1

The Problem of Development

> It is not good enough to answer [questions regarding development] by saying it is simply a matter of turning some genes on and others off at the right times. It is true that molecular biology provides numerous detailed precedents for mechanisms by which this can, in principle, be done, but we demand something more than these absolutely true, absolutely vacuous statements.
>
> – *Sydney Brenner* (1974)

The central problem of developmental biology is to understand how a relatively simple and homogeneous cellular mass can differentiate into a relatively complex and heterogeneous organism closely resembling its progenitor(s) in relevant respects. This is not a new problem. It has been with us since Aristotle, at least. However, it is only recently that we have established a handle on how possibly to solve it. I am not convinced that we have yet grasped the right handle, though.

A decade ago, an advertisement for *The Encylopedia of the Mouse Genome* appeared in a biotechnology serial. The tagline read: 'The Complete Mouse (some assembly required)' (cited in Gilbert and Faber 1996: 136). The parenthetical clause refers, of course, to development. As those of us who have purchased ready-to-assemble furniture know all too well, this is indeed an onerous requirement, for the assembly process may very well have the greatest impact on final outcome! What is true of ready-to-assemble furniture is also true, I contend, of organisms believed to be 'ready-to-assemble' from DNA and assorted other material.

No one honestly believes that development can be achieved unilaterally by genes acting alone or in concert. Rather, everyone agrees that genes are important to, but not sufficient for, development. This is so, ontogenetically at

least (and perhaps also ontologically, for those concerned with ontology), and serves as the basis for the recent 'interactionist consensus' on development: the view that neither genes nor environments, neither nature nor nurture, suffices for the production of phenotypes.

I want to take this further: genes are important to, but not sufficient for, not only development but also the *explanation* of development. This epistemic and methodological claim is more controversial than the ontogenetic truism at the core of the interactionist consensus. My burden is to diminish the controversy surrounding this claim, in part by unpacking the interactive assembly of organisms.

In this chapter, my strategy is to explore a number of methodological principles used in biology; the first two of them are general, and the next three are used specifically in the context of understanding development. I provide arguments, abstracted from the biological and philosophical literature, for both the use of heuristics as such (the first principle) and for the use of particular heuristics (the second principle). For rhetorical purposes, I interpret the five principles as premises in an argument aimed at explaining development. I then illustrate how variance in the interpretation and application of the second principle yields inconsistent results and biases our biological knowledge in various ways. I argue in favour of an unorthodox reading of one of the heuristics, but a reading required by the imperative to take development seriously. In the chapters that follow, I further explore this imperative.

HEURISTICS

It is fair to say that biological phenomena are a messy lot. Though this may often be true in other domains as well, in biology, at least, a staggering number of simplifying assumptions must be made just to get a research programme off the ground. Historically, the most significant simplifying assumptions (or heuristics) employed in genetics and developmental biology have resulted in the elision of the organism as both nexus and nadir of developmental interactions. For the most part, these heuristics are well justified; they are, at least, widely accepted. Nevertheless, differences in how they are interpreted and applied generate differences in what we can claim to know about development.

Let us define 'heuristics' as *simplifying strategies to be used in situations of cumbersome investigational complexity* (Wimsatt 1980, 1986c; Gigerenzer et al. 1999). One crucial caveat about heuristics is that they are purpose

relative. As Wimsatt notes, 'all instruments in the natural, biological and social sciences are designed for use in certain contexts and can produce biased or worthless results if they are used in contexts that may fail to meet the conditions for which they were designed' (Wimsatt 1986c: 297). Examples might include the use of analysis of variance as a surrogate for the analysis of causes (Lewontin 1974; Sober 2000); the application of the methods of quantitative genetics where the assumptions of quantitative genetics (linearity, additivity, constancy, and so on) do not hold (Pigliucci and Schlichting 1997); or the use of linkage analysis in psychiatric genetics where the conditions of successful linkage (single gene of major effect, clear diagnostic criteria, known pattern of inheritance, and clinical homogeneity amongst affected family members) are not met (Robert 2000a). In using heuristics, then, we must be careful to select the right one(s).

