editorship of Gabbay, D., Thagard, P., and Woods, J.. Elsevier, Amsterdam.

Traits, Genes, and Coding

Michael Wheeler

m.w.wheeler@stir.ac.uk

Department of Philosophy, University of Stirling
Stirling, FK9 4LA, UK

Abstract

According to the received view in biology, genes code for phenotypic traits during development. However, there are reasons to think that the massively distributed character of the causal systems underlying development is in tension with such representational talk about genes. The main contenders from the literature that purport to establish that genes are genuine coding elements in development fail to meet this challenge. An alternative and superior strategy for understanding and justifying coding talk in development turns on the fact that the process of protein synthesis exhibits the interlocking architectural features of arbitrariness and homuncularity. However, this proposal turns out to have the radical implication that it is mRNA, not DNA, that codes. Moreover, for any of the available strategies, including the one recommended here, there is a serious and unresolved issue surrounding the attempt to extend the reach of coding talk from proteins to traits.

Key words: genetic code, genetic information, representation, development, causal spread, arbitrariness, homuncularity.

Contents

1. The Uniqueness of Genes
2. Cause for Concern
3. False Starts and Dead Ends
4. A Better Idea
5. A Bullet to Bite
6. The Reach of the Code

1. The Uniqueness of Genes

Although, in most biological circles, talk of the causes of phenotypic structure naturally invites talk of genes, *everyone* knows (or ought to) that biological development is a mightily complex process involving a vast array of causal factors, some of which are genetic and some of which aren't. Elements with developmentally decisive effects are easily discoverable in non-genetic constituencies such as the gene's surrounding metabolic context and the developing organism's environment (examples below). Thus *everyone* knows (or ought to) that genetic and non-genetic factors *interact* during development, thereby causally *combining* to produce the phenotype. Such is the received wisdom in contemporary biological thought. Nevertheless, and in spite of this *interactionist consensus* (a term I have borrowed from Sterelny and Griffiths 1999), the fact is that among all the co-contributing developmental factors, genes remain special. That, anyway, is what we're told. So what mandates this prioritizing of the gene?

This is the point at which the concept of genetic coding makes its entrance onto the theoretical stage. The view that genes, or complexes of genes, code for phenotypic traits is just as much a part of the current biological orthodoxy as the interactionist account of development, and goes hand in glove with it. Such coding talk, which is of course a species of *representational* explanation, is, if not ubiquitous, overwhelmingly common, both within the scientific community and beyond. Indeed, it is the keystone of popular views according to which the genotype as a whole should be conceived as a set of *instructions* for, a *blueprint* for, a *plan* for, a *specification* of, or a *program* for, the building of the phenotype. All the highlighted notions, while perhaps subject to subtle differences in meaning that might be important in particular contexts, depend conceptually on the idea that genes make a representational contribution to development. The idea that genes code for phenotypic traits is thus an ineliminable component of such views. Moreover, one conceptual stage back, representation-talk gets a grip only where it makes explanatory sense to think in terms of structures that carry, are vehicles for, exploit, or in some other way trade in, *information*. Whether or not one can think of structures as information-carrying (in a rich semantic sense, rather than merely in information-theoretic terms – see below) without thereby thinking of those structures as representational is, I think, a moot point. In any case I shall take it that representation-talk requires information-talk, so establishing that the latter makes sense is a significant step towards establishing that the former does too.

Among other things, coding talk about genes is supposed to help us make good on the claim that genes are special developmental factors, that they count as being *privileged* causal elements in the developmental process. The way that coding talk is supposed to achieve this feat is nicely captured by Lorenz's (1965) image of the non-genetic causal factors in development as nothing more than the building blocks out of which organisms are systematically constructed according to a blueprint stored in the genes. On such a view, the real challenge for developmental biology is to understand how genetically specified instructions organise those available developmental materials into an organism. This way of looking at things really does make genes special.

There are many gene concepts in the literature, ranging from the essentially abstract, generically Mendelian notion of a gene as a trait difference marker to various

attempts to give molecular substance to the idea.¹ There are even some who argue that most of the assumptions that historically have underpinned the term ‘gene’ have been shown to be problematic, meaning that the very concept of a gene is now, in many ways, a misleading one that perhaps biology could do without (e.g., Dupre 2005). For the purposes of the present investigation I intend to put both definitional diversity and strategic critical eliminativism aside, and stipulate that we should be thinking of a gene as an entity with some sort of molecular unity, that is, as a stretch of DNA that possesses some sort of ontological integrity. To make this idea firm enough for the job at hand, we need to resist the tempting thought that the way to establish the molecular unity in question is by holding that genes are those parts of the genome that code.² Why this is should be clear enough: I have been assuming that there is conceptual space for the following result: there are genes but they don’t code for anything. If genes simply are the coding parts of the genome, then this result is not available. A negative answer to the question ‘do genes code?’ would imply that there are no genes. So we need to achieve the desired molecular unity without appealing to coding. But how? One answer would be to appeal to causally underpinned structural isomorphisms that exist between (a) sequences of DNA and (b) certain developmental elements that are causally downstream of those sequences. The most likely candidates for the latter are proteins, since the claim that systematic causal mappings exist between sequences of DNA and amino acids in proteins is not generally thought to be controversial (see the description of protein synthesis in section 4 below). Of course, if (i) all there is to coding is some sort of systematic causal dependence, and (ii) genes may rightly be said to code for proteins, whether or not they also code for traits, then the recognition of systematic causal mappings between (a) and (b) would herald significant progress in our investigation. Genes would code for proteins (at least). However (i) is implausible, as we shall see in section 3, so even if (ii) is true, it can’t be on the basis of (i). The upshot is that, at a programmatic level, we are in a position to identify genes in advance of settling the coding issue.

If the primary goal of introducing the concept of genetic coding is to single out genes as privileged causal elements in the developmental process, then it might well seem that any successful account of coding talk must have the consequence that, of the many causal factors that combine causally during development, it is the genes *alone* that end up coding for phenotypic traits. Let’s call this *the uniqueness constraint*. (Griffiths and Knight 1998 introduce what is essentially this very constraint in terms of what they call the ‘parity thesis’; see also Griffiths 2001.) The uniqueness constraint will not be met if either (a) the account of genetic coding under consideration fails to deliver the result that genes code for traits, since if genes don’t code for traits then they can’t do so uniquely, or (b) that account does deliver the result that genes code for traits, but its conditions for what it is to do this are met by other elements in the extended developmental system, since then genes won’t be the only developmental elements that code for traits.

It’s an irritating but undeniable fact that the natural world rarely plays ball with neat philosophical distinctions and categories, so the uniqueness constraint, in the strict

¹ For a recent review, see Griffiths and Stotz forthcoming. See also the *Representing Genes* project at <http://www.pitt.edu/~kstotz/genes/genes.html>.

² Contrary to some accounts, an organism’s genome is not simply its complete set of genes, but much more besides (see Dupre 2005).

form just stated, is very likely to be violated by any non-question-begging account of genetic coding on which we settle. Still, as long as such violations are not the norm, they are of no great matter. The background methodological thought concerning the genetic target of coding talk in biology can surely tolerate the odd non-genetic interloper. To be sensitive to this state of affairs we can modify the uniqueness constraint slightly, to require only the following: any successful account of genetic coding must have the consequence that those non-genetic elements for which it would be unreasonable, extravagant, or explanatorily inefficacious to claim that their contribution to development is representational in character do not count as coding for developmental outcomes. Call this the *weakened uniqueness constraint*. The weakened uniqueness constraint still has teeth, since the overwhelming majority of non-genetic developmental factors surely belong in the non-representational category. So perhaps it's acceptable for, say, an antero-posterior gradient of the bicoid protein in the *Drosophila* egg to be a vehicle of representational content (see Maynard Smith 2000b), but not, say, environmental temperature or the force of gravity (see below). There will no doubt be borderline cases to be fought over. Let's use the term *illegitimate non-genetic elements* to label those non-genetic factors for which it would be unreasonable, extravagant, or explanatorily inefficacious to claim that they code for developmental outcomes. So *legitimate non-genetic elements* are those non-genetic elements for which it would be reasonable, prudent, and explanatorily efficacious to claim that they code for developmental outcomes. We can now state an important principle: if one is considering the proposal that meeting certain specified conditions is sufficient for representing phenotypic structure, and it turns out that adopting those conditions would allow not only genes and legitimate non-genetic elements, but also illegitimate non-genetic elements to qualify (that is, there is a transgression of the weakened uniqueness constraint), then one should conclude that the proposed conditions are in fact *not* sufficient for representation.³

Using the benchmark of meeting the weakened uniqueness constraint as a sign of success, is it possible to give an adequate account of genetic coding? What follows is an attempt to answer this question. I should warn you that it won't exactly be a stroll in the park. Here's the route: Having set things up by saying more about exactly why the massively distributed character of the causal systems underlying development might actually be in tension with coding talk about genes (section 2), I shall consider the main contenders from the literature that purport to be not only plausible reconstructions of the character of such talk, but also justifications of its explanatory efficacy, and I shall find each of them wanting (section 3). At that point in the proceedings I shall lay out an

³ See Wheeler 2005, pp.208-209, for similar moves in the case of the neural target of the concept of representation in cognitive science. For what I take to be a similar weakening of (what I am calling) the uniqueness constraint in the case of genetic coding, see Stegmann 2005. To keep a sense of balance, it is worth noting that Sarkar (2000, 2005) explicitly recommends that a constraint which is closely analogous to the uniqueness constraint be dropped (at least for eukaryotes), on the grounds that no conceptually respectable concept of genetic information is available which doesn't have the consequence that that constraint is violated. I suggest, by contrast, that any notion of genetic information which has the consequence that the (weakened) uniqueness constraint is violated thereby loses its claim to conceptual respectability.

alternative and, I suggest, superior strategy for understanding and justifying coding talk in the relevant area of biology (section 4), but argue that that strategy has at least one quite radical implication that is, I think, a bullet that we just have to bite (section 5). In the final section (section 6), I shall consider an objection to the claim that there is coding for traits, an objection that applies to all the candidate strategies on the table, including the one I favour.

