The Mirage of a Space between Nature and Nurture

EVELYN FOX KELLER

DUKE UNIVERSITY PRESS Durham & London 2010

Printed in the United States of America on acid-free paper ∞
Designed by Jennifer Hill
Typeset in C & C Galliard by Keystone Typesetting, Inc.

Library of Congress Cataloging-in-Publication Data appear on the last printed page of this book.

CONTENTS

included here

INTRODUCTION

One of the most striking features of the nature-nurture debate is the frequency with which it leads to two apparently contradictory results: the claim that the debate has finally been resolved (i.e., we now know that the answer is neither nature nor nurture, but both), and the debate's refusal to die. As with the Lernian Hydra, each beheading seems merely to spur the growth of new heads. In the case of the Hydra, Hercules managed to definitively vanquish the beast. I do not pretend to the prowess of a Hercules; my aim is not even to crush the nature-nurture debate. Rather, it is to understand what it is about that debate that causes so much trouble, that so stubbornly resists resolution.

Part of the difficulty comes into view with the first question we must ask: what is the nature-nurture debate about? There is no single answer to this question, for a number of different questions take refuge under its umbrella. Some of the questions express legitimate and meaningful concerns that can in fact be addressed scientifically; others may be legitimate and meaningful, but perhaps not answerable; and still others simply make no sense. I will argue that a major reason we are unable to resolve the nature-nurture debate is that all these different questions are tangled together into an indissoluble knot, making it all but impossible for us to stay clearly focused on a single, well-defined and meaningful question. Furthermore, I will argue that they are so knitted together by chronic ambi-

guity, uncertainty, and slippage in the very language we use to talk about these issues. And finally, I will suggest that at least some of that ambiguity and uncertainty comes from the language of genetics itself.

For example, we often assume, and indeed often read, that the nature-nurture debate is about sorting out the contributions of nature from those of nurture, and trying to estimate their relative importance. But what exactly is meant by nature and nurture? Sometimes the distinction is between what is inborn and what is acquired after birth; more commonly, it is between genes and environment. Moreover, these terms are themselves ambiguous: what exactly *is* a gene, and what does it do? Even more troublesome is the ambiguity of the term *environment*. Do we mean it to refer to everything other than DNA, to the milieu in which the fertilized ovum develops, or to the factors beyond the organism that affect its development? Finally, there is also the question, contributions *to what*? This, alas, we almost never ask, either as readers or as writers. Yet here we can find what may be the most commonly encountered and the most recalcitrant source of trouble with the entire nature-nurture debate, for what is at issue—the subject of debate—depends critically on our tacit assumptions about how that question is to be answered.

By far the most common assumption—at least in the popular and semipopular literature—is that what is at issue is a comparison of the contributions of nature and nurture to the formation of individual traits. For example, this is the assumption that underlies much of the argument of Matt Ridley's widely read book, *Nature via Nurture* (2003). Ridley's central thesis is that modern genomics has shown us that the nature-nurture debate, as traditionally framed, is premised on a meaningless opposition. He writes:

> The discovery of how genes actually influence human behavior, and how human behavior influences genes, is about to recast the debate entirely. No longer is it nature versus nurture, but nature via nurture. Genes are designed to take their cues from nurture. (2003, 5)

In other words, what matters for development is not so much what genes an organism has, but how and when these genes are expressed—and to be

expressed, they need to be activated by environmental stimuli. His take-home message: nature depends upon nurture to be realized.

But in a review of Ridley's book, the evolutionary geneticist H. Allen Orr argues that Ridley misses the main point of the nature-nurture debate. Orr's chief complaint is that Ridley "seems to have the right answer to the wrong question" (2003). What Orr refers to as the "traditional question" of this debate is altogether different from Ridley's concern with how genes respond to experience:

> The first question is statistical. It asks about the percentage of variation in, say, IQ, that arises from inherited differences among individuals (do some parents pass on smart genes to their kids?) versus the percentage that arises from environmental differences (do some parents pass on books to their children?). The second question is mechanistic. It asks about how genes behave within individuals . . . The fact that genes respond to experience is certainly interesting and important . . . But it's the wrong *kind* of fact to settle the nature-nurture debate. (ibid.)

To Orr, the difference between the two questions seems clear, and we might ask (as in fact he does), how so sophisticated a science writer as Matt Ridley could make so elementary a mistake: "why does Ridley reach for the wrong level of analysis, confounding statistics and mechanisms?" Orr suggests that the explanation is as plain as the mistake: Ridley, he writes, "a self-styled champion of 'techno-optimism,' seems to have succumbed to genome hype" (ibid.).

I disagree. What Orr describes as Ridley's confusion between statistics and mechanism is simply too widespread, too difficult for both readers and authors to detect, and too resistant to clarification to be explained by excessive "techno-optimism." The conflation is everywhere, in popular and technical literature alike. It may well be that the distinction seems clear to Orr, but if so, if he himself never slides from one meaning to another, then he is truly an exception.

