CHAPTER FOURTEEN

Genes do not Encode Information for Phenotypic Traits

Peter Godfrey-Smith

14.1 Introduction

How can a philosopher sit here and deny that genes encode information for phenotypic traits? Surely this is an empirical, factual, scientific matter, and one that biology has investigated and answered? Surely I must realize that biology textbooks now routinely talk about genes encoding information? Isn't the discovery of the "genetic code" one of the triumphs of twentieth-century science? Indeed, doesn't much of the research done by biologists now focus on working out *which* genes encode information for *which* phenotypic traits? So is this one of those unfortunate cases of philosophers, in their comfortable armchairs, trying to dictate science to the scientists?

I can understand if that is your initial reaction to my claim. I hope to persuade you, however, that it is indeed a mistake to say that genes encode information for phenotypic traits. My aim is not to overturn empirical results or deny core elements of the textbook description of how genes work. Instead, I will be trying to argue that if we look closely at the details of what biology has learned, it becomes clear that it is a mistake to talk about genes encoding phenotypic traits. This common way of talking is not, as it seems to be, a straightforward and concise summary of how genes work. Instead, it involves a distortion of what biologists have actually learned.

My denial that genes encode information for phenotypic traits is based, as you might expect, on a particular way of interpreting some key terms. My argument is that *if* we understand "encode information" and "phenotypic traits" in a particular way, then genes do not encode information for phenotypic traits. I do think that my interpretations of these terms are reasonable ones, but I will also discuss other interpretations of them.

More specifically, my claims in this chapter rest on two distinctions. One is a distinction between two senses of "information." This distinction will take some time to

outline but, roughly speaking, there are two ways of using the language of "information" to describe events and processes in the world. One way is relatively harmless and uncontroversial. In this sense, any process at all in which there is a reliable correlation between two states can be described in terms of information. This is the sense in which dark clouds carry information about bad weather, and tree rings carry information about the age of the tree. The second sense of "information" is more controversial and problematic. Here, when we speak of information we mean that one thing functions in some way as a representation of another, and has a definite semantic content. Something more than a mere correlation or ordinary causal connection is involved. The rings in a tree are not functioning as representations or signals in the life of the tree; they are just indicators that happen to be useful to observers like us. The words on this page, in contrast, make use of a conventional symbol system to convey particular meanings.

The second crucial distinction that I will make is less complicated and less controversial. What counts as a "phenotypic trait" for the purposes of this discussion? Clearly, the properties of whole organisms, such as size, shape, and behavioral patterns, count as part of the phenotype. But might we also include the more immediate causal products of particular genes, which are mRNA molecules and protein molecules? In one very broad sense, the phenotype might be seen as including all of the physical properties of an organism other than its genetic composition. In this sense, even individual proteins are part of the phenotype. But that is *not* the sense that I will use in this chapter. Instead, when I say "phenotypic traits" I mean the familiar properties of whole organisms studied by biologists – size, shape, color, structure, behavior patterns, and so on.

Now I can say something about the sense in which I deny that genes encode information for phenotypic traits. It is certainly possible to describe the connection between genes and phenotypic traits in informational terms, in the first of the two senses of "information" above. This applies however one interprets the term "phenotype." But now consider the second sense of "information," the semantic sense. Do genes have a role in the production of the phenotype that can be described in terms of messages, representations, meaning, or codes?

Clearly, genes are not sending deliberate signals, and they do not use a representation system maintained by social convention, like the words on this page. But we should not get too hung up on these issues. We can ask a different, and more difficult question: Do genes operate in the production of phenotypes in a way that involves some kind of special process that is at least *strongly analogous* to familiar cases of representation? Might there be a useful and general sense of "representation" or "instruction" in which both genes and human sentences play a similar and special kind of role? In answering this question, a lot hangs on how we understand "phenotype." If we were to understand individual protein molecules as part of the "phenotype" of an organism, then I think a good case can be made for the genes having a special informational role. But that is not the ordinary sense of "phenotype." If we ask whether genes have a special informational role in the production of the phenotype in the ordinary sense – the structure and behavior of a whole organism – then the answer is "No."

14.2 More on the Two Senses of Information

Let us look more closely at the concept (or concepts) of information.

The first sense of information outlined above is basically the sense used by Shannon in the original formulations of information theory (Shannon, 1948). I can be fairly brief here, because this topic is covered in more detail in Sarkar's companion chapter 13. Information, in Shannon's sense, exists just about everywhere. A "source" of information is anything that has a range of possible states (such as the weather, or the age of a tree). And whenever the reduction of uncertainty at one place is correlated to some extent with the reduction of uncertainty at another, we have a "signal" that is carrying information. The sky, for example, can have a range of different states (cloudy, clear, and so on). The weather two hours from now also has a range of different states. If the state of the sky now is correlated (perhaps not perfectly, but at least to some extent) with the future state of the weather, then the sky carries information about the future weather.

