Genes and Codes: Lessons from the Philosophy of Mind?

Peter Godfrey-Smith
Department of Philosophy
Stanford University

Appears in V. Hardcastle (ed.), *Where Biology Meets Psychology: Philosophical Essays*. MIT Press, 1999, pp. 305-331.

1. Introduction

Do genes really code for biological traits? Of course genes have an important causal role in development and the production of the traits of organisms, but is this causal role a matter of genes coding for their effects?

Some would say that there is not much left to argue about. The view that the relation between DNA and some traits is a coding relation is part of basic textbook biology. A philosopher might disagree with the textbook view, but then that is a rejection of some very well established science -- not something for a philosopher to do lightly.

On another view, the talk of codes and programs in molecular biology has no genuine theoretical role. Although this talk appears constantly in textbooks and popularizations, and even in research articles for illustrative purposes, it is not a real part of the theory. Rather, talk of codes and programs is just a picturesque way of talking about certain causal relations (or perhaps correlations) between genes and traits. This talk could be dropped or denied without loss of explanatory power. So according to this second view, genes do not really code for traits, but to say this is not to break with biological orthodoxy. Philip Kitcher holds a view of this kind.

My own view is opposed to both of these. In contrast to Kitcher, I believe that the idea that some genes code for some traits is a real part of current biological theory. According to the standard picture, as I understand it, both genes and environmental conditions have causal effects on phenotypic traits, but only the genes code. And in contrast to the first view outlined above, I do not think that the idea that genes code is off-limits for philosophical discussion. It

would certainly be folly for a philosopher to deny the standard biological account of the chemistry of DNA and the mechanisms through which it affects the production of proteins. But the concept of genetic coding seeks to add something to that basic picture; it seeks to add a claim about the special nature of some kinds of genetic causation, and a theoretically important analogy between these genetic processes and processes involving symbols and messages in everyday life. Further, the idea of coding itself -- both in general and in genetics -- is not a straightforward concept that everyone understands in the same way. Dissenting voices within both biology and philosophy have claimed that it is a mistake to see genes as coding.[1] A philosophical discussion could be useful at least as a contribution to clarity on this issue, and perhaps to help settle the question.

The question of coding has recently become linked to debates around "developmental systems theory" (Oyama 1985, Griffiths and Gray 1994). Advocates of developmental systems theory claim, among other things, that the idea that genes code for traits is part of a picture that assigns to genes a false causal priority in development and evolution. Developmental systems theory opposes the idea that some of the factors that contribute to development are sources of information or form, while other factors are mere background, support or raw material. For the developmental systems view, it is a mistake to think that genes have a causal role that is different in kind from that of non-genetic factors. At least in principle, both kinds of factor can carry information, both can be inherited, and so on. So developmental systems theory is associated with various claims of symmetry for genetic and non-genetic factors.

As far as possible, in this paper I will discuss coding in isolation from questions about the causal priority of genes.[2] Someone could claim that genes really do code for traits while denying that genes have pre-eminent causal importance. The other combination of views (genes as pre-eminent, but not coding) is possible also. But the more general issue of symmetry between the causal roles of genetic and non-genetic factors will be discussed in some detail.

If genetic coding is recognized in some cases, it need not be recognized in all. One also has to decide exactly what messages genes can contain. Do genes only code for proteins, or can a gene also code for penicillin resistance, if the protein produced has a key role in producing this resistance? That question is a question about "how far out" coding can reach. A different but related question arises when a single gene makes a difference to a complex trait that can only be

built with many other genes and developmental factors (a trait like delayed sexual maturity, perhaps). The first of these two questions (the "how far out" question) will be discussed in this paper; the second (the one involving complex traits) I leave for another day.

The question of whether genes code would be easier if there was a widely accepted theory of coding and representation which we could apply to the problem. There is, indeed, a widely accepted framework in which some questions about messages can be addressed -- the mathematical theory of information. And some writers have applied information-theoretic concepts directly to the problems involving genes (Maclaurin forthcoming). But I will argue that this approach is unlikely to furnish a full solution to the problem.

Within philosophy, perhaps the most promising place to look is the literature on physicalist theories of representation in the philosophy of mind.[3] These theories aim to state, in physicalistically acceptable terms, necessary and sufficient conditions for an internal state of an organism being a representation of some specific object or state of affairs in the organism's environment. If this debate had produced a clear winner we might be able to apply this winning theory (or adaptations of it) to a range of other problems involving representation and meaning -- including the problem of the relation between genes and traits. The discussion has not, however, produced a consensus.

Even though there is no consensus, perhaps this work can at least provide some clues as we wrestle with the genes. I will attempt to make some connections in this paper.

This paper does not, however, give a full solution to the problem of genetic coding. It does not even reach a firm decision on whether genetic coding is real, and whether the idea is useful. This is an exploratory discussion. Much of the aim is just to outline some of the available options.

2. The Genetic Code and the "Gene for" Concept

In this section I will discuss some common ways of talking about the genetic code within biology, and then look at how some recent philosophical literature has handled both the question of coding, and the idea that a gene can be "for" a phenotypic trait. The first part of this section will also give a refresher on basic facts about the role of DNA in the manufacture of proteins.

We start with two statements from biology textbooks about the role of DNA.

> The information dictating the structures of the enormous variety of protein molecules found in living organisms is encoded in and translated by molecules known as nucleic acids. (Raven, Evert and Eichhorn 1992 p. 59)

> [DNA] contains a coded representation of all of the cell's proteins; other molecules like sugars and fats are made by proteins, so their structures are indirectly coded in DNA. [DNA] also contains a coded set of instructions about when the proteins are to be made and in what quantities. (Lodish et al. 1995 p. 10).

These are the sorts of statements I take as evidence for the claim that contemporary biology attributes to DNA a special set of properties that are described in semantic terms. While standard views in contemporary biology certainly see these coding relations as fundamentally causal in nature, these views also hold that among the various causal relations involved in development and metabolism, some causal relations are special because they involve the interpretation of a message or the "expression" of coded instructions.

Let us look briefly at the relevant biological processes, and the terminology used for them within biology. Two main steps are distinguished in the causal chain between DNA and a protein. "Transcription" is the process in which DNA gives rise to mRNA ("messenger RNA"). Then "translation," which occurs at the ribosomes, generates the protein itself. The molecule of mRNA produced during transcription is formed using a stretch of DNA directly as a template, and the mRNA contains a sequence of "bases" which corresponds, by a standard rule, to the sequence of bases in the DNA from which it was derived. In organisms other than bacteria, the mRNA is usually processed (in ways I will discuss later) before it is used in translation. Then, at the ribosomes, the processed mRNA is used to direct the formation of a chain of amino acids -- a protein.

In this process of translation, a crucial role is played by another kind of RNA molecule, tRNA (or "transfer RNA"). Molecules of tRNA bind to particular amino acids (of which there are 20 kinds), and at the ribosomes these tRNA

molecules bind to specific three-base sequences in the mRNA. So each triplet of bases in the mRNA is associated, via the chemical properties of tRNA, with a particular amino acid. The "genetic code" is, strictly speaking, the rule linking RNA base triplets with amino acids, and this interpretation of the RNA determines the interpretation of the DNA from which the mRNA is derived.