For another example of the same slippage, and to illustrate its ubiquity, consider the numerous arguments currently being made that invoke epi-

genetics (rather than genomics) as the crucial new science that resolves the nature-nurture debate. In an article from a recent issue of *Scientific American Mind*, we read:

> Psychologists, psychiatrists and neuroscientists have jousted for years over how much of our behavior is driven by our genes versus the environments in which we grow up and live. Arguments have persisted because there has been little hard evidence to answer basic questions . . . [But today] a field called epigenetics has finally begun to address some of these issues. (Steinberg, 2006)

In a similar vein, an article entitled "Genes or Environment? Epigenetics Sheds Light on Debate," published in a newsletter of the National Institutes of Health, asks:

> Which is more important in shaping who we are and what we will become—our genes or the environment around us? For centuries, people have debated whether nature or nurture decides how we look and act. Now, a field of research called epigenetics is showing that we can't really separate one from the other. (*News in Health*, February 2006)

What is it that makes these claims startling? Not that we have finally found a way to resolve the nature-nurture debate—that claim is about as old as the debate itself. Nor is it that we have finally learned that nature and nurture cannot really be separated. Rather, it is, first, the supposition that that lesson is somehow new, and second, that epigenetics—which I'm all for—is being invoked as the agent of resolution.

Epigenetics is a term that Conrad H. Waddington (1942) coined to refer quite generally to developmental processes (i.e., how we get from genotype to phenotype), and we have known for a long time that such processes involve far more than DNA. In this sense of the term, epigenetics is not a new field. Also not new is the recognition that the various factors involved in development—nucleic acids (DNA and RNA), metabolites, and proteins; nuclear and cytoplasmic factors; genetics and environment—are so deeply intertwined, so profoundly interdependent, as to make any attempt to partition their causal influence simply meaningless. Long

before the discovery of DNA, the geneticist Lancelot Hogben was obliged to caution his readers that "genetical science has outgrown the false antithesis between heredity and environment productive of so much futile controversy in the past" (1933, 201).

What is new today involves an altogether different reference of the term *epigenetic*. The "field of research called epigenetics" in the NIH newsletter refers primarily to the discovery that not only are changes in various extra- (or epi-) genetic factors affecting phenotype routinely passed on in cell division, but also such changes can often be transmitted through the generations, despite the fact that they do not involve changes in DNA sequence. Examples of epigenetic inheritance might involve changes in methylation patterns on DNA, or changes in chromatin structure, metabolic requirements, feeding patterns, or even modes of symbolic communication. These alternative "systems of inheritance"[1] are of immense importance to development; they radically change our understanding of inheritance; and they can also have a profound effect on evolution. The important point here, however, is that in themselves, such effects have little if anything to do with the nature-nurture debate as formulated in the examples I cite.

There is surely a muddle here—actually, there are several. There is the muddle about the meaning of *epigenetic* that I have just described; there is another muddle about what kinds of things can be separated, what kinds of questions can be answered; and there is a third muddle about what we take the nature-nurture debate to be about. The casting of the debate as an effort to determine "how much of our behavior is driven by our genes versus the environments in which we grow up and live" poses a question that is not only unanswerable but, as I have already indicated, is actually meaningless. Indeed, we scarcely need the new sciences of genomics and epigenetics to teach us this lesson.

In thinking about development, the causal components belonging to nature are conventionally taken to be the hereditary units that modern biology calls *genes*. The concept of the gene is in considerable disarray these days, but let us for the moment accept the hypothesis that such discrete entities exist as units of inheritance. Just what the components of

nurture (environment) might be, I have no idea, but even if we could identify them, we would still have a serious problem. What is the causal role of a gene in the absence of environment? None is clearly the answer. Absent environmental factors, genes have no more power to shape the development of an individual than do environmental factors in the absence of genes. Let us take the simplest possible contemporary definition of a gene on offer: a sequence of DNA nucleotides that codes for a protein.[2] By themselves, such sequences of nucleotides don't do anything: DNA is an inert molecule. What we think of as its causal powers are in fact provided by the cellular complex in which it finds itself. It is this complex that is responsible for both the code that enables a sequence of nucleotides to be translated into a sequence of amino acids, for the replication of DNA, and for the intergenerational fidelity of replication; it is the cellular complex that makes possible all the chemical reactions on which these processes depend. By themselves, the entities we call genes do not act; they do not have agency. Strictly speaking, the very notion of a gene as an autonomous element, as an entity that exists in its own right, is a fiction. In order for a sequence of nucleotides to become what is conventionally called a gene requires that the sequence be embedded in a cellular complex that not only reads, translates, and interprets that sequence, but also defines it, giving it its very meaning.

Not only is it a mistake to think of development in terms of separable causes, but it is also a mistake to think of the development of traits as a product of causal elements interacting with one another. Indeed, the notion of interaction presupposes the existence of entities that are at least ideally separable—i.e., it presupposes an a priori space between component entities—and this is precisely what the character of developmental dynamics precludes. Everything we know about the processes of inheritance and development teaches us that the entanglement of developmental processes is not only immensely intricate, but it is there from the start. From its very beginning, development depends on the complex orchestration of multiple courses of action that involve interactions among many different kinds of elements—including not only preexisting elements (e.g., molecules) but also new elements (e.g., coding sequences) that are

formed out of such interactions, temporal sequences of events, dynamical interactions, etc. Compounding the entanglement between genes and environment yet further, as Ridley so clearly explains, biologists now recognize that the development of phenotypic traits is guided not so much by the actual sequence of nucleotides as it is by patterns of gene expression that are themselves products of an immensely complex web of interactions between environmental stimuli (both internal and external to the cell) and the structure, conformation, and nucleotide sequence of the DNA molecule.