In the literature on genes and information, this is sometimes called the "causal" sense of information. Or people say that one thing carries "causal information" about another (Griffiths, 2001). This terminology is fine, although causation *per se* is not what is required here. What is required is a correlation (Dretske, 1981). So my preference is to talk of this first sense of information as information in the *Shannon* sense, after the main inventor of information theory, Claude Shannon.

Information, in this sense, is everywhere. It does connect genes with phenotypic traits, but it connects them in both directions and connects both of these with environmental conditions as well. Here it is important to remember that all phenotypic traits are causally influenced by both genes and by environmental conditions. In some cases, a trait might be much *more* sensitive to one or the other of these – to slight differences in the genes, or in environmental circumstances. But that does not change the fact that all traits arise from an interaction, of some kind, between genetic and environmental conditions.

So let us look at how we might use informational language of this kind to describe genes. Whenever we describe any process using this sort of informational language, we treat some factors as background conditions or "channel conditions." If we want to treat genes as carrying information about phenotypic traits, we will regard the role of the environment as a mere background condition. Then if the presence of gene G correlates reliably with the appearance of phenotype X, we can say that gene G carries information about the phenotype, and carries the information that X will be present. But we can also turn the situation around; we can treat genes as background, and look at the correlations between different environmental conditions and phenotypic traits. From this point of view, it is environmental conditions that carry information, not the genes (Oyama, 1985; Griffiths and Gray, 1994). And informational relationships run "backward" as well as "forward" in these cases. The phenotype carries information about the genes, just as the genes carry information about the phenotype.

As the reader will see, I am trying to suggest that when someone claims that genes carry information about phenotypes in the Shannon sense, there is very little being

claimed. In this sense of information, it is not a novel and important part of biological theory that genes and phenotypes have an informational connection between them. But I do not want to suggest that this sense of information is useless or unimportant in biology. On the contrary, Shannon's concept of information provides a very useful and concise way of describing masses of physical associations and connections, especially because information in the Shannon sense can be described mathematically. A lot of discussion in contemporary biology is facilitated by this conceptual framework. For example, biologists often talk about the information that genetic variation within and between populations carries about evolutionary relationships. It turns out that "junk" DNA carries more of this information than does DNA that is used to make proteins. So informational description of this kind has definite utility. But, obviously, in cases such as these, the information in the genes is something that only *we*, the observers and describers of genetic systems, use. This information is not part of any explanation of the biological role that the *genes* play within organisms. Here, genes are being used by us just as we might use tree rings.

Some of the talk by biologists about how genes carry information about phenotypes can be understood as using the Shannon sense of information only. But not all of it can be interpreted in this way. For example, biologists are often (though not always) reluctant to say that the environment contains information that specifies phenotypes, in the same sense as genes. And the word "code" is important here; there is not an environmental code in the way there is a genetic code. Biologists often treat genes as having an informational role in a stronger sense than the Shannon sense.

So let us look at the other sense of "information." I should say immediately here that the word "sense" is a bit misleading, because it suggests there is a *single* additional sense of the word "information." Some might think that this is true, but I do not. When we leave the precise Shannon sense of information, we encounter an unruly collection of different concepts. We encounter the large and controversial domain of *semantic* properties – properties that involve meaning, representation, reference, truth, coding, and so on. Despite a massive effort by philosophers and others over many years (especially the past 100 years), I think we do not have a very good handle on this set of phenomena.

The two domains where the idea of meaning and representation seem most familiar to us are the domains of *language* and *thought*. Philosophy of language and philosophy of mind have both seen attempts to formulate general theories of meaning and representation, that might be used to help us describe semantic phenomena in other areas. (For theories in the philosophy of mind, see Stich and Warfield (1994); for philosophy of language, see Devitt and Sterelny (1999).) People who are trying to distinguish genuine semantic properties from information in the Shannon sense often point to the capacity for error, and the ability to represent nonactual situations, as marks of semantic phenomena. In the simple clouds-and-rain case, there is no sense in which the clouds could *misrepresent* the weather. The correlation between clouds and rain might fail to hold in some particular case, but that does not imply that the clouds said something *false*. If I tell you a lie in this chapter, however, my words have indeed been used to say something false. I can also use these words to describe a situation that I know does not obtain, such as my having won the lottery. These are not features of information in the mere Shannon sense. Despite much effort, philosophers

have not made much headway in giving a general theory of semantic properties. (Some would disagree, but that's how it seems to me.)