2. Cause for Concern

In recent years some of the most persistent critics of the idea that genes are informational entities that code for traits have come from the ranks of the developmental systems theorists. (For classic statements of the developmental systems position, sometimes just called developmentalism, see, e.g., Oyama 1985; Griffiths and Gray 1994; Griffiths and Knight 1998; and various papers in Oyama, Griffiths and Gray 2001.) Developmental systems theorists hold that the fundamental unit of evolution is the life cycle (a process that reconstructs itself from one generation to the next using a suite of developmental resources). Given that they take the life cycle to be the basic evolutionary unit, developmental systems theorists object to any view that understands development in terms of some basic dichotomy between genes and the rest of the extended developmental system. Thus they reject (what they see as) the massive over-emphasis on genes in (what they see as) mainstream neo-Darwinian evolutionary biology. It is important to be clear here that developmental systems theorists are not denying that there are any interesting empirical differences between the ways in which, say, DNA sequences and, say, parental scaffolding of language learning during early childhood contribute to development. What they deny is that these empirical differences should be turned into what Griffiths (2001, p.406) calls a “scientific metaphysics.” As Griffiths and Gray put the point:

[G]enes are just one resource that is available to the developmental process. There is a fundamental symmetry between the role of the genes and that of the maternal cytoplasm, or of childhood exposure to language. The full range of developmental resources represents a complex system that is replicated in development. There is much to be said about the different roles of different resources. But there is nothing that divides the resources into two fundamental kinds. The role of the genes is no more unique than the role of many other factors. (Griffiths and Gray 1994, 277-304)

One sure-fire route to the sort of scientific metaphysics that developmental systems theorists reject would be to adopt coding talk about genes alongside the uniqueness constraint (in either its full-strength or its weakened form), and to suggest that (all or the vast majority of) non-genetic developmental factors should, in a Lorenzian fashion, be relegated to mere genetically assembled building blocks. With this line of thought in their critical sights, Griffiths and Knight (1998) claim that “DNA does not contain a program for development” (p.253) and deny that there are “pre-formed blueprints or representations of traits in DNA” (p.255).

This is not the place to become over-focused on the details of the

developmentalist agenda. Our concern will be with a general way of motivating anti-representationalism about genes that is often at work in developmental systems thinking, as well as in the arguments of other prominent genetic coding sceptics who lay stress on the distributed character of the causal processes underlying development (for example, Maturana and Varela 1987, more on whom below). To bring things into focus, it will be useful to highlight a phenomenon that Andy Clark and I have dubbed *causal spread* (Wheeler and Clark 1999; see also Wheeler 2003, 2005). Causal spread obtains when some phenomenon of interest turns out to depend upon causal factors external to the system previously or intuitively thought responsible. Thus the identification of causal spread depends on the previously accepted explanation of the phenomenon of interest. Of course, given some default view of the world, even the most mundane examples of representational systems might display some degree of causal spread. For example, we might reasonably think of a C program as a set of instructions for (i.e., as a set of representations of) computational outcomes. The fact is that a C program is nigh on useless without certain ‘environmental’ (with respect to the program) features, such as a working operating system. However, nothing about the positive representational status of the C program would be threatened by the discovery of the essential causal contribution of the operating system.

Having said that, not all modes of causal spread are quite so obviously harmless to representational explanation. Consider what one might call *non-trivial causal spread*. This phenomenon arises when the newly discovered additional causal factors reveal themselves to be at the root of some distinctive target feature of the phenomenon of interest. In effect, where one confronts non-trivial causal spread, a new sharing-out of the explanatory weight is mandated. Call this *explanatory spread* (Wheeler and Clark 1999). Mameli (2005) explains the key points in this way.

Causal spread occurs when we discover some new factor causally involved in the occurrence of a phenomenon. Explanatory spread occurs when we realize that some factor that was not considered to be necessary in the explanation of a phenomenon is instead explanatorily necessary for that phenomenon. Or, to put it differently, explanatory spread occurs when we realize that some factor that was not taken to be part of a sufficient explanation of a phenomenon needs to be included in such explanation. Since the fact that something is causally required does not entail that it is also explanatorily required, causal spread does not necessarily lead to explanatory spread. But in cases where the newly discovered causal factor is deemed to be an important one, causal spread is likely to generate the inclusion of the newly discovered factor in any sufficient explanation of a phenomenon to which this factor causally contributes. That is, in these cases, causal spread leads to explanatory spread. (p.388)

In the present paper, the phenomenon of interest is organismic structure, and the default position is that such structure is down to genetic coding (on something like a Lorenzian model according to which the non-genetic material causes in development are the bricks and mortar out of which the organism is assembled according to the genetic blueprint). Against this background, one would have non-trivial causal spread where one discovered

a distributed developmental system in which non-genetic organismic and/or wider environmental factors made explanatorily non-negligible contributions to phenotypic form. So, is there non-trivial causal spread, and thus explanatory spread, in (our theories of) biological development? The answer, surely, is yes. Developmental explanatory spread is common. I shall give just a few brief illustrative examples, but the biological literature is simply brimming over with others.

First, consider the process of determination during cell specialization. In vertebrates, prior to the third cleavage stage, the cells in the developing embryo retain the possibility of achieving any of the full range of developmental outcomes available to the original zygote. The process of determination, in which the future course of development in the cells is differentially restricted, depends on a process in which the nuclei of the various cells become embedded in different cytoplasmic environments which in turn have different regulatory effects on the genes within the various nuclei. The sources of this differential embedding are a range of non-genetic factors, including pH balance and gravity, which result in a non-homogenous distribution of cytoplasmic materials within the egg. The inclusion of such non-genetic factors in our explanation is thus necessary if we are to account for the phenotypic phenomenon of interest.

Now consider the Mississippi alligator. These creatures lay their eggs in a nest of rotting vegetation which produces heat in varying quantities. Eggs that develop at lower temperatures (within some overall range) end up producing females, whilst those that develop at higher temperatures end up producing males. Eggs in a clutch will pass through the critical developmental window at various different temperatures, meaning that a mixture of females and males will be born. This environmental method of regulating sex ratio (a ratio which, for reasons of population-survival, needs to stay somewhere near 50:50 in the population) might seem a little hit and miss, but it works well enough (for more details, see, e.g., Goodwin 1994, p.38). Environmental temperature is a non-genetic factor, the inclusion of which in our explanation is necessary if we are to account for the phenotypic phenomenon of interest.

Finally, turning to human development, there is the much-studied phenomenon of scaffolding, in which a caregiver provides an on-line support system to enable a child to complete a task. As the child displays improving competence at the task, the caregiver gradually withdraws the support system, transferring responsibility for the completion of the task to the child (see, e.g., Wood, Bruner, & Ross 1976). Scaffolding is a key feature of child development, in areas such as discourse participation, literacy, and self-regulation, although the style and extent of the caregiver intervention varies among cultures. Scaffolding is a non-genetic factor, the inclusion of which in our explanations is necessary if we are to account for a range of phenotypic phenomena of interest.

Taking the foregoing examples as paradigmatic of development, we can conclude that explanatory spread is rife in that arena. What we can't conclude right now is that this generates a problem for coding talk about genes. For even though a new sharing out of the explanatory weight is mandated, such that non-genetic elements such as pH balance, gravity, environmental temperature and caregiver scaffolding become part of the relevant explanatory matrix, we haven't yet found out exactly why that fact might undermine the positive representational status of the genetic contribution.⁴ So let's turn now to an

⁴ To be clear: it is highly plausible that the kinds of non-genetic factors highlighted are

explicit argument against the view that genes code for phenotypic traits, one that appeals to (what I am calling) developmental explanatory spread.

We have often heard it said that genes contain the “information” that specifies a living being ... [but] when we say that DNA contains what is necessary to specify a living being, we divest these components ... of their interrelation with the rest of the network. It is the network of interactions in its entirety that constitutes and specifies the characteristics of a particular cell, and not one of its components. That modifications in the components called genes dramatically affect the structure is very certain. The error lies in confusing essential participation with unique responsibility. By the same token one could say that the political constitution of a country determines its history. This is obviously absurd. The political constitution is an essential component in any history but it does not contain the “information” that specifies that history. (Maturana and Varela 1987, p.69)

Maturana and Varela’s claim is that the fan of genetic information mistakenly confuses “essential participation with unique responsibility.” This suggests that for genes to count as carrying the information that specifies phenotypic traits, and thus for genes to be in the right conceptual ballpark to code for such traits, genes would need to bear *sole responsibility* for phenotypic form. But if, as the examples discussed earlier suggest, biological development is a playground for explanatory spread, then any such description of the genetic contribution here looks to be unwarranted. In general, DNA will *not* meet the sole responsibility condition. So it seems that if the representational theory of genes is tied to this condition, then that theory is straightforwardly undermined by the presence of developmental explanatory spread. And that, in essence, is Maturana and Varela’s point when they say, with respect to the cell, that it is “the network of interactions in its entirety that constitutes and specifies the characteristics of a particular cell, and not one of its components.”

But now surely something has gone wrong. Given my opening remark that every biologist understands (or ought to understand) development as involving a vast range of genetic *and non-genetic* causal factors, Maturana and Varela’s argument seems to do no more than set up a straw man for summary execution. However, things are not quite that simple. Indeed, despite the pretty much universal acknowledgement that there are extra-genetic causal contributions to development, the fact is that many theorists fall prey to the following, seductive thought: if one could find out the complete sequence of an organism’s DNA, then, in principle, one would be able to use that information *alone* to compute the adult organism, such that one would be able to predict, in every relevant detail, that adult’s phenotypic form. As DeLisi puts it:

illegitimate non-genetic factors, in the sense that *if* they counted as representations of developmental outcomes according to some account of what it is for an element to play that role, then that in itself would be grounds for rejecting the proposed account, since the weakened uniqueness constraint would have been violated. But we don’t as yet have such an account on the table. Our investigation hasn’t progressed that far.