Accordingly, as the Swiss primatologist Hans Kummer remarked some years ago—and as Frans de Waal (2002) reminds us—trying to determine how much of a trait is produced by nature and how much by nurture, or how much by genes and how much by environment, is as useless as asking whether the drumming that we hear in the distance is made by the percussionist or his instrument. Richard Lewontin offered another metaphor: "If two men lay bricks to build a wall, we may quite fairly measure their contributions by counting the number laid by each; but if one mixes the mortar and the other lays the bricks, it would be absurd to measure their relative quantitative contributions by measuring the volume of bricks and of mortar" (1974, 401). All of these are intended to make the same point: the effort to separate "under two distinct heads the innumerable elements of which personality [or any other aspect of phenotype] is composed" (as Francis Galton posed the question in 1874) makes no sense; neither personality nor, for that matter, any other feature of an organism is composed of separable elements.

The point is a logical one about which there ought, at least in principle, to be no debate: causes that interact in such ways simply cannot be parsed; it makes no sense to ask how much is due to one and how much to another. As even Steven Pinker (2005) seems at times to recognize, "nature and nurture are not alternatives." This is true whether we read nature as innate and nurture as acquired, or nature as genes and nurture as environment. To drive the point home, my colleague Ned Hall has developed a cartoon version of the philosophical conditions that must be met in order to parse the causes of a phenomenon, a representation he refers to as

1 The Bucket Model. Adapted from cartoon by Ned Hall.

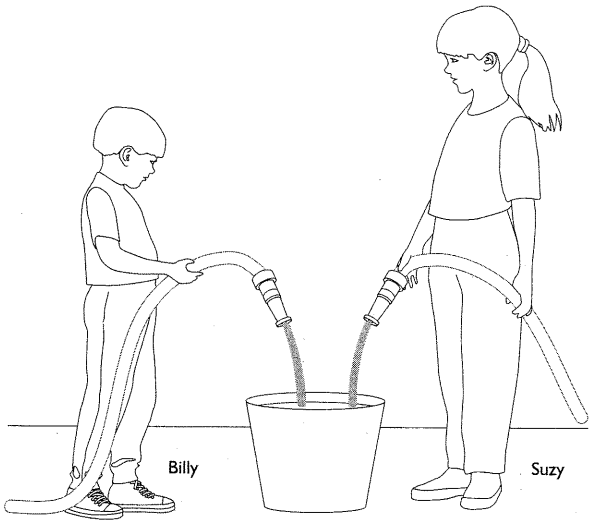

Here is a bucket: Billy fills it with 40L of water; then Suzy fills it with 60L of water. So: 40% of the water in the bucket is due to Billy, 60% to Suzy.

2 When Causes Interact. Adapted from cartoon by Ned Hall.

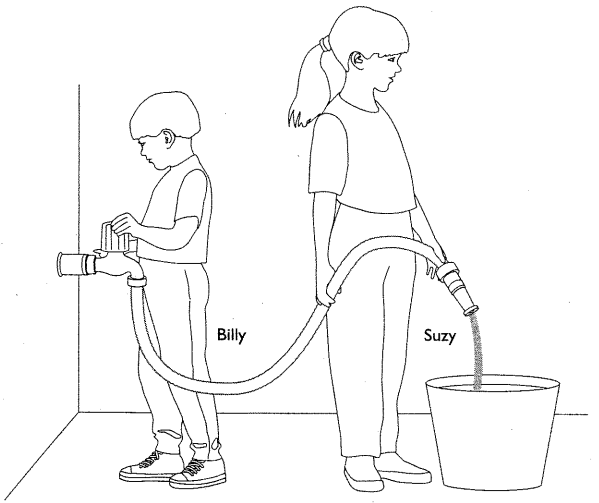

But suppose instead that what happened was this: Suzy brought a hose to the bucket; then Billy turned the tap on. Now how much of the water is due to Billy, and how much to Suzy?

Answer: The question no longer makes any sense.

the bucket model. I illustrate with two figures borrowed (and adapted) from Hall:

In figure 1, innate qualities (or genes) fill the bucket to height x, and personal (acquired) experiences (or environment) add an amount y. Put in these terms, it is easy to see that the categories of innate and acquired cannot be represented in this way, that a kind of logical error has been committed. We do not have to think of the neonate as a blank slate in order to appreciate the fact that many, if not all, of the traits we think of as innate depend upon the acquisition of nourishment, parental care, and socialization in order to develop. As I will show in a bit, even Locke did not harbor such a view. The true relation between innate and acquired, or between genes and environment, is more like that illustrated in figure 2.

Simple, right? But perhaps too simple. If all that was at issue in the nature-nurture debate was a comparison of the contributions of nature and nurture to individual development, then both Ridley and the commentators on epigenetics I cited are of course correct: this question is meaningless ("we can't really separate one from the other"), and the debate could indeed be said to be over. But unfortunately, the question of what the nature-nurture debate is about is not so easily settled. As Orr makes abundantly clear, to population geneticists, the debate is not about relative contributions to individual traits, but about contributions to the variation within a population. Still others think of it as being about the relative importance of the contribution of nature and nurture to differences between individuals. Furthermore, not only do different people have different questions in mind, but individual authors (e.g., Ridley) themselves tend unwittingly to vacillate between the various options.