Biologists began to use the terminology of information and coding back in the middle of the twentieth century (for the history of this trend, see Kay, 2000). Biologists do not all have the same thing in mind when they talk of genes carrying information. Some see it as a mere metaphor, or as loose and picturesque talk. Others mean no more than information in the Shannon sense. But in recent years, the idea of genetic information has been taken more and more seriously by some biologists (Williams, 1992; Maynard Smith, 2000), and these biologists have clearly had more than the simple Shannon sense of information in mind.

So do genes encode information for phenotypic traits in something more than the Shannon sense? Answering this question is made awkward by the absence of a good philosophical theory of semantic properties, the absence of a good *test* that we could apply to genes. Most philosophers will agree that we do not have a clear diagnostic question to ask. Here is the method that I will apply in this chapter. My approach will be to see whether there are properties of genes that have an important analogy with familiar and central cases of languages and symbol systems. The test is not so much whether these analogies *look* striking, but whether the analogies are *important within biological theory*. To what extent does describing genes in informational terms help us to understand how genetic mechanisms work? That will be how I approach the question addressed by my and Sarkar's chapters.

14.3 What Genes Do and How They Do It

My argument about genes and information will be based on some detailed facts about how genes exert their influence on the phenotype. My argument is not directed against the idea that genes have an important *causal* role in producing the phenotypes of organisms; no one can deny that. My argument concerns *how* they exercise that causal role. I claim that this causal role does not involve the interpretation of an encoded message in which genes specify phenotypic traits. So let us look in some detail at what genes do and how they do it.

What genes do, fundamentally, is make protein molecules, make a few other molecules, and contribute to the regulation of genes in these same activities. And it is not really very accurate to say that genes "make" proteins. Genes act as templates that specify a particular crucial feature of each protein molecule – the linear order of the protein's building blocks, which are individual amino acids.

Genes are made of DNA, which is, of course, a long two-stranded molecule arranged as a double helix. Each strand of DNA is a chain of four building blocks (the "bases" *C*, *A*, *G*, and *T*), held together by a backbone. Two main steps are distinguished in the causal chain between DNA and a finished protein molecule. "Transcription" is the process in which DNA gives rise to mRNA ("messenger RNA"). Then "translation" generates the protein itself, using the mRNA.

mRNA is a similar molecule to DNA – another chain of four building blocks, but single-stranded rather than double-stranded. 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 that 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 before it is used in translation. Then, elsewhere in the cell (at the ribosomes), the mRNA is used to direct the formation of another long-chained molecule. This time, the chained molecule is 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"). These are shorter chains of the RNA building blocks, which also bind to single amino acids. Different tRNAs bind to different amino acids (of which there are 20 kinds used in protein synthesis). At the ribosomes, tRNA molecules carrying amino acids with them bind temporarily to specific three-base sequences in the mRNA molecule. So each triplet of bases in the mRNA is associated, via the special chemical properties of tRNA molecules, with a particular amino acid.

When a protein is assembled at a ribosome, a chain of amino acids is produced whose sequence corresponds, by a near-universal rule, to the sequence of bases in the mRNA. "The genetic code" is, strictly speaking, this rule connecting RNA base triplets with amino acids. This interpretation of the RNA determines the interpretation of the DNA from which the mRNA is derived. 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 six different triplets. The three remaining triplets are "stop" signals. The chain of amino acids then 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 linear sequence of amino acids. The tertiary structure is the three-dimensional folded shape of a single amino-acid chain. The causal role of a protein within an organism depends greatly on its tertiary structure.

There is much more to all these processes, of course. But this outline describes the core of the process by which any genetic message is expressed. I should also add a quick remark about how genes operate in the *regulation* of these same processes. Generalizing greatly, the main way this occurs is by genes producing proteins that physically interact with particular stretches of DNA itself, in a way that either impedes or facilitates the process of transcription. So the process of protein synthesis feeds back on itself in complex ways. There are also other mechanisms by which genes can regulate the activities and products of the genetic system.

In this sketch, I have used standard terminology from genetics, including some standard terminology that seems to be asserting an important role for something like information. We have "transcription," "translation," and a "code" linking DNA/RNA sequences to amino-acid sequences. And I do not object, or at least do not object much, to this collection of terminologies (for a few minor complaints, see Godfrey-Smith, 1999). There are indeed some very special features of the role of genes within protein synthesis, features that seem to me to justify, to some extent, a description in terms of codes, messages, and so on. I will focus on these in the next section. The point to make here is that everything I have said concerns the connection between a gene and the protein molecules that it produces. Nothing has been said yet about the "downstream" effects of the protein itself. And here, a rather different story must be told.