The ribosome moves along the mRNA chain and as it goes a chain of amino acids is assembled, with the sequence of amino acids corresponding exactly to the sequence of bases in the mRNA by the rule comprising the genetic code.[4] As there are four bases in the mRNA (almost the same four as in DNA) there are 64 possible triplets. Of these, 61 specify particular amino acids; some amino acids are specified by as many as 6 different triplets. The three remaining triplets are "stop" signals. The chain of amino acids folds (and may be processed in other ways) to produce a finished protein. Protein structure is described at four different levels, of which the primary and tertiary are most important for our purposes. The primary structure of a protein is its sequence of amino acids; the tertiary structure is the three-dimensional folded shape of a single amino acid chain. The causal role of proteins depends greatly on their tertiary structure.

There is much more to all these processes of course, and on some views the extra detail is essential to an understanding of what is going on. But on more standard views, the information given above outlines the core of the process by which the genetic message is expressed.

Philosophically, the term "translation" seems a strange one, even within the standard picture of the genetic code. Translation, in ordinary usage and in philosophical theory, takes a message from one symbol system or language to another. But while the standard view sees the DNA sequence as a sort of language, amino acids and proteins are not usually understood as coding for anything (unless that is their job elsewhere in the body). On standard views, DNA and RNA are messenger molecules but the series of messages ends when the protein is formed. So the process of "translation," as it is usually conceived, would be more accurately described as "interpretation." And sometimes biological discussions do use that term, although it is not nearly as standard as "translation."[5]

Turning to recent philosophical discussion of the relation between genes and traits, a natural place to start is with a work that does not make any claims about coding at all. Sterelny and Kitcher (1988) give an analysis of what is involved in some gene being a "gene for" a particular trait, but they do not

express any part of the analysis in terms of coding. Roughly, to talk of a "gene for X" in the sense of Sterelny and Kitcher (1988) is to talk of a reliable correlation, in normal genetic and non-genetic environments, between the gene and the trait.[6]

Griffiths and Gray (1994), arguing for the explanatory symmetries characteristic of developmental systems theory, claim that an analysis of "gene for X" in the style of Sterelny and Kitcher must allow that environmental conditions or cytoplasmic factors can be "for" particular traits in the same sense that genes can. This is because when we call a gene a gene for X, we hold certain environmental factors constant, as a background condition. But we can also hold genetic factors fixed, as a background condition, and speak of environmental or cytoplasmic factors "for" traits. Griffiths and Gray took Sterelny and Kitcher to be seeking an analysis which would retain the idea that genes code for traits as well as causing them, so Griffiths and Gray took this symmetry between genes and environment to be a problem for Sterelny and Kitcher.

In their responses to Griffiths and Gray, Sterelny and Kitcher diverge in interesting ways. Kitcher (forthcoming) accepts Griffiths and Gray's point about the explanatory symmetry of genetic and non-genetic factors, and accepts that there are "environments for" as well as "genes for" traits. He claims that his and Sterelny's original reconstruction of the concept of a "gene for X" is entirely compatible with this move. But when Griffiths and Gray suggest that Kitcher would not want to say that an environmental feature codes for a trait, Kitcher dismisses their talk of coding as "a rhetorical flourish irrelevant to the discussion." (ms p.19).

When I objected to Kitcher that standard views in genetics do see genes as coding while environmental conditions cannot code, Kitcher replied (in correspondence) that there is no need to make literal sense of claims about genes coding for traits. It is just a colorful mode of talk that has no role in the explanatory structure.

While Kitcher steers clear of coding and all properties akin to it, Sterelny takes the opposite route in his response to the symmetry arguments of the developmental systems literature. Sterelny, Smith and Dickison (1996) accept that genes can be ascribed semantic properties as well as causal properties. Their preferred term is a philosophically strong one: "the genome does represent developmental outcomes" (1996 p. 387, emphasis in original). In response to the symmetry arguments, Sterelny, Smith and Dickison claim that some non-genetic factors in development have the same kinds of properties that genes are usually

taken to have; there are both genetic and non-genetic replicators. Replicators have been shaped by selection for their developmental role, and replicators "represent phenotypes in virtue of their functions" (p. 387). I will discuss Sterelny, Smith and Dickison's view in more detail in section 5 below.

Sterelny, Smith and Dickison also talk a good deal about information, as many others do. Griffiths and Gray, who oppose many standard ways of talking about genes, do not object to the idea that genes contain information, so long as the use of informational concepts is not restricted in its application to genes, and used to "privilege" genes over other developmental resources (p. 283). I will discuss the possible role of the concept of information in the next section.

3. Indicative and Imperative

As we have seen, philosophers and biologists use a range of semantic expressions when talking about what genes do. Sometimes it is said that genes carry information about traits, or represent the outcomes of developmental processes. Sometimes genes contain coded instructions. We find picturesque terms such as "blueprint," and far more neutral terms such as "specify."[7] The differences between these formulations are important, as some claims are empirically more plausible than others, and some attributions of meaning or content raise more philosophical problems than others.

A detailed analysis of how exactly genes code can be expected to choose some one of these formulations as best. As I will not defend any particular analysis in this paper, I leave some of these issues fairly open. But some preliminary points can be made. In this section I will discuss one important distinction between different kinds of semantic properties that genes might have.

Ruth Millikan, adapting older terminologies, distinguishes in her general theory of signs between "indicative" and "imperative" representations (1984). Roughly, an indicative representation is supposed to describe how things are, and an imperative one is supposed to bring something about. So a declarative sentence is an example of an indicative representation while a command is an example of an imperative representation.

If genes are representations at all, which kind are they? I claim they must be seen as imperative representations. Their role is to prescribe rather than describe. So within the family of semantic terms that philosophers and biologists have used about genes, the most appropriate ones are those that suggest

imperative rather than indicative contents. Viewing genes as containing "coded instructions," as "prescribing" or "dictating," has more chance of being right than viewing them as "describing" anything.[8]

The difference between assigning indicative and imperative contents is readily seen in cases of error, non-compliance and misrepresentation. Suppose you order a pizza, but what arrives is pasta. In such a situation, your message is not faulty or erroneous. The mistake was made by the people who received the message and filled your order. If, on the other hand, you did not order but predicted (or guessed, or claimed) that they would bring pizza, and they brought pasta, then your claim is where the error lies.

In the case of genes, which party is "at fault," according to the standard picture, if the protein produced on some occasion does not match the specifications of the DNA? I claim that on standard views about genes and coding, it is not the gene that is at fault for misdescribing the protein (for guessing pizza when they brought pasta). Rather, the interpreting mechanisms in the cell are at fault, for failing to comply with what the DNA instructed.

So even before the details of an analysis of coding properties of genes have been worked out, it is clear that the aim should be an analysis of DNA's capacity to carry messages with imperative semantic content. DNA, if it contains a message at all, contains instructions rather than descriptions.

Partly because of this, one of the most popular ways to ascribe semantic properties to genes -- a way using the concept of information -- is not a good approach to the problem.

There is a variety of ways in which the concept of information is used in describing what genes do. It is common to say that genes carry within them information, in coded form, about the proteins made by the organism, and perhaps information about whole complex traits the organism exhibits. Here the genes are playing a role like that of a message. At other times the genes are said to be a source of information used by the cell or by the organism in development. This may (although it need not) mean something different from the idea that the genes are a message.