Also, as I have tried briefly to indicate, there is more than a single source of confusion here—in fact, trying to make sense of how arguments about nature and nurture proceed quickly reveals a morass of linguistic and conceptual vegetation grown together in ways that seem to defy untangling. Indeed, it is precisely this morass that is the subject of my book. I want to explore both its conceptual underpinnings and its history. I want

to understand how a debate that is so frequently claimed to have been resolved manages, after each such claim, to survive and persist; and how so many readers and authors of manifest intelligence, scientific sophistication, and good faith continue to be enmeshed by the tangle of meanings that make up the nature-nurture debate.

The question that faces us is therefore twofold. First, how did we come to be so entangled, and second, how and why do we persist in making the very same errors, over and over again? The question I am asking is not only how (when, where) a sense of opposition, a versus, came to be inserted between nature (understood either as innate or as genetic) and nurture (understood either as acquired or as environmental), but, more specifically, how did the separation on which opposition is predicated—the *and* on which a *versus* depends—come to be so taken for granted? Whence came the mirage of a space between nature and nurture? How did this illusion become so deeply entrenched in our thinking, and why is it so resistant to dissolution?

The persistence of the nature-nurture debate has been a source of considerable puzzlement to many scholars. Many years ago, the developmental psychologist Daniel S. Lehrman had this to say:

> When opposing groups of intelligent, highly educated, competent scientists continue over many years to disagree, and even to wrangle bitterly, about an issue which they regard as important, it must sooner or later become obvious that the disagreement is not a factual one, and that it cannot be resolved by calling to the attention of the members of one group . . . the existence of new data which will make them see the light . . . If this is, as I believe, the case, we ought to consider the roles played in this disagreement by semantic difficulties arising from concealed differences in the way different people use the same words, or in the way the same people use the same words at different times; [and] by differences in the concepts used by different workers . . . (Lehrman 1970, 18–19)

It is the aim of this book to follow up on Lehrman's suggestion that we look to semantics—to the vicissitudes of the language we use—for an understanding of our impasse. But before proceeding, a word about the peculiar difficulty of my task: in making this analysis, I employ the same language that I critique, and necessarily so. Indeed, I have no other. My description is therefore prone to the very same kinds of slippage I am attempting to diagnose. Even my best efforts do not totally guarantee that the writing of this text has entirely avoided the problem, and for readers who may be less attentive to the problem, the difficulty is that much greater. The net result, I fear, can be somewhat dizzying. Perhaps a schematic outline may help.

Chapter 1 is historical. It focuses on the emergence of a bucket model of the roles of nature and nurture in the formation of traits—i.e., on the assumption already implicit in Francis Galton's catchy phrase, "Nature and Nurture" (1874) that there exist two domains, each separate from the other, waiting to be conjoined. Galton was hardly the first to write about nature and nurture as distinguishable concepts, but he may have been the first to treat them as disjoint. As far as I can tell, such an assumption of mutual exclusivity was not made by earlier writers. For those who used the terms, nurture was rarely, if ever, seen as separable from nature; instead, it was referred to as helping and assisting, or as responding to, nature; nurture was more of a verb than a noun. But those writing after Galton did tend to disjoin the two, increasingly so over time. What is especially noteworthy to me is that the shift in formulation followed directly on the heels of the introduction of a particulate theory of inheritance in the last third of the nineteenth century. Indeed, I argue that this shift was greatly assisted by the arrival of a new way of conceptualizing heredity, and perhaps even dependent upon it.

This remainder of the book brings the story into the modern era. Contemporary biology of development has clearly exposed the assumption in Galton's formulation as meaningless, as not making sense. Indeed, the problem was already evident early in the twentieth century, and in an effort to salvage the questions that had interested Galton, a reformulation of his project was soon provided. The English statistician R. A. Fisher was one of the first who taught us the necessity of making two fundamental distinctions if we wished to address Galton's concerns: we must distin-

guish first between traits and trait differences, and second between individual and population. Although we may not be able to parse the causal contributions of genetics and environment to individual traits, we *are* able, at least under some circumstances, to statistically parse the causal contributions of differences in genetics and environment to differences in traits averaged over a population.

The following two chapters focus on the nature of these distinctions, the enormous difficulty we seem to have in holding to them, and at least some of the reasons why. In chapter 2, I try to clarify the nature of the distinction between traits and trait differences. I claim that this distinction poses particular difficulty for us when formulated in terms of genetics, and I proceed by locating at least part of this difficulty in the discourse of genetics and its history, briefly indicating at the end of the chapter the ways in which contemporary genetics might offer us a way out of our difficulty. In chapter 3, I focus on the distinction between individual and population that was of such particular importance to Fisher. But here too we seem to have enormous difficulty in keeping the two categories apart, and here too I locate at least part of the problem in language—this time around, in the polysemy of the word *heritability*. In the technical literature of population genetics, heritability was defined as referring to a statistical measure that has meaning only in relation to populations. Unfortunately, however, the word was already in use, but with another, simpler meaning—namely, transmissibility from parents to offspring. The double meaning of heritability has been frequently noted, but, in my view, its role in the continuing confounding of the two meanings, and accordingly of individual and population dynamics—both in the technical and popular literature—has not been adequately pursued. Most scientists tend to underplay the importance of ambiguity (and the equivocation it invites) in scientific argument; they are wont to insist that they know what they mean, and perhaps they do. And certainly, in the immediate context of their experimental practice, where precision is mandatory, scientists generally do know exactly what they mean. But in other contexts (e.g., where words are serving as placeholders for connecting one set of observations with another), there may be considerably less need for, or even less possibility of, precision. In any case, heritability would seem to be a subject about which it is not only easy but virtually inescapable to mean more than one thing at the same time.