There is not a single, unified biological process by which proteins have their effects on the organism, in the way in which there is a fairly unified process by which genes give rise to proteins. Instead, proteins have a host of different effects, and a host of different ways of having these effects. Some proteins operate as parts of the physical structure of the organism; others act as enzymes controlling key reactions between other chemicals; others act as hormones; and so on. Proteins are the fundamental "working" molecules in all organisms, but their work involves a tremendously complex mass of cascading and interacting processes. The actions of proteins are part of a huge and varied causal network, some parts of which we know a lot about, and other parts of which we know very little about. Being blunt about it, once a gene has produced a protein molecule, the protein generally goes off and enters a kind of causal soup, or causal spaghetti, where it might affect, and be affected by, a great range of other proteins and other factors, including environmental factors.

Given this picture, how should we describe the causal connection between a gene and a particular aspect of the whole-organism phenotype? It depends very much on the particular case. In some cases, despite all the intervening complexity, a trait can be under rather tight causal control by specific genes. In any environment in which the organism can develop at all, if it has gene G (and a reasonably normal set of genes elsewhere in its genome), it will develop with phenotype X. Slight environmental differences have little effect on the phenotypic outcome. (Traits of this kind are sometimes called "canalized," although that term has several related uses.) In other cases, phenotypic traits have much more sensitivity to particular environmental conditions. Many writers like to describe this second set of cases by saying that the genes here determine a *range* of possible values for a trait, and the environment fills in the *specific* value. I have always been a bit suspicious about this formulation, although it does seem to be helpful sometimes. In general, I think that we have not really developed much of the causal language that we need for describing the relations between genes and phenotypic traits. The most everyday causal language, of the "this causes that" type, is often very misleading, especially because it suggests an overly simple picture in most cases. There are various more specialized forms of language used in genetics, including statistical talk about the "percentage of phenotypic variation explained by genetic variation," and causal talk about the "norms of reaction" of particular genotypes (Lewontin, 1974). Finding the right framework for the causal description of genes is a difficult and interesting business, but it is not the task that we confront here. The task here is the assessment of a particular way of describing what genes do, a way using the language of information and coding.

14.4 The Real but Restricted Role of "Genetic Coding"

I said earlier that philosophers have not succeeded in providing a good analysis of meaning, representation, and other semantic properties. Biologists, however, have become enthusiastic about the description of genes and gene action in term of coding and information. How should we respond to this situation? Should we just accept the biological description uncritically, especially given the shortcomings of philosophical analysis in this area? Or should we suspect that biologists have become careless (or

worse) in their enthusiasm for "the informational gene"? Here is the approach that I will take to the problem. I ask: Which attributions of coding and/or informational properties to genes have a *useful theoretical role* within biology? I will argue in this section that there is a real, but restricted, domain in which the attribution of coding properties to genes does help us to solve problems and understand how organisms work. This domain is the explanation of protein synthesis itself.

The picture sketched in the previous section was roughly like this: genes make proteins, and proteins go off and enter a kind of causal spaghetti in which they interact with all sorts of things and may have a variety of roles (including feedback on the operation of the genetic system itself). To fully understand how organisms work, we need to understand all parts of this process. And to some extent, different concepts might be useful in understanding different parts of the whole. Let us focus again on protein synthesis. What is being explained here is how the cell makes proteins, where proteins are long chains of elements that must be put into exactly the right order. How to cells do this? They use the linear order of bases in their DNA to specify the order of amino acids in a protein, and this is done via a two-step *templating* operation. In a way, the concept of a template, or "template surface," is *the* crucial one for understanding how cells use genes to specify the primary structure of proteins (Watson et al., 1987, ch. 3). But the nucleic acids (DNA and RNA) act as a very specific *kind* of template in protein synthesis. Three properties make the nucleic acid templates special.

First, the templates used in making proteins are not proteins themselves, but another kind of molecule. DNA templates for mRNA which, after processing, templates for the protein. Given that the building blocks of a nucleic acid sequence (DNA and RNA bases) are not identical to the building blocks of the protein (amino acids), there must be a rule of specificity linking the two. In actual organisms, this rule is largely (though not completely) fixed and invariant. So a great variety of proteins are constructed from different DNA templates, in the same and in different organisms, via the same compact and general rule.