The collection of chromosomes in the fertilized egg constitutes the complete set of instructions for development, determining the timing and details of the formation of the heart, the central nervous system, the immune system, and every other organ and tissue required for life. (DeLisi 1988, p.488)

At work here is a deceptively tempting view of outcome-directed representation that Clark and I have previously dubbed *strong instructionism* (Wheeler and Clark 1999; see also Wheeler 2003, 2005). Strong instructionism is the claim that what it means for some element to code for an outcome is for that element to fully specify the distinctive features of that outcome, where ‘full specification’ requires that the kind of exhaustive predictive power just indicated may, in principle, be achieved on the basis purely of what may be known about the putatively representational factor. In the present context, strong instructionism amounts to the claim that what it means for a gene (or a complex of genes) to code for a phenotypic trait is for that gene (or complex of genes) to fully specify the form of that trait. (Here we finally see the true colours of that compelling Lorenzian image of blueprints and materials.) However, given the presence of developmental explanatory spread, the fact is that knowing the entire sequence of an organism's DNA will *not* be sufficient to predict phenotypic form. So it seems that if coding talk about genes is tied to strong instructionism, then such talk is unsustainable.

Still, when it comes to providing a satisfactory account of genetic coding, there's *something* right about strong instructionism, namely that it respects the following, eminently plausible principle: in counting some target factor as a representation, in an appropriate outcome-directed sense, one buys into a crucial asymmetry between, on the one hand, that putatively representational factor and, on the other, the ecological backdrop against which that factor operates. Indeed, in all cases of algorithms, programs, instruction-sets, and other action-producing codes, those representational states and processes are able to perform their outcome-generating functions only given some assumed backdrop of other causally active states and processes. To build on a previous example: try running a C program without certain ‘environmental’ (with respect to the program) features, such as a working operating system. Moreover, where the right kind of asymmetry exists in the extended causal system, the discovery of causal spread, *even of the non-trivial variety that generates explanatory spread*, will not undermine representationalism. Thus we may conclude that it will be legitimate to treat genes as coding for traits, even in the face of developmental explanatory spread, just so long as we can legitimately regard the rest of the extended developmental system as the ecological backdrop against which genes make their representational contributions to phenotypic outcomes.

Notice that nothing about this suggestion requires that the crucial asymmetry be established independently of whatever detailed account we give of genetic coding. Rather, an adequate account of genetic coding should have the consequence that the right kind of asymmetry is manifest. We can now see how our overall benchmark for success, meeting the weakened uniqueness constraint, fits into the current dialectic. As I argued earlier, any satisfactory account of the concept of genetic coding must have the following consequences: (a) if any non-genetic factors count as coding for traits, then such violations of the uniqueness of genes in being representations of developmental outcomes should not be the norm; and (b) where such violations do occur, it should be neither

unreasonable, nor extravagant, nor explanatorily inefficacious to claim that the developmental contribution of the non-genetic factors in question is representational in character. In singling out genes as the predominant causal elements in the extended developmental system that code for traits, we simultaneously earn the right to treat the rest of that system as an ecological backdrop against which those genes (along with perhaps certain legitimate non-genetic elements) operate. Strong instructionism meets this demand through the full specification condition and the associated Lorenzian claim that non-genetic developmental factors in general are no more than biological bricks and mortar. But this view of non-genetic factors is not available once developmental explanatory spread is in the picture. So we are left with a challenge. What we need is an account of genetic coding that, without imposing the full specification condition, meets the weakened uniqueness constraint. In the next section I discuss a number of (ultimately unsuccessful) ways of addressing this challenge.

3. False Starts and Dead Ends

Here's a seductive first shot: genes code for traits because they *causally co-vary* with traits. In other words, appropriate causal co-variation is sufficient for genetic representation. One reason why this suggestion is provisionally attractive is that it makes contact with well-established views from elsewhere in science and philosophy that treat information in purely causal terms, or at least that might be used to explicate such an idea. Thus, at a first pass, causal information might, in part, be cashed out by way of mathematical information theory (Shannon and Weaver 1949), according to which (roughly) the quantity of information in a system is identified with the amount of order in that system. I say 'in part' because, strictly speaking, Shannon information supposes only correlation rather than causal correlation, so the causal nature of the correlation is an extra feature. I say 'at a first pass' because, for the purposes of genetic information, where we mostly want to talk about the *content* of the information in a system, rather than how much of it there is, the notion of causal information is more usefully explicated in the light of Dretske's (1981) influential philosophical treatment. Here is the resulting picture. Where there exists a sending system and a receiving system, connected by a channel such that the state of one system is causally related, in a systematic way, to the state of the other, then we have a signal – a flow of information – between the two systems. The causal information content of the signal is the source with which it is reliably correlated. This account is straightforwardly adapted such that entities carry information about causally downstream states with which they co-vary.

So how useful are causal information concepts in the present context? Mahner and Bunge (1997) question their applicability. First they point to the largely noiseless character of the (so-called) genetic code, noting that, practically speaking, the presence of noise is a standard issue when deploying Shannon information. Second, they claim that chemical processes cannot be thought of as signals that carry messages. In response, Maynard Smith (1999) argues (rightly in my view) that typesetting is largely noiseless, yet causal information concepts would surely be applicable there, and that it's hard to see why chemical processes couldn't be vehicles of causal information content, since all manner of other physical media, such as fluctuating currents in wires and sound waves, are standardly thought to be good for the job. A more serious barrier to the use of causal

information concepts in genetics is that, given the standard conception of the genome as specifying phenotypic outcomes in a disjunctive manner (i.e., develop like *this* under *these* environmental circumstances, like *this* under *these* environmental circumstances, and so on), the causal information view licences us to speak about genetic coding in ways that biologists don't. For example, to use an example due to Griffiths (2001), on the basis of a purely causal notion of information, the human genome would encode the instruction "when exposed to the drug thalidomide grow only rudimentary limbs." But biologists are unlikely to be tempted by such a claim. What this tells us is that the notion of causal information fails to capture the standard usage of informational terms in biology.

The most substantial problem confronting claims that appropriate causal co-variation is sufficient for genetic representation (or for genetic information), however, is one of *excessive liberality*. It is indeed a familiar point from the literature that genes are not the only factors in the developmental system that might be identified as causally co-varying with traits. Of course, it seems clear enough that if one could hold non-genetic causal factors in the developing body and the environment constant, while varying the genotype, then one would find causal co-variations between genes and phenotypic traits. However, if one could hold the genotype and the non-genetic causal factors in the developing body constant, while varying environmental factors, then one would find causal co-variations between environmental variables and phenotypic traits. Similarly, if one could hold the genotype and the environment constant, while varying non-genetic causal factors in the developing body, then one would find causal co-variations between those factors and phenotypic traits. But now if causal co-variation is a sufficient condition for a developmental factor to be representational, and if non-genetic causal factors in the developing body and the environment can causally co-vary with phenotypic traits, then those extra-genetic elements will sometimes count as coding for traits. This spells trouble because, given that many of the non-genetic factors here will be illegitimate ones, it falls foul of the weakened uniqueness constraint. In short, as a sufficient condition for coding, causal co-variation is excessively liberal, in that it licences explanations in which *too much* of some extended developmental system might emerge as coding for traits. So while it is eminently plausible that appropriate causal co-variation is necessary for genetic representation, it cannot be sufficient; genetic representation must be appropriate causal co-variation *plus something else*.⁵

What might that something else be? Here is a suggestion: genes code for traits because they (additionally) *set certain parameters* for the developmental systems that generate phenotypes. Perhaps then we can say that while genes do not fully specify the final phenotypic form (strong instructionism is false), they nevertheless code for developmental parameters and, by extension, for phenotypic traits (see Maynard Smith 1998 for a version of something like this strategy). The claim that genes might broadly be conceptualized as setting developmental parameters ought not, I think, to be particularly controversial. As Goodwin (1994, p.102) puts it, "[d]uring reproduction, each species

⁵ The fact that any systematic causal co-variation account of genetic coding will be excessively liberal (in the sense identified in the main text) is widely appreciated; see, e.g., Griffiths and Gray 1994, Maynard Smith 2000a, Griffiths 2001, Sarkar 2005.

produces gametes with genes defining parameters that specify what morphogenetic trajectory the zygote will follow.” However, it is quite another matter to claim that developmental parameter-setting is sufficient for representation, in the relevant sense. Indeed, this idea suffers from a version of the very excessive liberality problem that dogged the causal co-variation proposal. There seems little doubt that certain non-genetic factors (e.g., environmental temperature in the case of sex determination in the Mississippi alligator) might, like genes, be treated as parameterizing developmental systems. These non-genetic factors would then co-specify, along with the relevant genes, exactly which possible trajectory of that system would finally be traversed by the developing organism. But if performing the function of parameter-setting is sufficient for some developmental factor to count as coding for a phenotypic trait, then these extra-genetic factors will qualify. And, given that many of the non-genetic factors here will be illegitimate ones (environmental temperature would be an example), that violates the weakened uniqueness constraint. So even if performing the function of setting developmental parameters is necessary for a causal factor to play a representational role in development (which it might be, if one conceives of developmental systems as dynamical systems), such a role cannot be *just* a matter of developmental parameter-setting; it must be developmental parameter-setting *plus something else*.

It is time for a tactical rethink. So far we have considered, only to reject, two versions of the view that the status of genes as coding for traits is secured by properties of the direct causal contribution of genes. Perhaps the problem is that we’re looking at things all wrong. Perhaps representation is a matter of *function* rather than (brute) causation. In evolutionary biology, function-talk naturally invites an appeal to Darwinian selection. On this view, the function of a developmental element (if it has one) is (roughly) the positive contribution to organismic survival and reproduction prospects that ancestors of that element have made within historical populations. This generates the following proposal: genes code for traits insofar as they have been *selected* precisely so that a particular trait should occur (see, e.g., Sterelny 1995; Maynard Smith 2000a).