My final chapter addresses the question, what is to be done? Is it possible to reformulate the questions at issue in the nature-nurture debate in ways that not only capture what is of primary interest to people, but also, at the same time, lead to meaningful and answerable questions? A central point of this book is to argue that at least part of the explanation for the unreasonable persistence of the debate is to be found in the language of particulate genes, a language that originally developed out of the hope that a particulate theory of inheritance might do for biology what the atomic theory had done for chemistry. For good or bad, biology has turned out to be vastly more complicated, and so has genetics itself. Not surprisingly, however, the language of genetics lags behind, and by doing so, helps to keep alive debates that no longer have meaning—indeed, that may never have had meaning. My hope is that a language that better reflects the contemporary science will be able to help us out of the morass in which we have been bogged down for so long, and help put us on track to providing information relevant to the legitimate concerns that many people have.

and not all of them did. Jane Hume Clapperton, for example, a prominent socialist, feminist, and eugenicist of her day, apparently felt no need to assume a disjunctive relation between nature and nurture. To her, the importance of eugenics followed easily from a more traditional view, one in which even though nature and nurture are not in conflict, one can nevertheless readily see that "the power of nurture is limited. It can direct the forces of nature, but it can not alter the intrinsic quality of the raw material which nature provides" (Clapperton 1885, 365–66).

Historically, Pearson's influence was of course infinitely greater than Clapperton's, and Pearson did embrace Galton's social and scientific agenda as a package. In particular, the need for "race improvement" was for Pearson, just as it was for Galton, an important backdrop to his own efforts to sort the effects of nature versus nurture — or, as he later came to put it, the effects of heredity versus environment.[23] I suggest, however, that — for the particular purpose of understanding the origin and persistence of the belief in the separability of nature and nurture (my original question) — the issue of eugenics might logically, even if not historically, be put aside.

In any case, it is to the persistence of this illusion — and the place of Galton's particular vision of the relation between nature and nurture in the modern imagination — that I now need to turn. Galton may not have been in a position to perceive the defects of his formulation, but the modern reader clearly is. Today, it is widely accepted by contemporary biologists and lay readers alike — as Daniel Dennett puts it, "everyone knows"[24] — that genes and environment must interact to produce any biological trait, that nature (understood as heredity) and nurture (understood as environment) are not alternatives. And yet. And yet, the image of separable ingredients continues to exert a surprisingly strong hold on our imagination, even long after we have learned better. Although "everyone knows" it not to be true, many authors continue to argue as if nature and nurture were (or, at the very least, could be regarded as) separable and clearly distinguishable causes of development.

TWO Changing the Question to One That Does Make Sense — From Trait to Trait Difference

Galton's question (how much of our personality is due to nature and how much to nurture?) may have been badly posed, but is there not something in it of interest to many people that *could* — perhaps in some other way — be legitimately asked? Galton worked hard to formulate meaningful measures of the relative importance of heredity, and his efforts helped launch the science of biometrics. Karl Pearson continued Galton's efforts, and, building on his work, founded the discipline of mathematical statistics. But it was probably the English mathematician R. A. Fisher (1890–1962) who contributed most to reformulating Galton's question. Fisher devoted his life to trying to distinguish genetic from environmental influences, and he was clearly aware of the difficulties involved. Like Pearson, he shared many of Galton's concerns; perhaps especially, he too felt that a science of eugenics was much needed, both socially and scientifically. In 1911, while still a student, he helped form the Cambridge University Eugenics Society (together with John Maynard Keynes, R. C. Punnett, and Charles Darwin's son Horace). And a few years later, in an effort to save what was meaningful in Galton's quest, Fisher (1918) published a reformulation of that quest in a paper that was to be enormously important in shaping the future of population genetics. Here he clearly recognized — indeed, he was one of the first to so clearly recognize — the point that if it were to be realizable, Galton's hope of sorting genetic from

environmental influences would need to be recast in two important ways. First, it was necessary to reformulate the question of causation in terms of trait differences rather than in terms of traits per se, and second, it was necessary to turn from the analysis of heredity in individual lineages to the analysis of heredity in populations. Only if we ask a statistical question about the relative contributions of variations in genetics and in environment to our differences from each other—rather than their relative contributions to the processes that make us what we are—would we have a question that makes sense, and furthermore, one that we might be able to answer.

First Questions First

Asking about contributions to phenotypic differences rather than to the making of a particular phenotype still leaves a potential ambiguity, and thus we need to ask, what exactly do we mean by difference? In one meaning, difference refers to pairs of particular individuals. Returning to the cartoon characters of my colleague, Ned Hall, we might want to know what makes Suzy taller than Billy.