I will argue that the only proper role the concept of information has here is a weaker, less interesting role than it is often taken to have. There is a weak sense of "information" in which anything is a source of information if it can occupy a variety of possible states. And in this sense of information, if the states of some X are reliably correlated with the states of some Y, then X carries

information about Y.[9] This is the sense of information dealt with in the mathematical theory of information (Shannon 1948), and refined into a semantic theory by Dretske (1981). Information in this sense has also been discussed more informally by philosophers as "natural meaning."

Information, or natural meaning, is everywhere. It does indeed connect DNA with proteins and phenotypic traits, but it connects them in both directions and connects both of these to environmental conditions as well. DNA sequences have many possible states, as do proteins. Given background conditions (which define a "channel") the genes carry some information about the proteins produced by a cell. The proteins produced also carry some information about the genes responsible for them. In both directions the transmission of information is imperfect, for a variety of reasons.

So we can regard the environment as a background condition against which genes carry information about phenotypic traits. But as defenders of developmental systems theory insist, we can also view genetic conditions as background conditions or part of the "channel." Against such a genetic background, we can see environmental conditions as carrying information about phenotypes. And we can also see phenotypes as carrying information about environmental conditions as well.

If any of these attributions of informational properties are acceptable, then all of them are. In some cases and in some directions there will be more information carried than there is in others, but that does not affect the basic point about the ubiquity of information.

Attributions of informational properties of this kind cannot be used to analyze the special role played by concepts of coding in molecular biology. This is clear from the fact that although information is ubiquitous and runs in these cases in both directions, the coding relationships discussed in molecular biology are not. Coding is (i) specific to the relationship between genes and phenotypes, and (ii) asymmetric, as genes code for phenotypes but not vice versa.

This is not an argument that talking about information can have no useful role in molecular biology. Talk of information is often a useful way of picking out correlations and causal relationships of various kinds. The argument is just that this concept of information gives no grounding for the asymmetry expressed by the idea that while genes and environments both causally affect phenotypes, only the genes have their effects by coding for phenotypic features.

Compare another usage of "information" in molecular biology. Biologists sometimes talk about the information that genetic variation within and between populations carries about phylogenetic relationships. We can learn from analysis of DNA sequences the order in which various species split off from one another, and other historical facts of this kind. Different parts of organisms' genomes carry different amounts of information about these historical relationships ("junk" DNA is especially helpful). Sometimes one kind of analysis will be useful for more recent history while another analysis will be used to reconstruct more distant events. All of this is often described in terms of the "information" carried by patterns of genetic variation. But it is obviously only information that we use, not information that is part of any explanation of the causal role that genes play in development or evolution. There is no more to this kind of information than correlation or "natural meaning"; the genes are not trying to tell us about their past. Genes in this sense are like tree rings (a standard example of natural meaning). Similarly, the fact that we might sometimes be able to read the composition of a protein off a DNA sequence does not imply that the cell literally reads the composition of the protein off the sequence. If the role played by the concept of genetic "information" in explanations of development is something more important than the trivial role it plays in phylogenetic reconstruction, then a novel and richer concept of information-transmission must be developed.

I do not say that this is impossible; there may be other, richer concepts of information that biologists and philosophers could develop and apply here, and some analyses might use the concept of information as one component in a theory along with others (as Dretske does in his philosophy of mind). I am not claiming that no analysis of the semantic content of DNA that uses the idea of "information" can possibly succeed.[10] But I do think this is an unpromising way to proceed. It is unpromising first because the idea that genes "carry information about" proteins or phenotypic traits is naturally understood in terms of indicative rather than imperative content.[11] And it is unpromising also because it will always be hard to keep a richer concept of information distinct from the original sense -- a sense with genuine usefulness in biology -- in which information exists whenever there is reliable correlation.

4. Analyses in Terms of Developmental Role

This section and the next will look at two distinct ways to develop an analysis in which genes are assigned coding properties. On the first approach, discussed in this section, genes code in virtue of their role in developmental and metabolic processes.[12]

In this paper I will understand "development" as a process that takes place strictly within a single generation. Development is set of local causal processes through which early stages give rise to later stages, something which could in principle be described without any reference to evolutionary history. Genes, as everyone agrees, play a causal role in such processes. One possible view on the question of coding is that the causal role DNA plays in developmental processes is one that can fairly be described in terms of its bearing coded instructions; the content of specific DNA sequences is determined by a rule of interpretation that derives from this causal role. I will call these "developmental role" theories. They have a loose analogy to "conceptual role" theories of meaning in the philosophy of mind (Block 1986).

As I envisage views of this kind, the peculiar characteristic of DNA that justifies its being treated as a code lies in the fact that its sequence is physically read by the cell during the construction of proteins. The cell first creates an mRNA molecule whose sequence corresponds to the sequence of bases in the DNA, and then part of the cell's machinery physically moves along the mRNA molecule, at each step interacting with the base sequence, producing with each step a chain of amino acids whose linear structure corresponds, by a standard rule, to the linear structure of the mRNA. This process, according to an analysis of coding in terms of developmental role, is one in which the mRNA is directly read by the ribosomal machinery, and whole process is one in which the DNA is read as well.

On the developmental role view, what makes genes into coded messages is not just the specificity of their causal role -- the fact that particular DNA sequences reliably give rise to particular products. The causal specificity of DNA is one important feature but not the only one. This is because "specificity" is a property that applies to a wide range of causal relations. A cutting enzyme might be highly specific in what it cuts. A raw material might only be usable in one specific building project. So specificity alone is not the issue; it also matters how this specificity arises. What makes the "genetic code" into more than just a set of causal associations is the nature of the processes that underly those associations.

DNA is causally specific through being read; other factors (like enzymes) have causal specificity of different kinds and for different reasons. A family of enzymes might have its causal role described by a general principle linking particular enzymes to particular reactions, but this will not be an interpretation rule in the case of the enzymes, because enzymes are not read by these processes.

As I stressed in the Introduction, the idea that genes code because of their developmental role may or may not be coupled with a claim about the preeminent causal importance of genes. One might hold that genes code without holding that "genes are destiny." The point of the concept of coding, on a developmental role analysis, is to pick out one particular causal role among many. Within developmental and metabolic processes there are raw materials (like amino acids), cutters and joiners (enzymes), stores of energy (like ATP), readers and assemblers (ribosomes)... and there are coded instructions as well (the genes). Raw materials and stores of energy might be just as important as messages, but they are different kinds of causal players.

An analysis along these lines will face a variety of challenges. One philosophical line of objection proceeds by claiming that cells cannot possibly "read" anything because the concept of "reading" is one that only has a place in a system of human conventions of public symbol use. So it is nonsense to say that genes contain a message read by the cell.

The ordinary sense of "reading" may well be one that requires that the reader be an agent with mental states. But machines which "read" in extended senses, like bar-code readers in supermarkets, are all around us. A biologist can reply that though cells read in an extended sense, this is a sense which has a useful role in our understanding of many kinds of machines. Possibly the "reading" of supermarket bar-code readers depends for its status on the machines being the products of human design, so the genes and ribosomes are in this respect even further from the ordinary use. But it can be argued that "reading" in this sense is still a distinctive kind of causal process, one with real similarities to ordinary human reading and interpretation. Reading as done by supermarket machines is physically different from weighing and imprinting; different also from "guessing" and other actions that machines might perform in extended senses.