Why might someone think that appealing to selection is a good way to go on the issue of genetic coding? One motivating thought is that the concept of information that matters to biology is not causal information, but *intentional (or semantic) information*.⁶ The intentional concept of information is modelled on the kind of information carried by human thoughts and utterances. And one of the standard philosophical tests for the presence of intentional information is to see if one can make sense of the phenomenon of misrepresentation. In cases of input-related mental representations, misrepresentation occurs when the content of the representational state fails to correspond to the state of affairs in the world that caused it (e.g., one’s ‘cow’ representation is activated by perceptual contact with a horse). In cases of outcome-directed mental coding, misrepresentation occurs when the content of the representational state fails to

⁶ For this distinction drawn in these terms, see, e.g., Sterelny and Griffiths 1999; Maynard Smith 2000a; Griffiths 2001. For scepticism about the applicability, within the genetic context, of the intentional concept of information, see Sarkar 2005. Sarkar presents his own account of genetic information, in terms of what he calls ‘semiotic information,’ which, while being deflationary with respect to intentional information, is richer than Shannon information.

correspond to the state of affairs that it helps bring about (one's grasp-controlling representations are activated but, due to intervening causes, fail to result in the beer glass leaving the table). Genetic coding, if it exists, is, of course, an outcome-directed form of representation. So misrepresentation would occur if the content carried by the gene (its developmental instruction) fails to correspond to the phenotypic state of affairs that it helps bring about (the gene coding for long legs is causally active but, due to intervening causes in the developmental system, the phenotype ends up with short legs).

In outcome-directed mental representation, misrepresentation is made possible because the content of the mentally represented action-oriented instructions remains the same, no matter what happens in the rest of the action-generating system. So, in the case of genetic coding, we need it to be the case that the content of the represented instructions remains the same, no matter what happens in the rest of the developmental system. A dramatic illustration of the intuitive plausibility of cross-context content within biological systems comes from some striking experiments due to Halder, Callaerts, and Gehring (1995). There is a particular gene that plays a causal role in eye development in the mouse. Transfer that gene to the fruitfly *Drosophila* and it will result in the development of an eye – a compound eye, a fruitfly eye. Indeed, activate the transplanted gene at various sites and one will get a fruitfly eye developing at the different organismic locations in question (e.g., at the usual site of a leg). So, if this gene codes an instruction, the content of that instruction is very plausibly something like 'build me an eye here'. That's the developmental instruction represented by that gene.⁷ Intuitive plausibility aside, the key point here is that we can make sense of intentional representation because we can make sense of the coding element in question having an 'intended' effect (which in turn determines the content of the represented instructions), even if that effect doesn't come about. Where information is interpreted merely in terms of systematic causal covariation, there is no room for this distinction between intended and unintended effects, hence the fact that causal information concepts fall prey to the thalidomide counterexample discussed earlier. As Griffiths (2001) notes, the notion of intentional information can handle this case, since growing only rudimentary limbs is not one of the intended effects of the genes concerned. But while the idea of the intended effect of a representation might seem straightforward enough in the case of human utterances, exactly how are we to secure that idea in the case of genes? It's here that the appeal to selection comes in. Intended effects are identified by reference to the developmental contribution for which the gene/genes in question was/were selected. So, there is some justification for the claim that an appeal to selection may secure the appropriate sort of informational content for genetic representations. (Whether or not it is the only way to secure such content is another issue – see below.)

Another key thought in the literature is that the appeal to selection will not result in violations of (what I am calling) the weakened uniqueness constraint. Thus Sterelny (1995) observes that the growth patterns of snow gums will differ depending on whether they are exposed to snow or wind. Both genotype and environment are necessary causal

⁷ My interpretation of this scenario follows that given and defended by Maynard Smith (2000a,b). For an alternative interpretation, according to which the gene in question should be seen as a reader of information carried by other genes, rather than as carrying information in its own right, see Sterelny (2000).

factors in determining the plant's final phenotypic form. But whereas the climatic conditions are, in a sense, 'just there', the genotype exists purely because of its role in producing the phenotype, and thus has the evolutionary function of producing the phenotype. And that, according to Sterelny, is why the genotype codes for the phenotype, whereas the environmental factors do not. Notice that, on this view, two genes could play the same brute causal role (say in the production of an eye), but one would rightly be said to code for the relevant property of the eye, while the other wouldn't, if the former had been selected for that job while the latter hadn't.

The suggestion on the table, then, is that genes code for traits insofar as genes, *unlike the rest of the developmental system*, have been selected precisely so that a particular trait should occur. The first thing to say is that selection is not necessary for (genetic) representation (Sarkar 2000, Wheeler 2003). To see why, consider the following argument (Wheeler 2003). Genes are sometimes linked physically, in such a way that the evolutionary fate of one gene is bound up with the evolutionary fate of another. This provides the basis for a phenomenon known as genetic hitchhiking. To see how genetic hitchhiking works, let's provisionally allow ourselves the language of 'genes for traits,' and construct a simple evolutionary scenario. Assume that, in some creature, the gene for a thick coat is linked to the gene for blue eyes. Let's also assume that this creature lives in an environment in which it is selectively advantageous to have a thick coat, and selectively neutral to have blue eyes. What will happen is that the gene for a thick coat will be selected for. But since the gene for blue eyes is linked physically to the gene for a thick coat, the gene for blue eyes will be inherited too, even though it bestows no selective advantage, has not been selected for, and thus has no evolutionary function. For present purposes, the key feature of genetic hitchhiking is this: the fact that the hitchhiking gene is not selected for does not in any way threaten, by making theoretically awkward, our description of it as coding for blue eyes. So the phenomenon of genetic hitchhiking tells us that selection is not necessary for representation.

There are two obvious responses that the selectionist about genetic representation might make. First, she might complain that even if the hitchhiking gene has not been directly selected for, it has a kind of honorary 'selected for' status, on account of the fact that it is linked to a gene that has been directly selected for. But this seems to be the wrong way to describe the situation. After all, in the foregoing example, the gene for blue eyes has certainly not survived because of its role in producing the phenotype. Thus it is hard to see how the notion of being selected for can get any sort of grip. A second response might be to concede that the blue-eyes-related gene does code for blue eyes, but to maintain (i) that selection is sufficient for, but not necessary for, representation, and (ii) that while selection explains why we should describe the thick-coat-related gene as coding for thick coats, some other explanation will be required in the case of the blue-eyes-related gene. But unless there are some powerful independent considerations in favour of clinging on the selectionist strategy (considerations that would have to be produced and judged), there is surely no reason to multiply explanatory stories in this way. What we really want, it seems, is a single account of genetic coding that covers both cases.

Anyway, the fact is that if we adopt the view that selection is sufficient for genetic representation, then, contra Sterelny's snow-gum-driven conclusion, we will fall foul of the weakened uniqueness constraint. To demonstrate this, we can call on a thought

experiment due to Mameli (2004). Consider a species of butterfly with the following properties: (a) all members of the species are genetically identical, and no genetic variation can be produced; (b) the butterflies eat a particular species of plant during the early stages of their life; (c) females lay their eggs on plants of the same species as the one on which they hatch; (d) they do this by eating the leaves of the plant on which they hatch, by imprinting on the taste of the leaves, and by laying their eggs on plants with the same taste. Now, as a result of a developmental accident, the imprinting mechanism in one female malfunctions. She lays her eggs on the ‘wrong’ plant which, as it happens, is a new species of plant in this species of butterfly’s environment. By chance, this new plant makes these butterflies bigger. Now assume that, in this species, bigger size confers a fitness advantage. Because of this, the lucky butterfly’s offspring grow up fitter than other butterflies of the species. The offsprings’ imprinting mechanisms work just fine. So they lay their eggs on the new species of plant. Given competition for resources, the lucky butterfly’s descendants will out-compete their conspecifics and, eventually, all the butterflies of this species will hatch on the new plant. This is a process of natural selection – there is heritable variation in size caused by variation in plant of hatching – but there is no genetic variation. Mameli introduces the term *envirotypes* to describe factors such as plant of hatching in the lucky butterfly scenario, factors that are intergenerationally stable (and which thus underwrite selection by guaranteeing a correlation between parental variants and offspring variants), but which are environmental rather than genetic in character. Given the possibility of envirotypes, “not all selection is at bottom genetic selection. Some selection is *nongenetic (or envirotypic)* selection” (Mameli 2004, p.41).

For present purposes, the principal message of the lucky butterfly is that if being selected for is sufficient for some developmental factor to qualify as coding for a phenotypic trait, then non-genetic factors will sometimes attain coding status, since non-genetic factors may sometimes be selected for. And, of course, if those non-genetic factors are illegitimate ones, as is plausibly the case with plant of hatching in Mameli’s thought experiment, then that violates the uniqueness constraint. In short, we confront yet another version of the excessive liberality problem.

It is worth pausing here to note two things. First, the potential existence of Mamelian envirotypes blocks the thought that it must in principle always be possible to trace the adapted character of non-genetic developmental resources back to prior genetic selection (that is, given the suggestion currently on the table, to genes that code for those resources). This is especially clear when the notion of an envirotypes is established in the case where genetic variation is ruled out. Second, the conceptual linking of selection to representation, plus the claim that direct selection for non-genetic developmental units is held to be possible, are points embraced by Sterelny and Kitcher in their *extended replicator* proposal (Sterelny and Kitcher 1988). According to the idea of extended replicators, *all* adapted developmental resources code for traits. Now, if one interprets the extended replicator proposal as an attempt to reconstruct a theory of *genetic* coding (which is how it seems to be presented by, e.g., Sterelny and Griffiths 1999, p.87, where it is described as providing a “formal reconstruction of the “gene for” locution”), then one can only assume that Sterelny and Kitcher (a) are unmoved by considerations of the uniqueness of genes with respect to coding status, and (b) do not believe that there is any independent way (independent, that is, of the criterion of selection) to determine whether

or not an environmental contribution to development might legitimately qualify as a coding element. My view, as should be clear, is that both (a) and (b) are errors. Of course, it would be entirely consistent to endorse the weakened uniqueness constraint, agree that there are extended replicators, but deny that being selected for is a sufficient condition for coding.