Historically, however, the interest in human differences that has repeatedly manifested itself is of another, far more general, sort. We want to know not simply what makes Suzy different from Billy, but also what makes people like us different from people like them. That is, we have a tendency to sort people into demarcated groups, with the express intention of comparing them to one another. For example, we like to ask what makes women different from men, whites different from blacks, Europeans different from Asians, achievers different from nonachievers, etc. And as I've already mentioned, in order to address questions of this sort, a second move is required. The interest in what makes classes of people different from one another requires us to shift our focus not simply from trait to trait difference, but also from individuals to populations. Of course, in doing so, we must also decide what the population of interest is. For example, are people like us Caucasians, females, septagenarians, Americans,

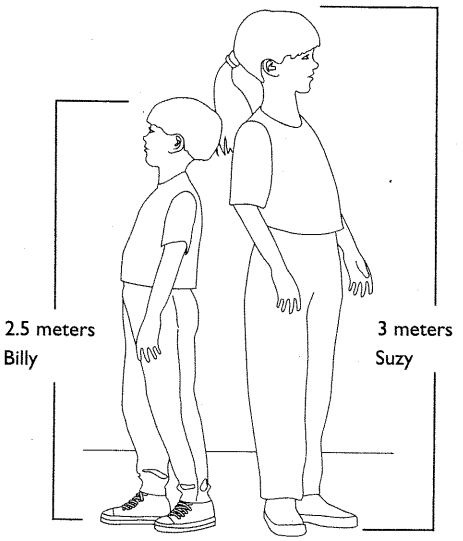

3 Comparing Traits. Adapted from cartoon by Ned Hall.

academics, hikers, or troublemakers? Needless to say, claims that address differences between whatever groups we happen to focus on are those that arouse the most intense and bitterest controversies of all.

In 2005, Lawrence Summers, then president of Harvard University, set off a firestorm with his speculation that innate differences between the sexes might account for women's underrepresentation in math and science. His remarks were ill advised for all sorts of reasons, scientific as well as political, and Summers was to live to regret them. But there are two reasons why I am not here going to address—at least, not directly—the question of why those claims (or the many other similar claims with which we are daily confronted in the mass media) are problematic. My first reason is that the most immediately obvious problem with such claims—namely, the conspicuous failure of the cases investigated to meet the basic requirement of noninteraction between nature and nurture that is required for any parsing of causal contributions to trait differences—has already been eloquently addressed by many others (for example, Lewontin 1974; Block 1996; Bateson 2004). The central issue in these discussions

is whether or not the interdependence of genes and environment in development necessarily undermines any measurement involved in the effort to partition the statistical variance.[1]

Perhaps more important, however, is my second reason for not focusing on these controversial claims: the issues I do want to examine here are, I claim, anterior. These are the underlying muddles we get into whenever these topics arise. They concern the difficulties we have, first, in making the distinctions that are required for turning the nature-nurture debate into potentially meaningful questions, and second, once those distinctions are made, the extraordinary difficulties that virtually all of us—scientists and lay readers alike—have in keeping to them.

Both historically and scientifically, the two shifts that Fisher's reformulation requires—from trait to trait difference, and from individual to population—have tended to occur together. But for the sake of trying to illuminate the conceptual aspects of the difficulties we have with them, I want to tease the two apart, and to examine each one separately. I will do so in order. In this chapter, I focus on the distinction between trait and trait difference, formulated not as a difference between populations but as a difference between individuals, deferring my discussion of the distinction between individual and population to chapter 3. In both chapters, I argue that a good part of the difficulty we have in maintaining the necessary shifts is to be found in the slippage of meaning, or polysemy, of our basic vocabulary. When the same term can be used to refer to two different concepts, it should be obvious that any attempt to distinguish one from the other becomes hopelessly confounded. Yet this, I claim, is precisely the situation for both of the distinctions with which we are here dealing.

Parsing the Causes of Trait Differences

What makes Suzy taller than Billy, Amy shyer than Mary, Richard more agile than Peter? There is nothing wrong with such questions: they are perfectly legitimate and, indeed, seem natural to ask. Sometimes it might even be possible to answer them. It turns out that, for quite straightforward technical reasons, questions like this about human behavioral traits can almost never be answered, but my point here is a different one. It is to call attention to the fact that such questions are of a different kind from those we might ask about what makes Suzy grow to almost six feet in height, or about what is responsible for Amy's stuttering, or for Richard's ability to leap five-foot hurdles.

In fact, this is precisely the point of Hans Kummer's analogy to drums and drumming. It is useless to ask whether the drumming that we hear in the distance is made by the percussionist or his instrument because each of the two variables on which the sound—the percussionist's performance (x) and the resonance of the instrument (y)—is influenced by the other in ways that simply do not permit separation. Yet if we hear two different sounds of drumming in the distance, we can ask and perhaps determine whether the difference between the two sounds is caused by a difference in drummers or by a difference in drums, or even how much of the net difference in sound is caused by the former and how much by the latter—provided, of course, the drummer's performance does not depend on the drum being used.

The questions differ not only in what is being asked, the circumstances in which they make sense, and the methods by which they might be answered, but also in the interests that give rise to them. Certainly, the new question—about the causes of difference—is one in which many people have considerable interest. We want to know what accounts for the difference in sound between different musical performances; even more, we want to know what makes people different. In the first case, the reason for our interest seems obvious: we want to find ways of making music sound better. In the second case, the link to practical intervention—to finding ways of making better (taller, smarter, more sociable) people—seems more problematic, both technically and politically. Although it may have seemed straightforward to Galton, the acknowledgment of such a link, especially in the case of the effect of genetic differences on behavioral traits, has in more recent times become so fraught with unacceptable political implications as to be generally avoided.[2]

In any case, it turns out that there is a rather stringent condition that must obtain for questions of this sort to be answerable (or indeed, for them to be well posed), and, using the terms of our example, it is this: that there be no interaction between the quality of the drummer's performance and the quality of the drum. But should this condition apply—i.e., should the performance of the drummer be precisely the same on any drum—it would be logically possible to parse the causal contributions of a change in drummers and a change in drums to the net difference in sound. Equally, if there were no interaction between changes in genetic and environmental (nongenetic) factors, it would be logically possible to parse the causal contributions of alterations in each of these factors to a change in phenotype. Were such conditions to be satisfied, we would be able to say, for instance, that 40 percent of the difference between the two sounds we heard was due to the change in drums, and 60 percent was due to the change in drummers. Or that 45 percent of the difference in height between Suzy and Billy was due to the difference between their genotypes, and 55 percent was due to nongenetic differences (e.g., nutrition) between the two. This can be important information, yet it still does not allow us to say anything about the proportion of genetic to nongenetic influence in the formation of the trait.