So a biologist will very likely accept that cellular mechanisms only read things in an extended sense of "read." That would imply that DNA only contains coded instructions in an extended sense of "coded instruction." But this

concession could reasonably be regarded as a minor one by a biologist; it is consistent with the claim that the processes in which DNA is involved have a remarkable and theoretically important similarity to ordinary processes of reading and interpretation. And that is what many biologists might regard as the important point -- it does not matter how exactly the genetic senses of "reading" and "interpretation" are connected to everyday senses.

A more important challenge derives from biological considerations. DNA is not "read" by the cell in exactly the ordinary sense, admittedly, but is it appropriate to see these processes as akin to reading at all? The suggestion I made about reading embodies quite a controversial claim about the causal processes in which DNA is involved. It might be objected that it is more accurate to see the DNA sequence as having the role of a template, rather than something that is read by the cell. Talk of "reading" might be harmless in certain informal contexts, but the processes of protein synthesis are not of a kind that will support the linguistic analogy once one looks closer.

When we look closer, we see that the sequence of bases in DNA is transferred onto an mRNA sequence by a process in which the DNA acts directly as a template for the synthesis of mRNA. Then the mRNA itself acts as a template along which a chain of amino acids forms. The ribosomes where the protein is formed are not much like readers; rather, they are an elaborate kind of scaffolding where certain reactions take place spontaneously. So according to this biological objection, the idea that DNA functions as a template is an alternative to the view that DNA contains instructions which the cell reads. Using a template and using instructions are two different ways in which structure can be transferred or conveyed. On such a view, the standard rule linking DNA base triplets to amino acids is seen, again, as describing a set of causal specificities and no more.

To this objection the advocate of genetic coding has two replies. It might be replied that a description in terms of reading is more accurate than a description in terms of mere templates, perhaps because of the combinatorial properties of the RNA/amino acid relationship. Alternately, perhaps the two descriptions are compatible. Watson et al. (1987), the king of molecular biology textbooks, puts a lot of theoretical stress on the concept of a template but evidently does not see this as antithetical to the idea of coding. Both issues are hard to settle and I will not attempt to solve the problem here.

Another feature of the developmental role approach to coding is that there will be interesting differences between different organisms, with respect to how appropriate it is to see genes as containing a code. All views on this matter must accept and make room for a range of phenomena that cause trouble for the most simplistic views of the coding properties of genes. These phenomena are sometimes specific to certain kinds of organisms. Viruses sometimes have the sequence of one gene completely embedded inside that of another, which is read differently. Both viruses and bacteria have genes whose sequences partially overlap. At the other end of the scale of economy, genomes of eukaryotes (roughly, all organisms except bacteria) contain introns. This is DNA sequence which is transcribed into RNA but then removed before translation. So the protein produced typically does not correspond to the sequence of any contiguous stretch of DNA. There can even be alternative ways of splicing together pieces of mRNA, to form different products from the same initial "primary" transcript.[13] But prokaryotic genomes do not contain introns.

Should a view that analyzes coding in terms of developmental role concede that prokaryote genomes more clearly or unproblematically contain coded messages than eukaryote genomes? For some this would be a problematic conclusion. If so, this suggests that developmental role analyses should not require too strict a correspondence between protein structure and DNA sequence.

Interestingly, Sarkar (1996) thinks that a different conclusion should be drawn; he thinks that the most important objections to the idea that DNA contains a code only apply in the case of eukaryotes. If what was true for bacteria was also true for elephants, Sarkar says, this would "make the linguistic view of genetics palatable" (1996 p. 860). So as I read Sarkar, he thinks that bacterial genes might reasonably be seen as coding for traits, as a consequence of their developmental role, while genes in eukaryotes do not contain a code.[14] He bases this argument not just on the presence of introns and other non-functional DNA, but on a variety of other complexities peculiar to eukaryotes. In particular, RNA is sometimes "edited" in ways that go beyond the removal of introns; sometimes the editing involves substitutions or additions of bases. In eukaryotes the sequence of amino acids in a protein need not correspond exactly by the standard interpretation rule to any stretch (contiguous or not) of the DNA sequence.

My own view is cautiously opposed to Sarkar's on this point. It seems to me unlikely that the differences between eukaryotes and prokaryotes are such that coding properties might be found in the DNA of one but not the other. And in his treatment of these issues, Sarkar seems sometimes to require far more transparency of interpretation for genetic codes than one would normally require for ordinary public messages. For example, Sarkar says "natural languages do not contain large segments of meaningless signs interspersed with occasional bits of meaningful symbols" (p. 863), whereas eukaryote DNA does contain such junk. But surely a great deal of symbolic interaction in everyday life is interspersed with meaningless extra elements. Admittedly, to reach a ratio of 95% junk to 5% real information (as is often claimed for human DNA) we might have to look at some unimpressive regions of social life (I leave it to readers to insert their preferred examples), but ordinary concepts of meaning and representation certainly do not preclude messages containing in their physical structure much that has to be edited out.

Within philosophy of mind, many theories of representation have difficulty accounting for the possibility of error and false representation.[15] This is thought important because it is taken to be essential to the concept of representation that wherever there is representation there is the possibility (in principle) of misrepresentation. So theories of genetic coding that rely purely on developmental role must face the problem of specifying how misrepresentation, error and non-compliance can exist. The problem is not one of showing how the DNA can err in description; the problem is showing how the cell might misread the DNA and fail to do what is instructed. An analysis drawing only on the actual causal role of DNA can distinguish between the common and the uncommon -- in some particular case a process involving the genes might give rise to a product that differs from the normal product. But philosophers of mind have labored long and hard over the fact that the ordinary concept of misrepresentation seems to be one in which the distinction between proper and erroneous in semantic contexts is <u>not</u> the same as the distinction between the common and the rare. It is also not the same as the distinction between the beneficial and the harmful. So for an analysis of the genetic code purely in terms of developmental role, either whatever is common is what is "supposed" to happen, or else there is no sense in which some things are supposed to happen while others are not.

Once the issue of error and malfunction is raised, some philosophers of mind will think that what is needed at this point is an appeal to a concept of "proper functioning" based in evolutionary history. There are several ways in which this might be done. A minimal way would be to add an evolutionary concept of normal or proper functioning to an analysis that relies in all other respects on the developmental role of DNA. This might be quite a promising way to proceed. But another possible response is to think that more extensive use of evolutionary concepts is needed. Such a move would build a more substantial bridge between genetics and the philosophy of mind.

5. Analyses in Terms of Evolutionary History

If genetic coding is analyzed in terms of evolutionary history, nothing about the pattern of interaction between DNA and proteins, considered just as a physical process, makes a DNA sequence into a message. Rather, the key fact is that the evolutionary history of these mechanisms is one that has given to DNA certain biological functions. DNA has the function of coding for amino acids, or the function of instructing the cell which proteins to produce (or some other functional property of that kind). This approach is taken by Sterelny, Smith and Dickison in "The Extended Replicator" (1996).

A view of this kind is able to draw on a range of influential ideas in the philosophy of mind. Millikan's view (1984), which I have drawn on earlier, is the most elaborate evolution-based theory of representation, and the most potentially adaptable to the case of DNA. Others include the theories of Papineau (1993) and Sterelny (1990). In philosophy of mind these are sometimes called "teleofunctional" theories.