Where next? One intuition that we haven't yet explored is that coding talk is conceptually intertwined with the notion of inheritance. Thus, one might claim that genes code for traits insofar as they are what is passed on from one generation to the next in evolution. Of course, genes *are* inherited. But, using a toy example, let's assume that eye colour can be traced to a single gene, and further that, in a particular offspring, the gene inherited at conception would, if expressed, produce brown eyes. Let's also say that psychology has shown blue eyes to be advantageous to getting on in life, by attracting the favourable attentions of others. This looks like bad news for our target offspring. However, a gene transplant is carried out, such that the inherited brown-eyes-related gene is removed, and a blue-eyes-related replacement inserted by doctors. If we deploy the same style of reasoning as we used in the hitchhiking example above, and provisionally allow ourselves the language of genes as coding for traits, we would naturally say that the inherited, but now removed, brown-eyes-related gene coded for brown eyes. But what about the non-inherited but functional, deliberately inserted, blue-eyes-related gene? As far as I can see, the fact that this gene has not been inherited does not seem to threaten, or make in any way theoretically awkward, the language of coding. This suggests that being inherited cannot be a necessary condition for coding-talk to get a grip within development.

Moreover, and perhaps more significantly, if we define inheritance without an antecedent pro-gene prejudice, as the biological like-begets-like phenomenon, and so as to fix on elements that are robustly and reliably replicated in each generation of a lineage, and that persist long enough to be the target of cumulative selection, then the fact seems to be that genes are not *all* that organisms inherit. For example, there are so-called *epigenetic inheritance systems*, such as the inheritance of methylation patterns via a separate (from the genetic, that is) copying system; and there is inheritance through *host imprinting*, as when Mameli's imaginary butterflies inherit increased size through imprinting on the taste of a new plant (see above); and then there is the phenomenon of inheritance via *niche construction*, as when beaver offspring inherit both the dam that was communally constructed by the previous generation and the altered river flow that that physical structure has produced. Moreover, as Mameli (2005) has argued, simply mentioning DNA-copying and DNA-transmission cannot be sufficient to explain the reliable trans-generational reoccurrence of some phenotypic trait, *if*, that is, one is compelled to mention more than DNA in one's explanation of the development of that trait. Thus:

If we want to explain why the shape and structure of the legs of human offspring reliably have the same shape and structure as the legs of human parents, we have to mention not only the reliable reoccurrence of the genes involved in normal human leg development, but also the fact that humans experience roughly the same amount of gravitational force from one generation to the next. And this means that, when we explain the reliable

reoccurrence... of legs with a certain structure and shape in human lineages, we have to mention not only DNA-copying and DNA-transmission, but also those processes that explain why human beings experience the same amount of gravitational force generation after generation. (Mameli 2005, p.389)

In short, Mameli's argument is that since there is explanatory spread in (our theory of) development, there is explanatory spread in (our theory of) inheritance.

The upshot of the foregoing observations is that if being inherited is sufficient for some developmental factor to qualify as coding for a phenotypic trait, then non-genetic factors will sometimes count. And if those non-genetic factors are illegitimate ones, as is plausibly the case with the processes that explain why human beings experience the same amount of gravitational force generation after generation, then that once again violates our old friend the weakened uniqueness constraint.

4. A Better Idea

Things are not working out, so let's switch tactics again, and focus our attention on the phenomenon of *protein synthesis*. The guiding intuition here is that something (or some things) about the contribution made by genes to this process will single them out as coding elements, in a way that doesn't contravene the weakened uniqueness constraint.⁸

We should begin by reminding ourselves of some familiar biological facts.⁹ In the first stage of protein synthesis, the organism's DNA acts as a template in the manufacture of molecules of *messenger RNA (mRNA)*. In prokaryotic gene expression, the initial RNA molecule generated by *transcription* (the process underlying templating) is equivalent to the mRNA. However, eukaryotic genes contain sequences of base pairs that are functionally redundant with respect to protein synthesis, sequences known as *introns*. In the initial transcriptional phase, all the DNA (redundant and salient) is transcribed into a complementary RNA copy called *nuclear RNA (nRNA)*. Then, in a post-transcriptional phase of so-called *RNA splicing*, the introns are subtracted so that only the functionally salient sequences, the *exons*, remain.

⁸ The thought that the concept of genetic coding will finally be vindicated by facts about the mechanisms of protein synthesis is shared by Wheeler and Clark 1999; Godfrey-Smith 2000b; Maynard Smith 2000a; Sterelny 2000; Sarkar 2000, 2005; Wheeler 2003; and Stegmann 2005. These alternative developments of the same basic idea contain some significant variations in the precise factors identified as the features of interest, and are occasionally accompanied by certain concessions regarding (a) the full-strength uniqueness constraint and (b) what exactly is represented. Here I shall not attempt to map out *all* the different features that characterize these different views, although it is worth noting at the outset that the concept of arbitrariness (understood one way or another) plays a central role in all of them. The nuances that matter will be discussed as I work towards and defend my own current view.

⁹ Protein synthesis is of course a complicated business, and I have no doubt that some readers will be unhappy with one or other aspects of the brief description that I shall give. Nevertheless, the simplified picture I shall paint is broadly correct and good enough for present purposes.

The second stage of protein synthesis is known as *translation*. This process is very similar in prokaryotes and eukaryotes, although in prokaryotes transcription and translation are closely coupled, with the latter beginning before the former is complete. In translation, the mRNA molecules produced by transcription (plus RNA splicing in the case of eukaryotes) determine the manufacture of different *proteins*, which are the building blocks of bodies. Molecules of mRNA are divided into triplets of nucleotide molecules known as *codons*, and (ignoring certain singular cases) every instance of a particular mRNA codon, as generated from its DNA template, is believed to result in an instance of the same amino acid being added to an emerging protein. However, this is, as Sarkar (2000, p.210) puts it, a “frozen accident.” In other words, there is nothing in current biological knowledge to suggest a convincing physical-chemical reason why the mappings could not have been set up differently.¹⁰ So what exactly goes on in translation? The cell’s cytoplasm contains protein-manufacturing-sites called *ribosomes*, along with molecules of another sort of RNA called *transfer RNA (tRNA)*. Molecules of tRNA are single nucleotide triplets attached to single amino acids. What happens during translation is that an mRNA molecule becomes attached to a ribosome, and then passes through it, one codon at a time. When a new codon moves into place, the ribosome (through trial and error) locates a molecule of tRNA that, according to the so-called base-pairing rules, features a particular nucleotide triplet. The ribosome then strips off the amino acid from the other end of the tRNA molecule, and adds it to the protein which is under construction. Stripped of its amino acid, the tRNA molecule floats off into the cytoplasm, to be ‘recharged’ with ‘the right’ amino acid.

Out of all this biological detail, two conceptually interlocking features of the architecture of protein synthesis strike me as representationally significant.

1. *Arbitrariness*: In the specific sense in which I am using the term, arbitrariness indicates that the equivalence class of different systemic elements (say nucleotide triplets) that could perform some systemic function (say, given other causal factors, produce a specific amino acid) is fixed not by any non-informational physical properties of those elements (say their shape or weight), but rather by their capacity, when organized and exploited in the right way, to carry specific items or bodies of information. The mappings from particular nucleotide triplets to particular amino acids are arbitrary, in this sense.
2. *Homuncularity*: The ‘right way’ of exploiting the systemic elements just highlighted is established where the system in question is *homuncular*. As I shall use the term, a system is homuncular just when it can be usefully compartmentalized into a set of communicating subsystems, each of which performs a well-defined subtask that contributes towards the collective achievement

¹⁰ The standard way of describing this frozen accident is to say that the genetic code is arbitrary. As will become clear, however, it is at least plausible that arbitrariness, understood a certain way, is a necessary condition on there being a code *at all*. If that is right, then if the mapping in question were not arbitrary, there would be no pressure to think of the system in question as one of encodings. So the right question is not “Is the genetic code arbitrary?,” but rather, “Is there a genetic code?.”

of a systemic outcome. In an homuncular analysis, the communicating subsystems are conceptualized as trafficking in the information that the inner vehicles carry. More specifically, certain subsystems are interpreted as *producing* information that is then *consumed* downstream by other subsystems. Of course, homuncular subsystems must not be thought of as being, in any *literal* sense, understanders of the information in question. (They are not really little people.) Nevertheless, the fact is that the ways in which the functionally integrated clusters of subsystems exploit inner elements, so as to collectively generate systemic outcomes, become intelligible only if we treat the subsystems involved as dealing in the information that those elements (organized and exploited as they are) carry, rather than as responding only to non-informational physical properties of those elements. The mechanisms underlying protein synthesis are most illuminatingly conceived of as being homuncular in the requisite sense. Thus mRNA molecules are assembled by a producer subsystem that *encodes* informational content in those molecules. And the translation-realizing machinery of ribosomes and tRNA constitutes a consumer subsystem that *decodes* (and thereby exploits) that same informational content.