The Elusiveness of Clarity

It may seem unduly cumbersome to fully specify the ways in which the question has been altered in each and every iteration, but failing to do so just makes it that much easier to lose sight of the difference. Even Wikipedia, although clearly registering the fact that the question has been changed, takes a shortcut: it casts the nature-nurture debate in terms that explicitly mark the shift of focus to trait differences, but not the shift from the respective roles of nature and nurture to the respective roles of *changes* in nature and nurture in causing trait differences. Thus it reports:

> *Nature vs Nurture* is a shorthand expression for debates about the *relative importance* of an individual's innate qualities ("nature") versus personal experiences ("nurture") in determining or causing individual differences in physical and behavioral traits.[3]

Furthermore, even for apparently meticulous formulations, a good deal of care is warranted on the part of us as readers. Kummer may be clear about the difference, but the rest of us are not. In fact, it has proven extraordinarily difficult for all of us, both as readers and as writers, to keep hold of the distinctions that have been made here. Indeed, even some of our most eminent philosophers, at times, write as if they too have lost sight of the difference. For example, Elliott Sober has given some of the clearest expositions of the confusions I have been discussing, yet even his lucid prose sometimes invites slippage. For example, consider his response to the claim that the entanglement of genes and environment makes it impossible to parse the causes of development. He writes: "Does this mean it makes no sense to ask whether genes or environment are more important causes of the resulting phenotype? No. It just means that we must pose the question in a different way" (Sober 2000, 357). My point is that we are not simply reposing a question, but substantively changing it.

On first reading, I assumed Sober's reference to "the question" reflected no more than a momentary lapse in his normally meticulous prose, but a subsequent paper made me begin to suspect that there may be more at issue here. "Separating Nature and Nurture" refers to the "attempt to disentangle the contributions of nature and nurture"—e.g., when breeders "think about what makes corn grow tall" (Sober 2001, 47). Sober's main point in this paper is that "the concepts of phenotypic, environmental, and genetic variance apply to human beings just as much as they do to cows and corn" (48), despite the fact that, for humans, significant technical difficulties impede our ability to sort out these different variances from one another. In other words, the difficulties we face are not just political, but also scientific. I agree. What troubles me about Sober's account is the failure to specify exactly what contributions are being disentangled, and to what they are contributing. Sober does not intend his question to collapse into Galton's—his disentanglement is of the contributions of nature and nurture (more precisely, of differences in nature and nurture) to pheno-

typic differences, not to traits themselves. But for the most part, this distinction goes unmarked, leaving the reader free to interpret as he or she wishes. Indeed, the same ambiguity also haunts Sober's effort to explicate the limitations of the associated statistical analyses. For example, he writes: "The analysis of variance permits one to infer *how much* a cause contributes to an outcome without understanding *how* the cause manages to have its effects" (italics in original). But the problem is not with the distinction between how and how much, but rather with the distinction between two different kinds of outcome. One (the question of how) refers to an individual trait, and the other (the question of how much), to a difference between traits.

A similar slighting of the difference between the two questions can be seen in C. Kenneth Waters's otherwise illuminating analysis of the relation between actual and potential causes (2007). Indeed, following James Woodward (2003),[4] Waters interprets the relation "x causes y" as meaning "a difference in x causes a difference in y." This interpretation enables him to slide easily from statements about genes' "causing" traits to statements about gene differences' causing trait differences, and vice versa.

The demands of experimental practice amply justify Woodward's (and Waters's) focus on manipulability, and I totally sympathize. Furthermore, where the outcome depends on only one variable, the relation between the two kinds of statements is relatively easy to see. Put in terms of elementary calculus, if we represent the dependent variable y as a function of x (the independent variable), we would say that y can be obtained from df/dx by integration—provided one knows the derivative df/dx for all values of x. But where more than one variable is at issue (as, e.g., in the case of the developmental processes that give rise to a trait such as height, where a vast number of factors participate), and we wish to assess the relative importance of each such factor (or variable), a huge problem arises. First, we must identify all the variables involved (itself a far from trivial task), and second, we must assess their independence, for it is only among independent variables that we can begin to partition the different causal influences, assigning independent weight to each. For the development of height, the list of variables surely includes both genetic and environmental factors, but we can rarely, if ever, be confident that we have identified all the factors that might be involved. Furthermore, few if any of these causal factors operate independently; indeed, their very demarcation is, like the demarcation of traits themselves, often quite arbitrary.