Two principles are basic to these theories, and both principles might be applied to genes. One is the idea that a rich concept of function is essential to an understanding of representation and other semantic properties. The other is the idea that functions are to be analyzed in terms of evolutionary history: functions are effects or capacities which have been selected for.[16]

DNA and the mechanisms that interact with it are certainly products of evolution, and some of their effects have been selectively important. So it should be possible to assign functions to these structures. And such an assignment of functions might provide a semantic analysis of what the genes are telling the cell.

There are many ways to develop the specifics of such a view. One dimension on which historical theories can be compared concerns how much of a role the theory gives to factors other than evolutionary history. Some theories might hold that genes code simply because of the special properties of their selective histories; others might see the appeal to history as one factor used along with others to analyze coding. I will discuss examples of both approaches below. Sterelny, Smith and Dickison's view puts almost total weight on historical properties in explaining why genes code. But in response to problems with that view, some might add non-historical factors to the story.

Another distinction has to do with how important the systematic properties of DNA are taken to be. Suppose we are seeking to assign biological functions to a DNA strand. There are two ways in which this could be done.

Option A: The particular DNA strand might have its own biological functions, as a consequence of the success under natural selection of strands with the same sequence (or almost the same sequence) that gave rise to the present strand. That is, we can look to the selective history of particular DNA sequences. Their interpretation is determined by the effects they have had that have led to their selective success.

Option B: It might be that DNA in general has the function of specifying proteins, as a consequence of the selective history of the entire machinery of protein synthesis. Then the interpretation of any particular sequence of DNA would be determined not by its own past effects, but by a standard interpretation rule that applies to all DNA, a rule determined by the evolutionary history of DNA in general.

Option A has the consequence that a new mutation has no natural interpretation; whatever the cell does with it, there is no reading/misreading distinction. On Option B, a new mutation does have a natural interpretation, as a consequence of the general rules of DNA interpretation, and a new mutation could in principle be misinterpreted even on its debut. I imagine that many readers will regard this as an argument against Option A. The DNA reading mechanisms are supposed to react in particular ways to any bit of DNA, new or old. That suggests that the new sequence does have a natural interpretation.

In general I think Option B is the right approach for a view that analyses coding in terms of evolutionary history. However, there might be some place in the picture for the reasoning behind Option A. Suppose a new mutation appears which for some reason always interferes with the reading mechanisms. The mutation does have effects on protein synthesis, but effects which are not part of the usual pattern of DNA interpretation. However, the effects are useful and the mutation is successful under natural selection. Call the protein it produces by this nonstandard route "protein Z."

Does this sequence of DNA code for protein Z? Inducing the cell to make protein Z is the sole effect the gene has that explains its success under natural selection. So on many theories of biological function, the gene will have the function of producing protein Z. But the causal route of this production is not one which involves the usual pattern of DNA interpretation. In a sense, this gene is being selected for despite being always misinterpreted. In a causal and functional sense, this is a gene for Z, but the gene does not code for Z. Some readers might find this an odd conclusion to draw, and I will discuss it further below.[17]

Incidentally, this case is not as biologically far-fetched as it might sound. "Frame shift" events sometimes occur during translation; the ribosomal machinery reads a quadruplet as a triplet, or backs up and reads a base twice, and then carries on reading triplets after that anomalous event. This happens through a variety of causes, but in some organisms there are "slippery" mRNA sequences which act, in the context of the cell, to induce frame shifting, and this frame shifting can be required for normal development.[18] I don't know if there are cases of frame shifting where a sequence is always read out of frame, however.

My response to the protein Z case is guided not just by a preference for Option B, but by a general assumption about the relationship between biological functions and semantic properties. The assumption is that not everything with a biological function in DNA replication and expression has semantic properties. This assumption about functions and representation is surely true in many contexts -- legs are for walking, but they do not represent walking. They do not (usually) represent anything, despite having biological functions. Something can have the function of producing a certain object or outcome in biological processes, without representing or coding for that object or outcome. I apply the same assumption to DNA.

I highlight this assumption about functions, because the solution to the coding question defended by Sterelny, Smith and Dickison (1996) works by denying this assumption for a special class of cases. Sterelny, Smith and Dickison claim (as I interpret them) that an analysis of the biological functions of DNA in developmental processes suffices to determine a semantic interpretation of genes; all factors which have biological functions within development represent the outcomes of those developmental processes. Although Sterelny, Smith and Dickison do not claim this tight connection between function and representation for all biological functions (so they would accept my claim above about legs), they view developmental functions as a special case. Here the connection between function and representation is especially tight.[19]

I do not think this is a satisfactory solution. There are entities which have the function of playing a certain role in development, where it seems quite implausible to assign them representational properties. For example, an enzyme such as an aminoacyl-tRNA synthetase has a certain function in protein synthesis, and this function can be understood in terms of a selective history. These enzymes have the function of attaching particular amino acids to particular tRNA molecules. That is the effect they have which explains why they are there. Though that is the function of these enzymes, and the enzymes are very causally specific, the enzymes do not code for or represent their products. There might be no representation without function, but there is function without representation.

Consequently, in understanding how it might be that genes code for traits, it is not enough to have an analysis that assigns them the biological function of having a certain causal role in the production of proteins. Something more is needed.

Here is one suggestion, a friendly amendment to Sterelny, Smith and Dickison's proposal. Perhaps DNA codes because the DNA sequence is supposed to help produce a product with a certain abstract relation to its own structure. DNA sequences have historically entered into causal processes in which a certain mapping relation between DNA sequences and amino acid sequences has been important in the evolutionary history of these mechanisms. That "mapping relation" is, of course, the familiar set of relations usually described as comprising the interpretation of the genetic code. A certain pattern in which DNA sequences have generated (with the help of other cellular machinery)

proteins with a certain abstract relationship to themselves, has been central to the evolutionary success of these mechanisms.

Would this move provide Sterelny, Smith and Dickison with a way to forge the link they want between function and representation? This amendment would certainly give them a way to respond to the problem with enzymes that I raised above. Although an aminoacyl-tRNA synthetase is supposed to have a certain causal role, this role is not one which involves production of products with abstract mapping relations to itself (except in trivial senses).

This suggestion might not take us far enough, however. I do not know of a real case that might provide a counterexample, but here is a hypothetical case. Suppose there is a family of enzymes whose function is to join identical protein subunits into larger structures. One enzyme joins pairs of the subunits together, and it has two binding sites at which the subunits attach when being joined together. Another enzyme makes units out of three of the subunits, and it has three binding sites. Another joins four subunit molecules... and so on. Here we have enzymes whose function is the production of certain products, where the products are supposed to have a certain abstract relation to the enzymes (an enzyme with two sites produces a double molecule, and so on). But would this be a case where the enzymes represent or code for their products? If intuitions are to be trusted in this arena, it seems to me that these enzymes would not represent their products, in any ordinary sense of "represent." Again we have function without representation.

What might be the next step down the same road? The previous suggestion was inspired by parts of Ruth Millikan's semantic theory (1984). The concept described above is related to Millikan's concept of an "imperative intentional icon." However, it is not quite the same as Millikan's concept. Maybe we should move still closer to her view?

If so, the most important further step would be to introduce what Millikan calls "producers and consumers" of the representation. For Millikan (and for some other views) all representations mediate between two specific kinds of functionally characterized things.[20] "Producers" are supposed to produce the representation, in the performance of their functions. "Consumers" are supposed to have their activities modified by the state of the representations they receive from the producers. In particular, in the case of imperative signs, the sign affects its consumers in such a way that the consumers acquire the function to produce a

certain state of affairs in the world. This state of affairs is the sign's compliance condition, its content.