Secondly, the specification of proteins by these templates is *combinatorially* structured; the rule has a definite part-to-whole organization, involving the free rearrangement of the same fixed components. This structure exists at two levels. Most obviously, the elements of the templates that are specific for particular amino acids are triplets of nucleic acid bases. And, in addition, during translation a given triplet specifies the same amino acid regardless of its neighboring triplets. (There are some exceptions to this principle, but it holds in general.) The interpretation of a long sequence of bases is (again, in general) a simple and fixed function of the interpretation of its component triplets.

Thirdly, the rule linking base triplets with amino acids is believed to be largely "arbitrary." Some controversy surrounds this point, and by "arbitrary" I mean something specific. I mean that nothing about the chemistry of a particular amino acid is responsible for it corresponding to a particular base triplet. Contingent features of the tRNA molecules, and the enzymes that attach the amino acids to tRNAs, determine which triplets go with which amino acids. If the tRNA molecules had different sequences, or if the enzymes that work with them operated differently, a given base triplet could have an entirely different interpretation in protein synthesis. This need

not have been true during the earliest evolution of the genetic code, but given the way protein synthesis works now, there is a kind of arbitrariness in the system.

So, my first main claim in this section is that the peculiar way in which proteins are made does justify talk of genes as "coding." This way of describing genes is, indeed, picking out some real and distinctive features of genetic mechanisms. Genes help make proteins by acting as templates, and this involves a combinatorial and arbitrary rule, largely static and universal, which connects nucleic acid sequences with amino-acid sequences. The analogy with human symbol systems such as languages is quite strong. So it is reasonable to say that genes encode information specifying amino-acid sequences in proteins, where this involves more than just the minimal, Shannon sense of information.

These features were not the ones picked out by the first discussions of "coding" in genetics. Indeed, the physicist Erwin Schrödinger, who was the first to use this terminology here, had nothing more than a *predictability* relation in mind when he initially introduced the idea of a "code-script" in his book in *What is Life?* (1992, p. 21). Tellingly though, Schrödinger made it clear in a later passage in *What is Life?* that he was thinking of systems with combinatorial properties, such as Morse code (p. 61). In any case, I claim that the positive, problem-solving content in the idea of "genetic coding" is its picking out these special features of the way nucleic acids act as templates for amino-acid sequences.

So, proteins are made by having their amino-acid sequence coded for. What does this conclusion tell us about the more general issues concerning the link between genes and phenotypes? It tells us that genetic causation has peculiarities, because at one specific place *within* the causal chains linking genes and phenotypic traits, we have some very special processes and mechanisms. But these unusual features of genetic causation concern *how* genes manage to have some of their *immediate* effects. These features of genetic causation do not extend beyond the local process in which the protein in question is being made.

In recent years, however, enthusiasm for the semantic characterization of genes has been unstoppable. This has led to the extension of "coding" talk, to the point at which the concept of genetic coding is now used to describe and distinguish the *entire causal paths* in which genes are involved. Such descriptions *seem* to bear real theoretical weight; the suggestion is that among all the causal paths leading to the development of an adult organism, some of these causal paths are distinctive because they involve the expression of a genetically encoded message specifying the process. But I suggest that once we leave the context of explaining protein synthesis, semantic descriptions of genes have a very different status. These further kinds of semantic description of genes have no empirical basis and make no contribution to our understanding.

Consequently, it is a mistake to use the idea of genetic coding or genetic information to pick out a distinction between the characteristics of whole organisms: a distinction between traits that are "coded for" and traits that are not. All that can be coded for is the primary structure of a protein molecule; not even higher-level protein structure is coded for, strictly speaking. To say this is not to deny the long causal reach of genes. A protein is made by being coded for, and the protein might have a key causal role in some far more complex and "distal" phenotypic trait. But I deny

that positing a *coding* relation, as opposed to an ordinary causal relation, helps us at all in understanding the downstream consequences of the production of a protein.

In a recent paper (Godfrey-Smith, 2000), I used an elaborate thought experiment to try to make this point persuasive. I will sketch the central ideas here (though if the reader is already persuaded, it is fine to skip to the start of the next section). The thought experiment is based on the history of twentieth-century biology. The point of the thought experiment is to ask: If there was no genetic coding, which parts of biology would have to be different?

For some time in the years before Watson and Crick's work, it was thought that proteins themselves were the natural candidates for being the molecules that made up genes. One way for this to work is for genes to be little *samples* of the protein molecules needed by the cell, stored in the chromosomes. A gene would make other proteins, and would be replicated, by acting as a template *for itself*. There seemed no feasible way for amino-acid chains to template for themselves directly. So there might be a set of "connector" molecules that could bind on each side to the same amino acid. There would have to be 20 types of these connectors, for the 20 amino acids. They would enable an indirect like-with-like template mechanism.