What we have in protein synthesis, then, is a producer-consumer economy of outcome-related, information-based transactions between homuncular subsystems. Such an arrangement surely warrants a representational interpretation, according to which the elements in which the homuncular subsystems deal are legitimately identified as coding for the outcomes in question.¹¹

If we add these observations, about the architectural conditions under which a representational interpretation of some system is mandated, to our previous thinking about the purely causal conditions for representation, then the following general principle suggests itself: the presence of (i) systematic causal co-variation between the putative vehicles of content and specific causally downstream structures, (ii) arbitrariness, and (iii) systemic homuncularity is sufficient for coding-talk. (If one conceives of

¹¹ In Wheeler 2005 (chapters 8 and 10) I argue that the interlocking architectural features of arbitrariness and homuncularity also form the basis of an adequate account of the notion of representation as used in cognitive science. In philosophy of mind and cognitive science, the connection between arbitrariness and representation has been made previously by, for example, Pylyshyn (1986), and the notion of homuncularity (or something very close to it) has been linked with representation before, by, for example, Millikan (1995). The conceptual interlock between arbitrariness and homuncularity is not part of these theorists' treatments, although it is anticipated by Wheeler and Clark's (1999) link between arbitrariness and information-based consumption. The sense of homuncularity that I have pressed into service in this paper is superficially 'thinner' than its cognitive-scientific cousin (at least as I develop the latter), since the present notion does not *explicitly* require that the subsystems concerned be organized in an hierarchical manner. In fact however, in any homuncular analysis there will always be a background commitment to the idea that subsystems that perform relatively complex subtasks could, in principle, be analyzed into further subsystems that perform relatively simpler subtasks, until the whole edifice 'bottoms out' in subsystems that perform primitive bio-chemical functions. Thus there is always a (perhaps weak) sense of hierarchicality in play.

developmental systems in dynamical systems terms, then one might replace the causal co-variation condition with one that explicitly mentions developmental parameter-setting. If so, then the causal co-variation condition will be implicit, since elements that are rightly conceived as setting developmental parameters will always causally co-vary in a systematic way with the outcome states of interest.)

As it happens, my view is that conditions (i)-(iii) are not only jointly sufficient for representation, but necessary too. It seems undeniable that systematic causal co-variation is necessary for representation. The additional necessity of arbitrariness is, perhaps, clear enough. Thus, to give an intuitive non-genetic example, where the outcome in question is, say, keying my actions to the door-stopping potential of some book on my office shelf, the equivalence class of neural states which may perform the right outcome-achieving role of selecting a suitable book will be fixed precisely by the fact that some of those elements are able, when organized and exploited in the right way, to carry some relevant item or body of information (e.g., that the book is heavy enough to hold the door open). Here it seems safe to say that the elements in question represent the associated worldly features. But now consider the outcome of simply holding my office door open. The equivalence class of suitable objects which may perform this role will be fixed by (roughly) the non-informational properties of being heavy enough and being sufficiently non-obstructive with respect to passing through the doorway. Here, where the equivalence class of different elements that could perform the function at issue is fixed by certain non-informational physical properties of those elements, there is simply no place for the language of representation. This suggests that arbitrariness is necessary for representation. And if, as my architecture-related reflections suggest, arbitrariness and homuncularity arrive on the explanatory scene arm in arm (conceptually speaking), then the claim that homuncularity is necessary for representation looks to be concurrently established.¹² So, if I am right, the joint presence of (i) systematic causal co-variation between the putative vehicles of content and specific causally downstream structures, (ii) arbitrariness, and (iii) systemic homuncularity is necessary and sufficient for coding-talk.

As my description of the machinery underlying protein synthesis indicates, that machinery satisfies conditions (i)-(iii). (Although I have not argued explicitly that there are appropriate causal co-variations in protein synthesis, it should be clear enough that there are systematic causal mappings between, on the one hand, both DNA and mRNA, and, on the other, proteins.) But how secure is the general account of representation that I have given? Here I shall consider four objections.

First, one might object to the claim that arbitrariness is necessary for representation, on the grounds that not all elements that we take to be representations have that property. This is the sort of complaint that needs to be settled on a case-by-case basis, but let's at least consider one of the more plausible candidates for positive representational status coupled with non-arbitrariness, namely onomatopoeiaic words.¹³ Since the pronunciation of such words suggests their meaning (e.g., meow), it might seem that they cannot be arbitrary. Yet we still think of them as representational, so they

¹² In Wheeler 1995 (chapter 10) I give independent reasons for thinking that homuncularity is necessary for representation.

¹³ This worry was put to me by Elliott Sober (in discussion). Thanks to Phyllis McKay and Peter Sullivan for helping me to think about the best way to repel it.

provide a counter-example to my suggestion that arbitrariness is necessary for representation. However, it seems to me that the intuition that onomatopoeic words cannot be arbitrary trades on a thought that is not reflected in the concept of arbitrariness, as I have unpacked that concept here. Many different physical sound patterns could realize the word ‘meow’ (compare the way in which native French and native English speakers pronounce the word), and what fixes the equivalence class of appropriate sounds is the informational content that they carry (roughly, this is a sound that cats make). Of course, the class of sounds that may be legitimate physical realizers of the word ‘meow’ is presumably not infinite, but then infinite realizability is not required for arbitrariness, in the sense that I am using that term. The class of legitimate physical realizers of the so-called ‘genetic code’ is certainly not infinite.

Second, one might object to the claim that conditions (i)-(iii) are sufficient for coding, on the grounds that what is additionally necessary for (any sort of) representation is the presence of combinatorial structure – perhaps of a mild kind – in the inner elements, enabling structurally related elements to guide different-but-related outcomes. It is worth pointing out that the system underlying protein synthesis would plausibly satisfy this condition (for related thoughts, see Godfrey Smith 2000b). However, as far I can see, and despite arguments to the contrary by, for example, Haugeland (1991), such systematicity concerns the *power of* a representational system, rather than its *status as* a representational system.

Third, one might complain that in moving beyond an austere causation-based story about coding, to one that is based on architectural features, I have introduced an ineliminable reference to function, and thus ultimately to natural selection. If so, then there would at least be a suspicion that I am open to the very criticisms of selection-based approaches that I myself have advanced. However although, in evolutionary biology, function-talk naturally invites an appeal to Darwinian selection, generating what we might call *Darwinian functions*, that is not the only way to think about functions in biological systems. *Causal role functions* (Cummins 1975), as studied by, for example, anatomists and physiologists, are identified not by evolutionary history, but by analyzing an overall task (thinking, swimming, digesting food, assembling proteins) into well-defined subtasks performed by well-defined parts or subsystems. Griffiths illustrates the distinction with an example germane to our project here. A “sequence of nucleotides GAU has the [Darwinian function] of coding for aspartic acid if that sequence evolved by natural selection because it had the effect of inserting that amino acid into some polypeptide in ancestral organisms” (Griffiths 2005, p.1). The same nucleotide sequence “has the [causal role function] of coding for aspartic acid if that sequence has the effect of inserting that amino acid into some polypeptide in the organism in which it occurs” (Griffiths 2005, p.2). Homuncular analysis naturally buys into the causal role sense of function, but it remains a further issue whether or not the causal role function of an homuncular subsystem is accompanied by a function in the selective sense. But notice, in this context, that the notion of causal role function (which is conceptually richer than mere causal information) supports talk of misrepresentation, and thus plausibly of intentional information. Without additionally appealing to selection, we can surely make sense of a scenario in which intervening causes prevent the subsystemic outcome that is related to a particular causal role function from coming about.

Finally, one might worry that conditions (i)-(iii) suffer from their own excessive

liberality problem, in that they will be met by inappropriate environmental factors. To see why this is plausibly not the case, we can build on an example due to Godfrey-Smith (2000a). Take a plant that responds to an increase in day length by starting to flower. According to Godfrey-Smith, the connection between the cause (the increase in day length) and the effect (flowering) here is arbitrary, because the cause could have been interpreted in many other ways by the flower. Thus the cause in this arrangement counts as arbitrary, it's environmentally located, and it looks like the kind of factor that really shouldn't count as a representation of a developmental outcome; so, with respect to the arbitrariness condition alone, excessive liberality emerges as a genuine danger. Of course, I have characterized arbitrariness not in terms of a cause potentially having a range of different effects, but in terms of the equivalence class of different physical factors that could have played the same causal role being fixed by informational rather than brutally physical considerations. Nevertheless, the content 'start to flower' could clearly have been carried by environmental factors other than increase in day length, so it looks as if arbitrariness in my sense is present too, and in the same worrying place. The solution (in the framework I am promoting) is to take seriously the conceptual interlock between arbitrariness and homuncularity with respect to the justification of coding-talk. For while, in the flower case, it might well be said that there is a consumer system that digests the putative information (by interpreting the increase in day length as an instruction to flower), it is hard to see how to make sense of the claim that the overall arrangement contains a producer system that has performed the role of *encoding that information in* the relevant causal factor, namely in the increase in day length. So the environmental factor in question does not emerge as being representational in character.

This response to Godfrey-Smith's example does not establish that non-genetic factors could not ever qualify as vehicles of representational content in development, once arbitrariness and homuncularity are plugged in as necessary conditions. Take animal signalling systems. If one could specify the appropriate causal co-variations (that is, between the signals and the construction of developmentally downstream structures), those systems will contain noises, marks, and so on, that will count as environmentally located vehicles of representational content. (Of course, the producer subsystem will be in one individual animal, while the consumer subsystem will be in another, but nothing I've said rules out such a state of affairs.) However, notice that the existence of such elements does not violate the weakened uniqueness constraint. It is neither unreasonable, nor extravagant, nor explanatorily inefficacious to claim that the developmental contribution of such factors is representational in character. What needs to be ruled out is the systematic inclusion of illegitimate factors (such as an increase in day length). And that, I think, is plausibly achieved by a proper recognition of the part played by the producer subsystem. However, that recognition also brings us to what, I suspect, is the most controversial claim that I shall make in this paper

5. A Bullet to Bite

Strictly speaking, according to the proposal currently on the table, it's not the DNA molecules that constitute the representational vehicles that play a coding role in development, but rather the nucleotide triplets (the *codons!*) that make up the mRNA

molecules.¹⁴ Genes don't code; mRNA does. Why is this the right unpacking of the proposal? On the view developed here, representation requires a producer and a consumer. The representations are the vehicles of content that support the communicative transactions between these systems. The producer encodes information into the vehicles in question, the consumer decodes information from them. In the case of protein synthesis, the consumer system is the distributed mechanism of ribosomes and tRNA that realises the process of translation in which mRNA determines the manufacture of proteins. So what is the producer system? The most compelling answer, it seems to me, is that it is the distributed mechanism underlying the process in which the organism's DNA acts as a template in the manufacture of *mRNA* molecules, that is, the producer system is the machinery of transcription and, in the case of eukaryotes, RNA splicing (as described above). It's that very machinery that encodes the information in mRNA molecules, the information that will later be decoded during translation. To see why this interpretation is the most compelling, we need to consider some objections.