If we introduced this idea, we could rule out the hypothetical case involving enzyme binding sites that I described above. The only problem is: who are the producers and consumers?

There are various options for answering this question, but none of them seem attractive. The consumers are probably easier to handle than producers. The most likely initial candidates for consumers are the ribosomes, the machinery where protein synthesis actually occurs.

A first problem with this idea is that it is mRNA, not DNA, which interacts with the ribosomes and functions in the assembly of proteins. So applying Millikan's view in this way, we could see the mRNA as instructing the ribosomes, and then see the DNA in the nucleus as the producer of the representation. In some ways that makes sense, but it is not a solution that vindicates standard views about coding. Those views, recall, assign content to the DNA in the nucleus, not just to the mRNA. So perhaps we might see the ribosomes as consumers of both the DNA and the mRNA -- the mRNA is an intermediary.

Another problem is that regarding the ribosomes as "consumers" in Millikan's sense might be giving them too much credit. If there are to be consumers, it probably makes more sense to see them as comprised by a set of interacting factors, including the ribosomes themselves, the tRNA molecules, and various proteins and other molecules that initiate, control and fuel the process.

So to see the DNA as interpreted by a Millikan-style "consumer," we must identify this consumer with a range of factors, including those that produce mRNA and those that use mRNA in building proteins. I do not see this "distributed" nature of the consumer as a problem for such an analysis.

But who is the producer of the DNA message? It might be the entire previous cell that gave rise to this one, in which case messages would get passed indefinitely back through cells across the generations. But then it is harder to see why these "producers" are giving particular instructions to the machinery of the new cell. We should not say (as people sometimes do) that the DNA is a message passed from one generation to the next, as the DNA is part of what gives rise to the next generation. The next generation is a product of the DNA (plus other inherited resources) rather than its consumer.

I will not follow up these specific ideas further here. Some might think that I have unfairly dismissed a viable alternative, or might think there are options I have not discussed that will do the trick. (Sterelny, Smith and Dickison say at one point, for example, "a gene can have the function of telling the developmental program how to build haemoglobin molecules" (1996 p. 338, emphasis added). Maybe something like that could work.)

In general terms though, if one accepts that Sterelny, Smith and Dickison's original 1996 view is unsatisfactory but on the right track, the obvious thing to do is to constrain the kinds of causal role that are associated (when suitably historically embedded) with representational properties. The problem will be to do this without forcing a misleading or ill-fitting set of distinctions onto our empirical picture of the causal role of genes. Insisting that we locate a set of "producers" and "consumers" might, for instance, be forcing onto the cell a framework derived from elsewhere and with no empirical motivation in this context, just to retain the idea that genes code for something.

Recall, in contrast, the first amendment I made to Sterelny, Smith and Dickison's view above. This was the suggestion that DNA codes for proteins because the DNA sequence is supposed to help produce a product with a certain abstract relation to its own structure. Even though this view might not make all the right intuitive distinctions, nonetheless this proposal does make use of empirically well-motivated distinctions. We might decide to only retain a modified or weakened concept of coding, one which can be understood in terms of this view.

6. Distality: Amino Acid Sequences, Folded Proteins and Traits

I will discuss one more topic before drawing some conclusions. This topic is: what exactly can the genes code for? As far as causal relations are concerned, we can trace a chain as far as we please, from the proximal effects of genes (amino acid sequences), through to folded proteins, and then further on to traits such as camouflage, blue eyes, penicillin resistance (in bacteria), and musical ability. All will agree that genes have some causal role in even the most complex traits. But as I have argued in earlier sections, the question of what is coded for by genes is not the same as the question as what is caused by genes.

One common view in this area is a permissive one, seen in Dawkins (1982) and Sterelny and Kitcher (1988). According to this permissive view, the

standard concept of "gene for X" recognizes no natural boundary beyond which it is false to associate effects with particular genes. If a gene has a systematic association with a trait that is complex and far removed causally from particular proteins -- even if it is manifested in another organism from the one bearing the gene, as in some of Dawkins' favorite cases -- still there is no problem in principle with saying that the gene is a gene "for" that trait or effect.

I suggest that if coding is taken seriously then this permissive view cannot be the whole story. If coding is a real relation linking genes and their effects, then not every causal consequence of a gene will be coded for. There is no indefinitely extended phenotype in the sense involving coding, even if there is in the sense involving causing.[21]

To see the point, consider some everyday cases involving messages with imperative contents. Here we certainly do not see all the effects of a message, even the reliable and systematic ones, as necessarily coded for by the message. Suppose I know that if I order the extra-large pizza, that will have the consequence that the delivery arrives late. This fact does not imply that when I order the pizza I am also ordering them to make the delivery late. The likely effects of a message, even an imperative message, are not all part of the content of the message. Not everything caused by a message is coded for.

So the question of how "distal" the content of genetic instructions might be has to be settled by specific analyses of coding. Suppose a gene produces an amino acid which folds to produce an enzyme that catalyzes a reaction that produces a pigment that makes the organism camouflaged from its predators. The amino acid is the most proximal of these effects of the gene, while camouflage is the most distal. Here there are no fewer than four possible degrees of distality that the content of the DNA might have. Even on the assumption that there is a simple causal chain from amino acid sequence to enzyme to pigment to camouflage, there is a further question which of these products is coded for. (The problem obviously gets harder in the case of traits that involve large numbers of genes and many non-genetic causes.) Different theories of coding will answer this question in different ways; in each case there are constraints on the possible contents of genetic messages deriving from factors used to analyze those messages.

In many cases the theories I have considered will apparently preclude assigning very distal contents to genetic messages. Theories relying on the developmental role of DNA, in particular, will have this consequence. If what

makes DNA a message is the fact that it is read, along with the facts concerning the specificity of base triplets to amino acids, then apparently the only thing DNA can code for is the sequence of amino acids (the primary structure of a protein). Not even the folded state of the protein is coded for, even in cases (if there are such) where the pattern of folding is fully determined by the amino acid sequence.

Within the family of theories of coding that make use of evolutionary history, there is not so definite a verdict. If an analysis of coding gives an important role to the "consumer" of the message, as in a Millikan-style theory, then the content of the code is constrained by the possible scope of the biological functions of the consumer. If the consumer of the genetic message is the ribosomal/tRNA machinery, then the genetic message can only have as a content something that this machinery can have the function of bringing about. Probably then, the content of the message is no more distal than instructing the production of a protein. (This claim might be contested, I realize, by people with ambitious views about functions.) At the end of the section on evolutionary theories of the genetic code I discussed the possibility of a theory that stays closer to the outlines of Sterelny, Smith and Dickison's 1996 view, and does not try to accommodate all intuition-driven counter-examples. I suggest that these views, too, will tend to have the consequence that the content of genetic instructions can be no more distal than the production of a protein. But here again, there might be other views possible.

Biologists, incidentally, do not exhibit consensus on this issue of distality. Some specifically restrict coding to the specification of the amino acid sequence (Sarkar quotes Crick saying this, for example: 1996 p. 858). In contrast, recall this sentence from the Lodish et al. textbook: "[DNA] contains a coded representation of all of the cell's proteins; other molecules like sugars and fats are made by proteins, so their structures are <u>indirectly</u> <u>coded</u> <u>in</u> <u>DNA</u>." (1995 p. 10, emphasis added).