Imagine an alternative possible biology in which this is what genes were like; genes would be little samples of protein. In such a situation, there would be no such thing as genetic coding. Instead of being a sort of *representation* of a protein, the gene would be a *sample* of the protein. So there would be no messenger, no translation, no expression, and no interpretation. Coding talk makes sense in genetics because the ordering of amino acids in proteins is done by means of *another* kind of molecule, the nucleic acids, which contain components that can be mapped with a special combinatorial rule onto the order of the amino acids in proteins.

Despite the absence of genetic coding in this situation, rather little else need be different. In the causal processes found in the development and metabolism of organisms, *one part* of every causal process would differ from our actual situation. The way in which genes place amino acids in the right order in protein synthesis would be different in each case. But the causal stories *from* the point at which each individual protein molecule has been produced could be exactly, or nearly, the same. Once a protein is produced, it makes no difference whether it was coded for or run off from a sample.

In my recent paper (Godfrey-Smith, 2000), I discuss the consequences of this thought experiment more fully. The main point is that the *overall* patterns found in the biological world need not be much different in the scenario we are envisioning, although many of the details would have to be different. The claim that I am making is that once the amino acids are placed in order, the role for genetic coding is over. After that, it is up to the proteins to do whatever they can do.

14.5 Two Objections and Replies

The previous section argued that the semantic description of genes is justified in one specific domain, but only in that domain. Basically, I broke the causal chain between a gene and its phenotypic effects into two parts. The first part connects a gene with

its protein products, and the second part connects the proteins with their myriad "downstream" effects on the organism's phenotype. The term "downstream" here does involve a simplification, because of the role of some proteins in regulating gene action, but I do not think this simplification affects my main points. I argued that genes do code for proteins, but this does not imply that they code for the proteins' further effects on the phenotype.

But what reason is there to insist that the "coding for . . ." relation only extends as far as I say it does? This is the first objection that I will discuss. In the case of *causal* description, at least of some kinds, we find something different. If A was a cause of B, and B was a cause of C, then A was also a cause of C. This principle is not completely uncontroversial, and it only applies to some kind of causal description. But it does seem to apply for much of the time. The usual term for this is "transitivity." Any relation R is transitive if: given that A has R to B, and B has R to C, then A must also have R to C. Causation is often said to be a transitive relation.

The situation we have in the case of genes, as I have described it so far, is that a gene codes for a protein and the protein goes off and affects various phenotypic traits. The protein might have its most obvious effect just on one trait, but it will have some causal involvement (if only as background) with many others. Let us look at a simplified case, though. Suppose that gene G makes protein P, which then reliably causes the organism to have trait X. This, I should emphasize, is a *highly* simplified scenario. Much of the point of the preceding sections was to argue that although a gene can specify a protein's structure in a fairly direct and straightforward way (at least in many cases), the downstream causal role of a protein is a complex and interactive affair, which usually cannot be described simply in terms of the protein "causing trait X." But consider the simplified case for now. What is to stop us from saying that because gene G codes for protein P, and protein P causes trait X, gene G codes for trait X? This is not exactly the same as a question about transitivity, because we do not have the same relation between G and P, and between P and X. But it is related, obviously.

In a sense, nothing stops us from saying that G codes for X, in a case like this. This will be a new, specialized sense of the word "code." It is possible to concoct all sorts of novel ways of describing the roles that genes have, by mixing causal and semantic language together and giving explicit definitions of what one means to say. But are there any good reasons to describe things in this way?

One thing to note is that this treatment of the idea of coding would not follow the general pattern that we find in the case of other kinds of semantic properties. Here, I have in mind the semantic properties of more familiar examples of representation, such as thoughts and sentences within public languages. For example, suppose you know that if you order the extra-large pizza from your favorite pizza place, your action will have the consequence that the delivery arrives late. (Smaller pizzas tend to arrive on time.) 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. The semantic content of your message is just: *bring me an extra-large pizza*. Your order is complied with if and only if the order leads to the delivery person bringing you an extra-large pizza.

Aside from being late, your ordering of the extra-large pizza might have all sorts of further downstream effects, both systematic and accidental. Who knows what might ultimately result from the lateness of a pizza, or (to use an older example) the loss of a horseshoe nail? Still, although the causal chain starting from your pizza order can extend indefinitely, you only ordered the extra-large pizza. The semantic content of the message is specific to a certain portion of the resulting causal chain. If you order P and P causes X, that does not mean that you ordered X. You might, in some cases, *intend* X to occur as well. But the fact that X is caused by P does not imply that it was part of the message itself.