That notwithstanding, without the use of heuristics, we would be much further from solutions to pressing biological problems than we currently are. Here, then, is a universally acknowledged premise of biological research:

1. Simplifying strategies and assumptions, as such, are absolutely necessary in biological science.

This is an heuristic dealing with the use of reductionistic heuristics. There are at least twenty reductionistic heuristics in widespread use today, including those used in conceptualisation, model building, theory construction, experimental design, observation, and interpretation; Wimsatt has documented these heuristics, and also their characteristic biases (Wimsatt 1980, 1986c).

Unlike Laplacian demons, human investigators of all stripes have limited intellectual, computational, temporal, and financial capacities. Any biological system to be studied must be simplified in various ways to make it tractable for agents like us. The very reason that we build simplified models is that we are limited beings, and most of the systems we want to understand are too complex in their natural state; thus we abstract from them what seem to be the most important or the most easily manipulated variables in order to generate a manageable representation of their workings.

One of the most common heuristic strategies is to simplify the *context* of a system under study. If we want to learn about *intrasystemic* causal factors – that is, if we want to learn about what's going on inside a particular system – we build a model or design an experiment wherein the context of the system is simplified rather than the system itself. Of course, we sometimes have to do both, especially if the system of interest is particularly complex; in such

a case, we might use another kind of reductionistic strategy. But a golden rule of experimental design is this: simplify the context first. Hence, a second general principle of biological methodology:

2. Simplifying the context of a system is advantageous if we want to learn about intrasystemic causal factors.

Amongst those who hold to the interactionist consensus, the strategy of context simplification is extensively employed in investigations of the role of genes in development, usually in the form of 'environmental control'. Here, one holds environmental variables constant across experiments or, worse, actually believes that the environment simply is invariant. One standard approach is to vary genetic factors against a common, invariant background of environmental factors – a standard environment. Context simplification, instantiated as environmental control, is the basic methodological framework of many researchers creating and employing genome sequence data, for instance. Sequence data are produced by isolating strands of DNA, cloning them, and employing a variety of techniques to ascertain the order of nucleotides and their physical relationship to each other. Genomes, or even individual strands of DNA – the systems under study – do not exist in isolation from natural environments except in the pristine artificiality of the lab; moreover, as we shall see in later chapters, there are good reasons to believe that even the structure (let alone the functions) of strands of DNA cannot be understood in isolation from their organismal context. Nevertheless, the environments, broadly construed, of DNA were abstracted away and held constant in the effort to generate the sequence of the human genome. (The same is true, of course, of the genome sequences of model organisms, such as the mouse and the nematode worm.) The context was simplified, the experimental work proceeded, and draft versions of the genome sequence are now at hand.

For the most part, and despite occasional slips to the contrary, biologists are careful in employing the strategy of context simplification. For instance, with rare but notable exceptions – such as Hamer and Copeland (1998), but see Hamer (2002) – very few scientists or commentators would today suggest that either nature (genes) or nurture (environments) is singularly decisive in organismal development. Despite the standard use of experimental or interpretive techniques to partition causation into internal (natural, genetic) and external (nurturing, environmental) components, techniques which may be unable by their very design to detect interactions between genes and environments (Wahlsten 1990; Sarkar 1998), most scholars grant that phenotypic

traits arise from complex, possibly nonadditive, interactions between multiple factors at many hierarchical levels.

However, not all varieties of interactionism are equivalent, and a vigorous debate has arisen over which varieties in fact take interaction seriously, and which simply pay 'lip service' to interaction in a reflexive refrain masking secret adherence to the old nature–nurture debate (Robert 2003). This debate will figure prominently in the paragraphs that follow, as well as in later chapters in the discussion of how best to interpret the second premise.

EXPLORING DEVELOPMENT

Let me now briefly spell out three additional premises, again universally granted, which are employed as additional steps, beginning with the first two premises, in (roughly) a chain of argument putatively leading to a conclusion about development.