The question of distality is made more complex by the role of the common phrase "gene for...." Does this phrase imply a relation involving coding, or just some relation of causal or statistical relevance? In earlier sections I discussed cases which suggest that there can be a useful concept of "gene for X" which does not imply a coding relation, and I think that many will agree that <u>if</u> coding is to be viewed as a real part of the theoretical structure of molecular biology, then it

will be useful to also have a sense of "gene for X" which does not imply that the gene codes for X.

It might even be useful to recognize several different senses of "gene for X" that do not imply coding. Along with the coding sense, there may be distinct statistical, causal and teleofunctional senses. In the statistical sense, a gene is a gene for X if it displays a certain pattern of correlation with X. A causal sense of "gene for X" will require that the connection between the gene and the trait have the right causal properties (there is likely to be much dispute about which properties these are).[22] And in the teleofunctional sense, a gene will be a gene "for" X if it has been maintained under natural selection because of its association with trait X.

I suggest that those who take coding seriously should want to recognize at least one of these other "gene for" concepts, and perhaps will find a use for all three.

If biologists recognize genetic coding along with one or more concepts of "gene for X" that do not imply coding, there will be a range of cases where we have a gene for X that does not code for X. For example, on a "gene for X" concept like Sterelny and Kitcher's, it is straightforward to have a gene for camouflage, penicillin resistance or even reading. A sense of "gene for X" that requires tighter causal connections might recognize genes for camouflage and penicillin resistance but not for reading -- there are various ways the details of such views could be developed. But most reasonable theories of genetic coding will probably not hold that there can be genes coding for camouflage or penicillin resistance, and almost certainly not for reading.

So many views will recognize genes for X that do not code for X. Is the converse possible? Could there be a gene that codes for X that is not a gene for X? On some views this is a possibility. I discussed a semi-hypothetical case in Section 5 where a gene is systematically misinterpreted by translation mechanisms, but succeeds under selection despite this. So we have a causal, statistical and even naturally selected association between a gene and a protein, where the standard interpretation of the gene does not associate that gene's sequence with the protein's primary sequence. The gene does not code for the protein, on any theory of coding that gives a central place to the standard rule linking base triplets with amino acids, but in every other sense the gene is associated with that protein.[23]

A central topic of this paper has been an asymmetry within standard views in biology: while both genes and environmental conditions have causal effects on development, only genes code for (some of) their effects. My various proposals for analyzing genetic coding have all been developed with this constraint in mind. What is the status of the three alternative concepts of "gene for X" with respect to this asymmetry? The statistical and causal concepts of a "gene for X" will clearly be as applicable in principle to environmental conditions and non-genetic inherited factors as to genes -- here I agree with Griffiths and Gray (1994). The status of the teleonomic sense is less clear; whether there can be "environmental conditions for X" in this sense will depend on the details of the analysis of evolutionary processes. On standard views, environmental conditions are not shaped by selection in the same sense that genes are, and there will be no "environmental conditions for X" in the teleonomic sense. But the "extended replicator" view developed by Sterelny, Smith and Dickison (1996) would certainly treat genes and environments symmetrically in this respect. Other unorthodox views might do the same.

To close this section I will introduce yet another concept of "gene for X," as it has a combination of features that some might find useful. Suppose it is accepted that genes can only code for the primary structure of proteins. Is it possible to describe a concept of "gene for X" that (i) includes coding (ii) is hence restricted to genes and not environments, but (iii) is more distal in the values it allows for "X"? One way would be to say that a gene is a gene for X if it codes for a protein that causes the distal trait. To pick a simple case, if a gene in a bacterium codes for an enzyme that causes penicillin resistance, that is a gene for penicillin resistance. This also provides a way to make some sense of the quote from Lodish et al. that I gave above: "... other molecules like sugars and fats are made by proteins, so their structures are indirectly coded in DNA" (1995 p. 10).

This last concept might be helpful to orthodox views, but I stress that if it is true that genes cannot code for anything other than amino acid sequences, that is an important fact which should be highlighted more than it often is by biologists. The importance of combating the mistaken idea that "genes are destiny" requires that the concept of coding, with which so many errors can be made, be kept strictly in its place.

7. Closing Summary

I have no firm conclusions to draw but here is a summary of some of my main points.

(i) A central feature of the coding problem is the status of the claim that while genes, environmental conditions and other factors can all causally affect traits, genes are distinguished from other factors because only they code for some of their effects.
(ii) The existence of a standard rule mapping amino acids to RNA and DNA base triplets does not solve the coding question. This rule could be seen as describing a set of causal specificities, without giving a rule of interpretation.
(iii) If DNA sequences have semantic content, this content is imperative rather than indicative.
(iv) An appeal to the concept of information, as understood in the mathematical theory of information, is unlikely to solve the problem, especially as such an approach is ill-equipped to solve the asymmetry problem described in (i) above.
(v) Two approaches to the problem of coding are analyses based on the developmental role of DNA and analyses based on evolutionary history.
(vi) Views based on developmental role will fight their main battles over the concept of "reading." They may also fight philosophical battles over the distinction between reading and misreading. If the idea that the cell reads DNA sequences can be defended, a satisfactory analysis might be possible by combining developmental role with a minimal appeal to a historical concept of function, to deal with the problem of misreading.
(vii) Within the evolutionary approach, the analysis of Sterelny, Smith and Dickison (1996) has problems with counter-examples, because it casts its net too widely. One way to avoid counterexamples is to move closer to Millikan's view in philosophy of mind, but this risks forcing an empirically unmotivated framework onto the biology.
(viii) All the genes can code for, if they code for anything, is the primary structure (amino acid sequence) of a protein.
(ix) If coding is taken seriously, there are good reasons to recognize one or more concepts of "gene for X" that do not imply that the gene codes for X. These concepts of "gene for X" might be analyzed in statistical, causal or teleofunctional ways.

There are clearly some things genes can do that environmental conditions cannot do -- act as a template in the constuction of amino acid chains, most notably. But despite the enthusiasm of biologists, whether or not this role is best understood in terms of a coding relationship is a harder, and to my mind unresolved, issue.

Notes

Acknowledgment: Much of this paper emerged from discussions with Richard Francis. Thanks also to Lori Gruen, Philip Kitcher, Susan Oyama, Kim Sterelny and Kritika Yegnashankaran for helpful correspondence and discussion.

1 For various kinds of dissent and unease, see Oyama (1985), Lewontin (1991), Moss (1992), Sarkar (1996), Griesemer (forthcoming), and Francis (forthcoming). In this paper I will use the language of coding in standard ways when I discuss examples.

2 This is a move whose importance has been stressed to me by Kritika Yegnashankaran.

3 The literature is large. Landmarks include Dretske (1981), Millikan (1984) and Fodor (1987). See Stich and Warfield (1990) for a collection of key papers and Sterelny (1990) for a review of the options.

4 I say "exactly" but there are exceptions in particular cases -- see below in section 5.

5 So I see this sentence, from one of the textbooks cited earlier, as making a very strange claim, even within a strongly symbolic view of molecular biology: "The synthesis of protein is known as translation because it involves the transfer of information from one language (nucleotides) to another (amino acids)" (Raven, Evert and Eichhorn 1992 p. 144). I see claims like this as some evidence for a view like that of Kitcher, who holds that talk of coding makes no real contribution to molecular biology.