So, in the everyday cases, we do not use semantic concepts in the way outlined in my statement of the objection to my position. That does not mean that we *couldn't* use the language of genetic coding in this way if we wanted to. All sorts of novel forms of description can be introduced, by modifying existing terms. But is there any good reason to describe genes in this way? The idea of a genetic code linking nucleic acid bases and amino acids is part of a theoretical framework that solved a real problem in biology. That problem is restricted, however, to the domain of protein synthesis. We do not gain any further understanding – and we make distortion and oversimplification quite likely – when we extend the semantic description of genes beyond this domain. Genes can have a causal role that extends far beyond the production of proteins, but proteins are all that a gene can code for.

I will discuss one other possible view here, which I will also express in the form of an objection to my position. My argument has been that there is no reason to see genes as coding for anything beyond the primary structure of proteins. Once the protein has been made, the downstream effects of the protein involve a mass of ordinary causal relations. But perhaps, it might be argued, the situation looks different once we look at the relation between gene and phenotype in an *evolutionary* context.

Particular genes are favored by natural selection because of their overall effects on the organism's phenotype, even if these effects are indirect and are dependent on many other factors. For some philosophers and biologists, the key to understanding the semantic properties of genes is to approach the question from an evolutionary point of view. They have argued that if a gene has been favored by natural selection because it tends to bring about a particular phenotypic effect, X, then the gene *represents* X, or carries the information that X is to appear in the phenotype. This suggestion has been made by Sterelny, Smith and Dickison (1996) and also by Maynard Smith (2000). When natural selection favors a gene because of some particular phenotypic effect, that effect becomes something that the gene *represents*, as well as merely *causes*.

One way in which to understand this proposal is to look at the concept of *biological function*. There are many senses of the word "function," and many even within biology. But in one sense of the term, the function of a structure is the effect it has that has been responsible for its being selected for. When someone claims that the function of a bird's long tail is to attract mates, or the function of a spider's web is to catch insects, it is not just being claimed that the tail or the web tends to have this effect. Rather, these are the effects that the tail and the web are *supposed* to have. Catching insects is what the web is *for*, in an evolutionary sense (Wright, 1973; for a collection on these issues, see also Buller, 1999). A number of philosophers have

thought that there is a very close relation between functional properties, in this special evolutionary sense, and semantic properties (Millikan, 1984). For example, this sense of "function" is one that permits claims about *malfunction*, or error (see section 14.2 above). When the bird's tail attracts a hunter rather than a mate, it is not performing its function.

So, the advocates of an evolutionary approach to genetic information argue that when a gene has been favored by selection because it causes X, it acquires the *function* of causing X. This, in turn, justifies the claim that the gene instructs, encodes, represents, or carries information specifying that the phenotype is to include X. This proposal would only assign informational properties to genes that have been favored by natural selection, but that might not be a bad problem for the view.

Might this be a good way to defend the idea that genes code for more than proteins? In reply to this proposal, I want to first emphasize the fact that lots of biological structures have functions, in the evolutionary sense, without carrying information or representing anything. The spider's web does not represent prey, although it has the function of catching prey. Legs are for walking, but they do not represent walking. Something can have the function of producing a particular effect in biological processes, without representing or coding for that effect. The suggestion made by Sterelny, Smith and Dickison is based on *denying* this principle, for a special set of cases. They claim that all structures whose function is to operate within biological development, and bring about a certain effect on the phenotype, *represent* the outcomes of those developmental processes. In these cases, they argue, the connection between function and representation is especially tight. But I don't see any good reason to believe this claim. For example, consider all the other structures that operate in biological development alongside genes, and are part of the machinery by which genes exert their effects. Consider, for instance, the tRNA molecules and the enzymes that operate with them in translation processes. These structures have functions, in the evolutionary sense, and their role involves the same processes that genes are involved in. But they are not usually said to represent anything, or carry information about anything, and there seems no reason at all to say that they do (I discuss this evolutionary proposal about information in more detail in Godfrey-Smith, 1999).

So I think the proposal of Sterelny, Smith, and Dickison and of Maynard Smith is mistaken. Some particular genes are favored by natural selection because of their overall effects on phenotype, and can hence be described using an evolutionary concept of function. Genes *also* have a role in protein synthesis that justifies a kind of informational description, because of the special features of the "genetic code." But I emphasize the "also" in the previous sentence; these are two quite separate facts about genes. It is a mistake to put them together and say that genes encode the phenotypic traits of whole organisms.