The third premise, already alluded to, states the following:

3. Genes by themselves are not causally efficacious, as genes and environments (at many scales) interact (differentially, over time) in the generation of any phenotypic trait.

Whereas, once upon a time, biologists and commentators may have been happy to claim that genes determine organisms, body and mind alike, just as other scientists (mainly social scientists) and commentators were happy to claim that the organism is a kind of tabula rasa to be inscribed, shaped, and structured entirely by experience, no one seriously (or, at least, no one justifiably) entertains either of those perspectives today. It is for this reason that scientists are happy to declare the nature–nurture debate dead, settled in favour of both (Goldsmith et al. 1997). There are no (overt) genetic determinists these days, even though some environmental determinists persist (usually in an effort to ward off the spectre of genetic determinism). As Russell Gray has put it, 'nowadays it seems that everybody is an "interactionist"' (Gray 1992: 172). So much so, in fact, that those perceived to be stirring the ashes of the nature–nurture debate are called nasty names and relegated to the periphery of accepted scientific practice. This is the legacy of the interactionist consensus.

The fourth premise is designed to permit investigation of interacting variables in development (in line with premises 1 and 2):

4. We decide to focus on the causal agency of genes against a constant background of other factors, for pragmatic or heuristic reasons.

Experimental tractability is a core scientific *desideratum*. It is nice to imagine the world as full of interconnected parts not meaningfully separable from each other; but just try to analyse the world so imagined and science grinds to a halt. It turns out that genes are much more experimentally tractable than a wide range of other interacting factors and agents. This may be, of course, simply because we have spent so many decades perfecting techniques for genetic manipulation, and that huge amounts of money are available for such activities compared with others (Griffiths and Knight 1998: 255; Robert 2001b). Given the enormous amount of money available to study gene sequences, it is little wonder that genetic manipulation is quite easy compared with the experimental manipulation of other factors in development.

Nevertheless, it is worth briefly describing two scientifically well-regarded philosophical analyses justifying premise 4, such that premise 4 is universally acknowledged. First, Schaffner has published a careful study of the role of genes in the behavioural development of the nematode worm, *Caenorhabditis elegans*. Though he (and the scientists he studies) is well aware that genes must be coupled with other molecules within an organism in order to be causally efficacious (premise 3), Schaffner contends (in line with premises 1 and 2, and in support of premise 4) that '*epistemically* and *heuristically*, genes do seem to have a *primus intra pares* status'. This is in part because 'methods have been developed to screen for mutants, map "genes for" traits (as a first approximation), localise those genes, clone them, and test their role as "necessary" elements for a trait using sophisticated molecular deletion and rescue techniques' (Schaffner 1998: 234). With such methods in place, not starting with genes seems methodologically foolhardy. The embryologist Ross Harrison aptly noted early in the twentieth century that 'the investigator enters where he can gain a foothold by whatever means may be available' (Harrison 1918; cited by Gilbert and Sarkar 2000: 4).

A second, and related, justification for premise 4 is laid out by Gannett. She has analysed how genes come to be identified as causes primarily for pragmatic reasons (Gannett 1999). Having ruled out as unsuccessful the efforts of those who attempt to apply objective criteria (namely, causal priority, nonstandardness, and causal efficacy) to single out genes as causes, she argues that practical, and not theoretical, considerations are at play. Drawing on the work of Collingwood and van Fraassen on the context dependence of causal explanations, Gannett shows that what we identify as 'the' cause, amongst competing, equally necessary causes, depends jointly on the capacity to manipulate it (scientists' 'handle' – or, in Harrison's term, their 'foothold') and also the specific purposes of investigators (what sorts of questions are found meaningful and worthy of attention).

Pragmatic factors structure both of these contingencies: the capacity for manipulation is a function of past choices in, for instance, the development of particular technologies, and the questions found meaningful are decided by investigative aims, the practical end sought – for instance, the treatment or prevention of disease. Both contingencies are also deeply influenced by the availability of research funds; with the Human Genome Project, countless lab scientists suddenly saw a need for expensive gene-sequencing machines. Gannett concludes that, given the (necessary) incompleteness of causal explanations, whatever causal explanation offered will be both partial and pragmatically determined.