The first is to claim that while there is a strict sense in which it's mRNA nucleotide triplets that code, the fact is that DNA codes *by extension*. Godfrey-Smith (2000b, p.32) puts it like this: "The "genetic code" is, strictly speaking, the rule linking RNA base triplets with amino acids. This "interpretation" of the RNA determines the "interpretation" of the DNA from which the mRNA was derived." This suggestion faces a serious difficulty. To see why, we need to consider an analogy with the cognitive science of visually guided action. In the broadest terms, according to much thinking in cognitive science, patterns of stimuli on the retina determine the structure of certain inner states that intervene between sensing and action. Strictly speaking, what determines the final outcome (the agent's behaviour) will be some action-specifying inner state that needs ultimately to be translated into physical movements. Now if, as seems warranted, we map (a) the pattern of stimuli on the retina onto DNA sequences, (b) the process by which those stimuli determine the structure of the outcome-specifying inner state onto transcription plus RNA splicing, (c) the outcome-specifying inner states onto mRNA molecules, (d) the process by which those states are turned into physical behaviour onto translation, and (e) the behaviour onto proteins, then by something like the reasoning that Godfrey-Smith advocates in the case of protein synthesis, it would be right to say, in cognitive psychology, that patterns of stimuli on the retina code for particular actions. And that doesn't seem right. There will, of course, be systematic correlations between both (i) the form of the action and the retinal patterns, and (ii) the content of the inner action-specifying state and the retinal patterns, but the fact is that as we travel causally downstream from the retinal input, extra content is introduced that is relevant to the exact form of the actions produced. Crucially this content is introduced during the construction of the output-specifying inner states (reflecting, e.g., the goals and interests that the agent is pursuing, and that determine how the agent should respond to the input). Interestingly, in theories of visually-guided action where environmental stimuli are said to specify

¹⁴ Essentially the same claim is made, on related but importantly different grounds, by Bullock (1998). Bullock treats genes as themselves encoders, a position which I reject (see later in this section), and he makes a pivotal appeal to natural selection in his argument that the machinery of protein synthesis contains a consumer system, an appeal which I think is unsustainable (see arguments in section 3 above).

actions more directly (e.g., in Gibsonian ecological psychology), those theories are often characterised as being non-representational in character. The consequence of these observations is that one wouldn't have a mandate to say that how we interpret the inner action-specifying states here determines how we should interpret the retinal input. If the analogy holds, then similarly we should not endorse Godfrey-Smith's suggestion that the interpretation of the mRNA determines the interpretation of the DNA from which the mRNA was derived.

The obvious counter-move here is to question the analogy by claiming that nothing approaching the complexity present in the psychological case is present in the process by which DNA sequences are transformed into mRNA base triplets. Thus Maynard Smith (2000a) draws his own analogy, this time with Morse code. In the use of Morse code the content of the message is, Maynard Smith claims, first encoded into phonemes by the original coder (a human being), and then merely *converted into* Morse code. He then argues that, in the case of DNA, the original coder is natural selection, which encodes developmental information into genes. That information is then merely converted into mRNA base triplets.

The first thing to say here is that we have found good reasons to conclude that, in the present context, selection is not necessary for representation (see above), so the appeal to natural selection needs to be treated with suspicion. However, the claim about 'mere conversion' could in principle be freed from the link with natural selection. One might try to argue, for example, that the way in which the interpretation of the mRNA determines the interpretation of the DNA obviates need for a producer system altogether. What really needs to be resisted, then, is the claim that the DNA-to-mRNA transition can be relegated to anything approaching mere conversion on the phonemes-to-Morse model. The second point to make is that, in *the case of eukaryotes at least*, there are events that occur between transcription and the beginning of translation that undermine any such relegation. I have already mentioned RNA splicing. Sometimes this takes the form of so-called *alternative splicing* in which the same initial RNA transcript gets spliced in different ways to generate several proteins. In addition, there are other complex processes of RNA editing, involving the addition, removal, or replacement of bases. So, *in the case of eukaryotes at least*, the analogy with the mechanisms by which sensory stimulation results in inner action-coding seems to hold, which means that one cannot deploy Godfrey-Smith's strategy to establish that eukaryotic DNA codes in protein synthesis. And, having blocked the use of that strategy in the case of eukaryotes, it seems to me that we have good methodological reasons to extend our preferred interpretation – that mRNA not DNA codes in protein synthesis – to prokaryotes too. As we have seen, Godfrey-Smith himself concedes that *strictly speaking* the so-called genetic code is the mapping between mRNA base triplets and amino acids, suggesting strongly that, *strictly speaking*, it's mRNA that codes for proteins. There seems little reason to speak loosely for prokaryotes if such talk has shown to be misleading in the case of eukaryotes.

What looks like an alternative way to resist my argument concerning the location of the coding entities in protein synthesis may be found in an argument due to Stegmann (2005). Stegmann identifies a notion that he calls *instructional content*, unpacked as the information for the synthesis of some outcome, such that that outcome is determined via the step-by-step realization of operations specified in advance. The thought is that this kind of content is familiar from everyday representational entities such as cooking recipes

and computer programs. Given this notion of content, Stegmann argues that if we look at the role of DNA in transcription, then we find that it carries instructional content, in virtue of the template-directed synthesis that produces (primary) RNA transcripts from DNA. Thus, if we take ‘code for’ to be equivalent to ‘carries the information for,’ genes get to code, *independently* of anything we might say about the relationship between DNA and proteins. The question then, is, can the coding relationship in transcription be extended *forwards*, so that it reaches proteins? Stegmann’s answer is yes, but only under certain conditions. Here’s the chain of thought: (a) just as DNA contains instructional content for synthesizing RNA transcripts, those transcripts contains instructional content for synthesizing proteins; (b) the bases in a DNA template stand in a neighbour relation to each other, in that C is next to T, T is next to G, and G is next to A; (c) the neighbour relation present in DNA is preserved in the RNA transcript, in that the base in the RNA product corresponding to C is next to the base in the RNA product corresponding to T, and so on; (d) this neighbour relation isn’t disrupted in translation; so (e) the linear order of the DNA template determines the linear order of both the RNA and the protein; so (f) DNA codes for (carries the instructional content for) proteins.

Stegmann’s argument, even if sound, is restricted in its scope. As we have seen, and as Stegmann himself notes, the no-disruption condition, (d), is typically not met for eukaryotes, so the putative result that genes code for proteins may well be restricted to organisms such as bacteria. Elsewhere the putative result is that genes code for RNA, RNA codes for proteins, but genes don’t code for proteins. However, this is by-the-by, because there is a problem with Stegmann’s argument. We think of cooking recipes and computer programs as having instructional content *only* because (i) a producer system – a cooking expert, a computer programmer – has encoded the instructional information in the physical vehicles which carry that information, and (ii) a consumer system – the cook using the recipe, the right compiler – interprets those physical vehicles as instructions. Notice that this is *not* a demand that there be sentient agents in the loop. As mentioned above, the systems that we rightly identify as producer systems and as consumer systems need not literally understand the content of the representations in question. They simply need to be play the right architectural roles. Nevertheless systems of this sort need to be part of the story. But if that’s right, then DNA doesn’t carry instructional information, since (and this point has been bubbling just below the surface of my recent discussion), *there is no relevant producer system in the case of DNA*. Replication is not the same as encoding, so one cannot think of DNA as somehow *self*-encoding (with, of course, the help of some complex supporting chemical machinery). And if one tries to recruit natural selection as the producer system (cf. Maynard Smith’s analogy with the Morse coder), one simply re-confronts the by-now familiar objection that factors which have not been selected for may sometimes qualify as coding within development. There would be no explanation for the positive representational status of such elements.

What this all suggests is that the part played by DNA in development is rather like the part played by sensory input in the perceptually guided action case. DNA doesn’t code for outcomes, but rather provides a causally critical stimulus for subsequent development, a stimulus that is, of course, both determined by the target system’s operational context (one which is environmental in the case of perceptual activity, and historical in the case of development), and partly predictive of the final outcome.

6. The Reach of the Code

The foregoing analysis of protein synthesis suggests that mRNA base triplets are rightly said to code for proteins. But do they also code for phenotypic traits? Some thinkers who have concluded that DNA codes for proteins have proceeded to worry that the reach of the code stops there, and that the claim that genes code for phenotypic traits is indefensible. Indeed, even prominent critics of the whole genetic coding bandwagon are often willing to grant that genes code for proteins, but not traits. Thus Griffiths claims that “the only truth reflected in the conventional view is that there is a genetic code by which the sequence of DNA bases in the coding regions of a gene corresponds to the sequence of amino acids in the primary structure of one or more proteins” (Griffiths 2001, p.395). I shall bring the present treatment to a close then by considering an argument due to Godfrey-Smith (2000b) which questions the extension of the coding relationship from proteins to traits. If this argument is sound, it would compel me to conclude that mRNA codes only for proteins and not also for phenotypic traits. Here is the argument:

The concept of genetic coding is now used to describe and distinguish the *entire causal paths* in which genes are involved. This use of the concept of genetic coding has, I claim, no empirical basis and makes no contribution to our understanding ...

To make this claim is not to deny that at least some causal relations are transitive, and so to deny that genes can causally affect complex traits of whole organisms... The long causal reach of genes is not at issue in this paper. What is at issue is the relation of “coding for...” ... A case from everyday life illustrates the point. Suppose you know that if you order the extra-large pizza, that will have the consequence that the delivery arrives late. This fact does not imply that when you order the extra-large pizza you are also ordering them to make the delivery late. The likely or inevitable *effects* of a message are not all part of the *content* of the message. Similarly, genes can have a causal role which extends beyond the production of proteins, but proteins are all a gene can code for (Godfrey-Smith 2000b, p.35)

Godfrey-Smith is surely right about at least one thing here. His pizza example does indeed show that the “likely or inevitable effects of a message are not all part of the content of the message.” But, on the face of it, this doesn’t provide a mandate for his conclusion that “proteins are all a gene can code for.” It establishes only that we need a way of conceptually screening off those causally downstream (in this case, phenotypic) effects which *don’t* count as part of the content of a particular coding from those that *do*. How might this be achieved, and the reach of the code thereby extended from proteins to traits? Notice that in providing an answer *this* question, we are now at liberty to appeal to factors that we rejected when our target was a *different* question, namely ‘Why should we use representational language *at all*, when trying to understand development? I have given an answer to this latter question, in terms of three conditions: appropriate causal

co-variation, arbitrariness, and homuncularity. What I haven't done yet is give an answer to the former question, the question highlighted by Godfrey-Smith's argument. What Godfrey-Smith's pizza undoubtedly shows us is that the answer to that question, in the case of any outcome-directed representations, cannot be "Whatever the effects are that the representation in question has."