14.6 Conclusion

I will end by quickly summarizing the main points of my argument. There are two senses of information, the Shannon sense and a stronger, "semantic" sense. Information in the Shannon sense is everywhere; it does connect genes with phenotypic traits,

but it connects them in both directions and it connects environmental conditions with phenotypic traits as well. I have no argument with the idea that genes contain Shannon-type information about the organism's phenotype, but this is a very uncontroversial claim that does not say very much.

No one has a good general theory of semantic properties, like representation and meaning. But the role of genes within protein synthesis is, in a sense, an informational or representational role. Genes specify the primary structure of protein molecules via a compact, code-like rule that is combinatorially structured and (in a sense) arbitrary. These facts provide some justification for describing genes as coding for the primary structure of protein molecules. But these facts provide no justification for the idea that genes encode the phenotypes of whole organisms, or indeed of anything beyond protein primary structure. It is not just loose talk but positively misleading to over-extend the informational description of genes, and see genes as encoding information for phenotypic traits.

Bibliography

Buller, D. (ed.), 1999: *Function, Selection, and Design*. Albany, NY: SUNY Press.

Devitt, M. and Sterelny, K. 1999: *Language and Reality*, 2nd edn. Cambridge, MA: The MIT Press.

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

Godfrey-Smith, P. 1999: Genes and codes: lessons from the philosophy of mind? In V. Hardcastle (ed.), *Where Biology Meets Psychology: Philosophical Essays*. Cambridge, MA: The MIT Press, 305–31.

—— 2000: On the theoretical role of "genetic coding." *Philosophy of Science*, 67, 26–44.

Griffiths, P. E. 2001: Genetic information: a metaphor in search of a theory. *Philosophy of Science*, 68, 394–412.

—— and Gray, R. 1994: Developmental systems and evolutionary explanation. *Journal of Philosophy*, 91, 277–304.

Kay, L. 2000: *Who Wrote the Book of Life? A History of the Genetic Code*. Palo Alto, CA: Stanford University Press.

Lewontin, R. C. 1974: The analysis of variance and the analysis of cause. Reprinted in Levins, R. and Lewontin, R. C. 1985: *The Dialectical Biologist*. Cambridge, MA: Harvard University Press.

Maynard Smith, J. 2000: The concept of information in biology. *Philosophy of Science*, 67, 177–94.

Millikan, R. 1984: *Language, Thought and Other Biological Categories*. Cambridge, MA: The MIT Press.

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

Schrödinger, E. 1992: *What is Life?* Cambridge: Cambridge University Press (originally published in 1944).

Shannon, C. 1948: A mathematical theory of communication. *Bell Systems Technical Journal*, 27, 279–423, 623–56.

Sterelny, K., Smith, K., and Dickison, M. 1996: The extended replicator. *Biology and Philosophy*, 11, 377–403.

Stich, S. and Warfield, T. 1994: *Mental Representation: A Reader*. Oxford: Blackwell.

Watson, J., Hopkins, N., Roberts, J., Steitz, J. A., and Weiner, A. 1987: *The Molecular Biology of the Gene*, 4th edn. Menlo Park, CA: Benjamin/Cummins.
Williams, G. C. 1992: *Natural Selection: Levels, Domains, and Challenges*. Oxford: Oxford University Press.
Wright, L. 1973: Functions. *Philosophical Review*, 82, 139–68.

CONTEMPORARY DEBATES IN PHILOSOPHY OF SCIENCE

Edited by

Christopher Hitchcock

Blackwell Publishing

350 Main Street, Malden, MA 02148-5020, USA
108 Cowley Road, Oxford OX4 1JF, UK
550 Swanston Street, Carlton, Victoria 3053, Australia

First published 2004 by Blackwell Publishing Ltd

Library of Congress Cataloging-in-Publication Data

Contemporary debates in philosophy of science / edited by Christopher Hitchcock.
p. cm. — (Contemporary debates in philosophy ; 2)
Includes bibliographical references and index.
ISBN 1-4051-0151-2 (alk. paper) — ISBN 1-4051-0152-0 (pbk. : alk. paper)
1. Science—Philosophy. I. Hitchcock, Christopher, 1964– II. Series.

Q175.C6917 2004
501—dc22

2003016800

A catalogue record for this title is available from the British Library.

Set in 10/12½ pt Rotis serif
by SNP Best-set Typesetter Ltd., Hong Kong
Printed and bound in the United Kingdom
by MPG Books Ltd, Bodmin, Cornwall

For further information on
Blackwell Publishing, visit our website:
http://www.blackwellpublishing.com