6 They add a causal constraint as well, but in a somewhat indirect way. An allele A is a gene for X if it is at a locus that causally affects X, and individuals with A (in normal genetic and non-genetic environments) have X. (Perhaps it would be better to build the causal constraint into the relation between A and X, or not have it at all.) Sterelny and Kitcher also stress that there are several different ways of handling the fact that a gene's role is dependent on its environment -- they do not think there is just one way to understand "normal environment."

7 See Oyama 1985 Chapter 5 for a menagerie of such phrases and formulas.

[8] For "prescribing" see Lodish et al. (1995) p. 101, for "dictating" see the quote from Raven et al. at the start of Section 2.

[9] Some analyses require that the correlation be based in natural law (Dretske 1981).

[10] For one attempt, see Maclaurin (forthcoming). In my view, proposal has problems deriving from his not taking the content of DNA to be imperative.

[11] Griffiths and Gray (1994 p. 281) give an interesting quote from Konrad Lorenz in which Lorenz says both that genes (i) contain a blueprint, and (ii) give the organism descriptive information about its environment ("a rival is red underneath"). I will not tackle this second kind of content-attribution here. Certainly not all claims about genetic coding can be understood as providing environmental information. I'm not sure how many could be handled this way.

[12] For most of this discussion I will just say "developmental," not "developmental and metabolic," because the philosophical literature on this point is mostly concerned with development.

[13] For a discussion of some of these phenomena which draws interesting philosophical conclusions, see Neumann-Held (unpublished).

[14] I say "might reasonably" because Sarkar's overall view is that the concept of coding plays little positive role in molecular biology, and while it is much less problematic in the case of prokaryotes, on balance we might be better off without the concept.

[15] See Fodor (1984) for an influential discussion.

[16] For a discussion of the exact relations between functions and evolutionary history, see Godfrey-Smith (1994). Dretske's view (1988) is the best representative of theories which use a biological concept of function for part of the analysis but use other concepts as well.

[17] Sarkar also discusses an interesting case. The RNA coding for part of an enzyme (NADH dehydrogenase subunit 7) in a parasite (Trypanosoma brucei) has hundreds of "U" bases inserted and some deleted before it is translated into protein. So there is no piece of DNA (contiguous or not) whose sequence corresponds, by the standard rule, to the amino acid sequence of the finished protein. Consequently, Sarkar says, "the DNA segment encoding the primary transcript can hardly be considered a gene for NADH dehydrogenase subunit 7" (1996 p. 861, emphasis added). I prefer to say this is a gene for NADH dehydrogenase subunit 7, but not a gene that codes for it.

[18] On frame shifting see Watson et al. (1987) p. 458, Lewin (1997) pp. 237-49.

[19] Sterelny, Smith and Dickison's main discussion of these issues is around pp. 387-89 of their paper. For example: "One element of the developmental matrix exists only because of its role in the production of the plant lineage phenotype. That is why it has the function of producing that phenotype, and hence why it represents that phenotype." (1996 p. 388)

[20] I use the term "representation" here although Millikan (1984) calls these signs "icons," and reserves "representation" for a richer concept.

[21] Related to this is the fact that causal relations are transitive in a way that coding relations are not. If A codes for outcome B and B codes for outcome C, that does not imply that A codes for C.

[22] For a good discussion of the differences between some of the relevant causal and statistical concepts in this area, see Block (1995), and a classic discussion in Lewontin (1974).

[23] Some might say that even the presence of introns generates this consequence, though I do not agree with that. Then there is the case discussed by Sarkar (note 17 above), and cases of frame shifting. Whether Sarkar's case fits might depend on the details of how the editing process is caused.

References

Block, N. (1986). "Advertisement for a Semantics for Psychology", reprinted in Stich and Warfield (1994).

Block, N. (1995). "How Heritability Misleads About Race", Cognition 56: 99-128.

Dawkins, R. (1982). The Extended Phenotype. Oxford: Oxford University Press.

Dretske, F. (1981). Knowledge and the Flow of Information. Cambridge MA: MIT Press.

Dretske, F. (1988). Explaining Behavior. Cambridge MA: MIT Press.

Fodor, J. A. (1984). "Semantics, Wisconsin Style", reprinted in Stich and Warfield (1994).

Fodor, J. A. (1987). Psychosemantics. Cambridge MA: MIT Press.

Francis, R. (forthcoming). Genes, Brains, and Sex in the Information Age.

Godfrey-Smith, P. (1994). "A Modern History Theory of Functions". Noûs 28: 344-362.

Griffiths, P. and R. Gray (1994). "Developmental Systems and Evolutionary Explanation", Journal of Philosophy 91: 277-304.

Griesemer, J. (forthcoming). "The Informational Gene and the Substantial Body: On the Generalization of Evolutionary Theory by Abstraction", in N. Cartwright and M. Jones (eds.), Varieties of Idealization. Amsterdam: Rodopi.

Kitcher, P. S. (forthcoming). "Battling the Undead: How (and How Not) to Resist Genetic Determinism". To appear in R. Singh, C. Krimbas, D. Paul and J.

Beatty (eds.), Thinking About Evolution: Historical, Philosophical and Political Perspectives. Cambridge University Press.

Lewin, B. (1997). Genes VI. Oxford: Oxford University Press.

Maclaurin, J. (forthcoming). "Reinventing Molecular Weismannism: Information in Evolution". To appear in Biology and Philosophy.

Lewontin, R. C. (1974). "The Analysis of Variance and the Analysis of Cause", reprinted in R. Levins and R. C. Lewontin The Dialectical Biologist. Cambridge MA: Harvard University Press, 1985.

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

Lodish, H., D. Baltimore, A. Berk, S. L. Zipursky, P. Matsudaira and J. Darnell (1995). Molecular Cell Biology, 3rd edition. New York: Freeman.

Millikan, R. G. (1984). Language, Thought and Other Biological Categories. Cambridge MA: MIT Press.

Moss, L. (1992). "A Kernel of Truth? On the Reality of the Genetic Program", PSA 1992, Volume 1: 335-348.

Neumann-Held, E. (unpublished). "Lets De-BlackBox the Gene!", presented at ISHSSPB Conference, Seattle 1997.

Oyama, S. (1985). The Ontogeny of Information. Cambridge: Cambridge University Press.

Papineau, D. (1993). Philosophical Naturalism. London: Blackwell.

Raven, P. H., R. F. Evert, and S. E. Eichhorn (1992). Biology of Plants, 5th edition. New York: Worth.

Sarkar, S. (1996). "Decoding "Coding" -- Information and DNA", BioScience 46: 857-864.

Shannon, C. E. (1948). "A Mathematical Theory of Communication", Bell System Technical Journal 27: 379-423, 623-656.

Sterelny, K. (1990). The Representational Theory of Mind: An Introduction. London: Blackwell.

Sterelny, K. and P. S. Kitcher (1988). "The Return of the Gene", Journal of Philosophy 85: 339-61.

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

Stich, S. P. and T. A. Warfield (eds.) (1994). Mental Representation: A Reader. Oxford: Blackwell.

Watson, J., N. Hopkins, J. Roberts, J. A. Steitz and A. Weiner (1987). The Molecular Biology of the Gene, 4th edition. Menlo Park: Benjamin/Cummins.