One initially attractive thought is that only the phenotypic outcomes that ensue in the normal developmental environment count as part of the content of the code. But that raises the thorny question of how 'normal' is to be interpreted here. It cannot be interpreted as 'statistically normal,' since one of the lessons of Peter's inevitably late pizza is that some of the effects that an instruction has in its statistically normal environment may not be part of its content. Here it is tempting to revive an appeal to the intended effects, which would succeed in screening off the lateness of the pizza. (The intended effect was an on-time extra-large pizza, not a late one.) As we saw earlier, in the biological case, the appeal to intended effects will be unpacked in terms of selection. But now recall, once again, our (made-up) example of genetic hitchhiking, in which a non-selected-for gene that is causally implicated in the production of blue eyes hitchhikes into the population by being physically connected to a selected-for gene that is causally implicated in the production of a thick coat. Ignoring, for a moment, the matter of whether it's genes or mRNA nucleotide triplets that code, an appeal to selection will straightforwardly deliver the result that our representational element codes not only for proteins, but also for a white coat, since that trait (and thus the related coding element) has been selected for. But if we turn now to the coding element that is causally implicated in the presence of blue eyes, the appeal to selection leaves it devoid of any post-protein content, since blue eyes have not been selected for. The discrepancy here indicates that the appeal to selection falls short of the explanatory mark, since it would surely be uncomfortable to be forced to conclude that the white-coat-related element codes for a white coat, whereas the blue-eyes-related element codes only for proteins. And thinking of the blue-eyes-related gene as being indirectly selected for certainly won't help, since the explanation for its presence is that having a thick coat is selectively advantageous in the environment in question; so that would make the content of the blue-eyes-related gene, 'build a thick coat,' which is surely not what we want.

Although there is undoubtedly more to be said on the reach of the code, the foregoing discussion indicates that it is a difficult and challenging issue. In view of the problems in extending that reach beyond proteins to traits, the default option ought to be to restrict it to proteins. Add this to the conclusion that the locus of coding talk in biological development is mRNA, the base triplets that determine the strings of amino acids constructed during protein synthesis, and the following picture emerges. The power of coding talk in development may be limited at both ends. Such talk doesn't stretch as far back as genes, and it may not stretch as far forward as phenotypic traits.

Acknowledgements

For important critical discussion of the ideas presented here, my thanks go to Andre Ariew, Andy Clark, Sharon Groves, Katie Kendig, Phyllis McKay, Matteo Mameli, Mohan Matthen, John Maynard Smith, Alan Millar, Marion Rosowski, Nicholas Shea, Elliott Sober, Ulrich Stegmann, Peter Sullivan, Terry Sullivan, Denis Walsh, and my 2005

Foundations of Evolution Masters class at the University of Edinburgh.

References

- Bullock, S. 1998. 'The Emptiness of the Self-Contained Coder: Commentary on an extract from 'Twisted Tales' by A. Clark.' *Connexions*, 3: 7-9.
- Cummins, R. 1975. 'Functional Analysis.' *Journal of Philosophy*, 72:741-765.
- DeLisi, C. 1988. 'The Human Genome Project.' *American Scientist*, 76: 488-93.
- Dretske, F., 1981. *Knowledge and the Flow of Information*, Cambridge MA: MIT Press.
- Dupre, J., 2005. 'Are there Genes?' In A. O'Hear (ed.) *Philosophy, Biology and Life*, Supplement to *Philosophy*, 56. Cambridge: Cambridge University Press, pp.193-210.
- Godfrey-Smith, P., 2000a. 'Information, Arbitrariness, and Selection: Comments on Maynard Smith', *Philosophy of Science*, 67: 202-207.
- Godfrey-Smith, P. 2000b. 'On the Theoretical Role of Genetic Coding.' *Philosophy of Science*, 67: 26-44.
- Goodwin, B. 1994. *How the Leopard Changed its Spots: the Evolution of Complexity*. London: Phoenix.
- Griffiths, P. E., 2001. 'Genetic Information: a Metaphor in Search of a Theory.' *Philosophy of Science*, 68: 394-412
- Griffiths, P. E., 2005. 'Function, Homology and Character Individuation.' Paper presented at 'The concept of function in biology and language', Università di Catania, May 2005. Available at:
<http://philsci-archive.pitt.edu/archive/00002351/01/Func.Hom.Chr.PDF>
- Griffiths, P. E. and R. D. Gray., 1994. 'Developmental Systems and Evolutionary Explanation.' *Journal of Philosophy*, 91:6, 277-304.
- Griffiths, P. E and R. D. Knight., 1998. 'What is the Developmentalist Challenge?,' *Philosophy of Science*, 65: 253-258.
- Griffiths, P.E. and K. Stotz., Forthcoming. 'Gene.' In M. Ruse and D. Hull (eds.), *Cambridge Companion to Philosophy of Biology*. Cambridge: Cambridge University Press.
- Halder, G. P., P. Callaerts and W. J. Gehring, 1995. 'Induction of Ectopic Eyes by Targeted Expression of the *Eyeless* Gene in *Drosophila*.' *Science*, 267: 1788-1791.

- Haugeland, J. 1991. 'Representational Genera.' In W. Ramsey, S.P. Stich and D.E. Rumelhart (eds.), *Philosophy and Connectionist Theory*. Hillsdale, New Jersey: Lawrence Erlbaum, pp. 61-90.
- Lorenz, K., 1965. *Evolution and the Modification of Behaviour*. Chicago: University of Chicago Press.
- Mahner, M. and M. Bunge. 1997. *Foundations of Biophilosophy*. Berlin, Heidelberg, New York: Springer-Verlag.
- Mameli, M. 2004. 'Nongenetic Selection and Nongenetic Inheritance.' *British Journal for the Philosophy of Science* 55: 1, 35-71.
- Mameli, M. 2005. 'The Inheritance of Features.' *Biology and Philosophy*, 20:2-3, 365-399.
- Maturana, H. and F. J. Varela. 1987. *The Tree of Knowledge: the Biological Roots of Human Understanding*. Boston: New Science Library.
- Maynard Smith, J. 1998. *Shaping Life: Genes, Embryos and Evolution*. London: Weidenfeld and Nicolson.
- Maynard Smith, J. 2000a. 'The Concept of Information in Biology.' *Philosophy of Science*, 67:177-194
- Maynard Smith, J. 2000b. 'Reply to Commentaries.' *Philosophy of Science*, 67: 214-218.
- Millikan, R. G. 1995. *White Queen Psychology and Other Essays for Alice*. Cambridge, Mass.: MIT Press.
- Oyama, S. 1985. *The Ontogeny of Information*. Cambridge: Cambridge University Press.
- Oyama, S., P. E. Griffiths, and R. D. Gray. 2001. *Cycles of Contingency: Developmental Systems and Evolution*. Cambridge, MA.: MIT Press.
- Pylyshyn, Z. W. 1986. *Computation and Cognition*. Cambridge, Mass.: MIT Press.
- Sarkar, S., 2000. 'Information in Genetics and Developmental Biology: Comments on Maynard Smith.' *Philosophy of Science*, 67: 208-213.
- Sarkar, S., 2005. 'How Genes Encode Information for Phenotypic Traits.' In his *Molecular Models of Life: Philosophical Papers on Molecular Biology*. Cambridge MA.: MIT Press, pp.261-283. This is a revised and expanded version of a paper that appears in C. Hitchcock (ed.), 2004, *Current Issues in the Philosophy of Science*, Oxford: Blackwell, pp.259-274.

- Shannon, C. E., and W. Weaver 1949. *The Mathematical Theory of Communication*. Urbana, IL: University of Illinois Press.
- Stegmann, U. E., 2005. 'Genetic Information as Instructional Content.' *Philosophy of Science*, 72: 425–443.
- Sterelny, K., 1995. 'Understanding Life: Recent Work in Philosophy of Biology.' *British Journal for the Philosophy of Science*, 46: 155-183.
- Sterelny, K., 2000. 'The "Genetic Program" Program: a Commentary on Maynard Smith on Information in Biology'. *Philosophy of Science*, 67: 195-2001.
- Sterelny, K., and P. E. Griffiths, 1999. *Sex and Death: an Introduction to Philosophy of Biology*. Chicago: University of Chicago Press.
- Sterelny, K. and P. Kitcher., 1988. "The Return of the Gene." *Journal of Philosophy* 85:7, 339-361.
- Wheeler, M., 2003. 'Do Genes Code for Traits?,' in A. Rojszczak, J. Cachro and G. Kurczewski (eds.), *Philosophical Dimensions of Logic and Science: Selected Contributed Papers from the 11th International Congress of Logic, Methodology, and Philosophy of Science*. Synthese Library vol. 320. Kluwer, pp.151-164.
- Wheeler, M., 2005. *Reconstructing the Cognitive World: the Next Step*. Cambridge MA: MIT Press.
- Wheeler, M. and A. Clark., 1999. 'Genic Representation: Reconciling Content and Causal Complexity,' *British Journal for the Philosophy of Science*, 50:1, 103-135.
- Wood, D. J., J. Bruner, and G. Ross. 1976. 'The Role of Tutoring in Problem Solving.' *Journal of Child Psychology and Psychiatry*, 17, 89-100.