What we identify as a cause has its causal effects only in combination with additional necessary conditions (which, for other pragmatic reasons, might have themselves been identified as causes). This idea is epitomised in a fifth and final premise, one that may seem more controversial than the first four but is nonetheless widely acknowledged:

5. A trait *x* is caused by a gene *y* only against a constant background of supporting factors (conditions), without which *x* would not be present (even if *y* is present).

Prima facie, given premise 2, this fifth premise is a close relative of premise 3. Variations on this fifth premise have been employed as definitions of a 'genetic trait'. Consider Sterelny and Kitcher's sophisticated treatment:

> An allele *A* at a locus *L* in a species *S* is for trait P^* (assumed to be a determinate form of the determinable characteristic *P*) relative to a local allele *B* and an environment *E* just in case (a) *L* affects the form of *P* in *S*, (b) *E* is a standard environment, and (c) in *E* organisms that are *AB* have phenotype P^*. (Sterelny and Kitcher 1988: 350)

In other words, as long as that particular allele, in genetic and standard environmental context, is associated with the relevant phenotypic outcome, then that particular allele may be deemed an 'allele for' that phenotype. Given the necessity of simplifying assumptions (premises 1 and 2), as long as we recognise the critical contextual qualifications (premise 3) and also that we focus on allele *A* for heuristic and pragmatic reasons (premise 4), then we may deem premise 5 to be a plausible singling out of a gene as a cause in organismal development. So far, so good.

To reiterate, the five premises we have before us are as follows:

1. Simplifying strategies and assumptions, as such, are absolutely necessary in biological science.

2. Simplifying the context of a system is advantageous if we want to learn about intrasystemic causal factors.
3. Genes by themselves are not causally efficacious, as genes and environments (at many scales) interact (differentially, over time) in the generation of any phenotypic trait.
4. We decide to focus on the causal agency of genes against a constant background of other factors, for pragmatic or heuristic reasons.
5. A trait *x* is caused by a gene *y* only against a constant background of supporting factors (conditions), without which *x* would not be present (even if *y* is present).

These five premises taken together are usually thought to justify the following conclusion:

6. Therefore, organismal development is a matter of gene action and activation, as particular alleles have their specific phenotypic effects against standard environmental background conditions.

This conclusion coheres nicely with the standard explanation for why organisms develop as they do: there is a programme or set of instructions for development inscribed in the genes. Of course, genes alone do not an organism make. The genetic program must be activated or 'triggered', as there is no unmoved mover in the world as we know it; and the DNA must be suitably housed in appropriate cellular and extracellular contexts, which may themselves be very complex, in order for development to proceed. However, given these caveats, the specificity of development – the reliable, transgenerational reconstruction of form – is widely held to be best explained as a matter of gene action and activation.

But is that in fact true? Is development in fact *explained* in terms of gene action and activation? My argument is that it is not, though we all happily agree, at least in the abstract, with the five premises thought to generate it. Are we then illogical or, worse, illogical because we are ideologically motivated? Or is it rather the case that the five universally acknowledged premises do not actually generate the inference to the usual conclusion? I interpret the inference to the orthodox conclusion as invalid: the conclusion does not follow from the premises we have before us, because there are two mutually exclusive possible readings of the second premise just detailed, only one of which could be taken to support the conclusion. (Even were the second premise perfectly straightforward, as it does, indeed, seem to be, and even were we therefore justified in asserting the conclusion on the basis of the five

premises, we would be mistaken to interpret the conclusion as specifying an *explanation of development* – a point to which I return in later paragraphs.)

A FLAWED HEURISTIC?

Recall that premise 2 stipulates that simplifying the context of a system is advantageous if we want to learn about intrasystemic causal factors. Context simplification is usually achieved by holding certain factors constant while solving for others, and decisions about what to hold constant and what to investigate are pragmatically motivated, as already explained. However, the pragmatic dimension of these decisions renders the second premise crucially ambiguous: what counts as a system is not a matter of objective determination but is itself influenced by pragmatic factors, such that what counts as intrasystemic or extrasystemic is decided by a range of considerations and not, as it were, thrust at us by nature. Accordingly, our results are constrained by the experimental design and not the facts of nature.

Several systematic problems (what Wimsatt calls 'biases') are associated with environmental control as a context simplifier. First, context simplification is biased toward lower explanatory levels, so simplifying the environmental context stems from, and leads to, focusing on simple components of a system. Higher-level components of systems, and higher-level systems, are legislated out of epistemological and methodological existence in favour of lower-level systems and their components. Consequently, an investigator who simplifies the context in line with premise 2 may well be guilty of simplificatory asymmetry (Wimsatt 1986c: 300, 301). Second, we may be prone, should we forget or fail to appreciate the gravity of the simplifying assumption, to draw unjustified causal inferences; it is remarkably easy to fall into the trap of generating causal stories about genes against a constant environmental background (which itself exists only in the laboratory) – hence our fifth premise. We must be eternally vigilant, in simplifying the context, not to exaggerate the conclusions we draw.

I suggested earlier that premise 5 strikes us as entirely justified by appeal to premises 1 through 4. However, there is no necessity in my particular formulation of premise 5, nor in Sterelny and Kitcher's instantiation of this premise. Consider that, by parity of reasoning, we might just as well have (again for some pragmatic reason) postulated not an 'allele for' P^* but rather an 'extracellular environment for' P^* given standard allelic, cytoplasmic, and other environmental contexts (Gray 1992; Smith 1992; Mahner and Bunge

1997; Robert 2000c). That we do not postulate such 'extracellular environments for' does not imply that they do not exist; it implies, rather, that we have decided, for whatever reasons, that 'alleles for' are more important to establish. We are thereby guilty of explanatory asymmetry inasmuch as we a priori construe the relevant system in strictly reductionistic terms, thereby inviting inference to the conclusion that development is a genetic affair.

This result is fostered by only one of the 2 possible interpretations of premise 2. Both interpretations are heuristics in their own right. I shall refer to the suspect one as the 'hedgeless hedge' heuristic (HHH); the other, to be explored and defended in later paragraphs, is the 'constant factor principle' heuristic.

The phrase 'hedgeless hedge' is attributed to Roger McCain, who diagnosed hedgeless hedging as a major limitation of early sociobiological thinking (McCain 1980; see also Neumann-Held 1999). The notion, though, is more broadly applicable than that. A typical definition of 'hedging' is protecting oneself from loss or failure by undertaking a counterbalancing action, as in hedging one's bets by not placing all one's eggs in a single basket (an awkward mixture of metaphors, to be sure!). Hedgeless hedging is a win–win strategy, denoting a fail-safe type of hedging: one puts virtually all one's faith in *A* and relatively little in *B* and then attempts to establish *A* but not *B*; but betting on *B* at all (say, by publicly announcing that *B* is true, likely, or possible) provides a measure of safety just in case *B* and not *A*. Less formally, in proceeding according to the HHH, 'one admits the existence of an anomaly or problem of theory and then proceeds as though one had not. If one is then accused of neglecting the anomaly, one then produces the admission of its existence as conclusive evidence of one's innocence of the charge' (McCain 1980: 126). The hedgeless hedge is well characterised as a simplifying assumption, in particular a simplification of context: one admits the implausibility of the simplifying assumption but proceeds with the simple model nonetheless, generating results inadequate to the reality of the situation; when challenged, one refers back to the original admission of implausibility for exoneration.

McCain's example of this strategy is sociobiologists' treatment of inheritance. Although complexes of many genes (polygenes) are involved in the generation of any trait, for purposes of tractability the early models of sociobiological inheritance – such as that advanced in E.O. Wilson's *Sociobiology: The New Synthesis* (Wilson 1975) – reverted to one-locus theory, according to which we assume that one and only one gene is associated with a given inherited trait. As Wilson's mathematical models depend so heavily on one-locus theory, and the assumption of single loci is so inadequate to the reality