Even situations where all the variables have been clearly demarcated and enumerated—even when, say, there are only two variables at stake—there is still a serious problem. While it is certainly legitimate to infer that a and b causally participate in the production of y from observations showing that both changes in a and changes in b cause change in y, unless a and b are the only causal factors, and unless they act independently, such observations do not permit us to answer questions like, how much of y is due to a and how much to b? They might enable one to ask (and answer), how much of the change in y is due to the change in a, and how much to the change in b? But the problem with which I am concerned arises when we attempt to infer from the answers to that question (*how much*) an answer to the earlier question of *how*, because in general, this is not possible. With Waters's interpretation of causal statements of the form "x causes y" as meaning "a difference in x causes a difference in y," it becomes all but impossible to mark, or even to recognize, the distinction between the two questions, and correspondingly easy to think that by answering the second, one has also answered the first.

The reader might wonder why I am belaboring the issue. How much does the difference between the two questions really matter? After all, knowing how changes in each variable alter the outcome must surely tell us something about the role each variable plays in producing the outcome—if it did not, why would scientists spend so much effort studying the effect of specific changes? Indeed, such measurements seem to lie at the very heart of the experimental method, and, as I have already acknowledged, a finite effect of a change in variable does identify that variable as a causal factor. The two questions are clearly related, but in systems as complex as the developing organism, that relation can be extremely opaque.

This is a lesson classical geneticists knew well. At least when pressed, they readily acknowledged that the study of phenotypic differences was of limited use in trying to tease out the influence of genotype on phenotype.

But as I discuss below, phenotypic differences were all they had access to, and the habit of conflating the two questions—of attempting to infer a causal relation between genes and traits from analyses of trait differences[5]—may well have begun with the limits of their methodology. But this habit has obvious dangers: just consider, e.g., the folly of attempting to understand the causal dynamics of vision by studying all the ways in which blindness (an extreme example of phenotypic difference) can be induced.

Molecular geneticists are in many ways far more fortunate than their classical precursors. They have direct access to the entity that we generally accept as being of primary importance to the development of traits—namely, the DNA. They can, for example, directly induce specific changes in the DNA and then study the effect such changes have on phenotype. But despite such direct access to at least part of the causal substrate of development, anyone expecting to be able to infer the causal dynamics of DNA from studying the effect of differences in that molecule on phenotype is subject to a few surprises. For example, the effect of changing a variable that is itself known to be causally important to the production of the phenotypic end product may be (in fact, often is) reduced or erased by a system of buffering that is built into the dynamical networks mediating between genotype and phenotype. Indeed, when such effects (or rather, the absence of effects) were first observed, geneticists greeted these results with surprise and consternation. More recently, however, such insensitivity to changes in contributing variables has come to be recognized as the hallmark of systems designed to be robust in the face of common fluctuations. In such systems, the difference effected by a change in variable is no guide at all to the causal importance of that variable.

Of course, genetics is not the only context in which questions about the causal dynamics of a process are conflated with questions about the effect of perturbations on that process, yet I strongly suspect that the tendency is particularly acute here, and for reasons that lie deep in the language of that science. Indeed, when talking about traits, the conclusion that it is possible to decompose a difference in phenotype, even when one can say nothing about the composition of the phenotype, itself seems to contradict basic sense. Imagine a situation in which we were able to conclude that drinking

four extra glasses of milk a day throughout her childhood was responsible for 70 percent of the difference in height between Suzy and her non-milk-drinking sister Mary, and still not be able to say anything about the relative contributions of genetic and nongenetic influences to Suzy's (or Mary's) height. That seems altogether counterintuitive, the result of some sleight of hand. But I suggest that at least part of the reason it seems so counterintuitive is that an actual sleight of hand—slipping from talk about differential effects to talk of underlying causes—has become so routine in discussions of genes and environment as to go unnoticed, even to the point of having become part of much conventional reasoning about genetics.

For an example of this sleight of hand, I turn to yet another philosopher of science, Neven Sesardic. Sesardic is known for his extensive critique of the arguments that Lewontin, Jencks, Block, and others have made, and that I have here taken as standard. Much of Sesardic's argumentation depends on close reading of the works that Lewontin et al. take to task, but Sesardic's own formulations unfortunately escape comparable scrutiny. Although his arguments center on the concept of heritability that I discuss only in chapter 3, I include them here because of their illustrative conflations between traits and trait differences on the one hand and between the causal impact of genes and gene differences on the other. For example, he writes: "The idea that heritability reflects the causal strength *of genetic influences on phenotypic differences* has been consistently opposed by a number of authors. It has been said, e.g., . . . that it is dubious whether a clear meaning can be given to *'genetic determination of traits'* (Burian); that inferences about genetic determination of traits should be 'disavowed once and for all' (Kitcher); . . . and so forth" (1993, 399, italics added). Notice that the difference between "the causal strength of genetic influences on phenotypic differences" and the "genetic determination of traits" is completely elided here. In a similar vein, Sesardic elsewhere asks: "Does the fact that a given phenotypic trait is heritable entail that it is genetic (i.e., that the differences in that trait are due to genetic differences)?" (2003, 1002). This question clearly implies that we can equate the claim that a phenotypic trait is genetic with the claim that differences in that trait are due to genetic differences. Sesardic's recent book (2005)

on the subject of genetic causation is permeated by the same slippage. Throughout his writing, Sesardic sets himself apart from virtually all other philosophers of science in insisting that measures of the causal factors giving rise to trait difference can and do inform our understanding of the causal dynamics of individual trait development. Unfortunately, however, his argument rests on a routine confounding of the two questions. Thus, in response to David Lykken's claim that "it is meaningless to ask whether Isaac Newton's genius was more due to his genes or his environment, as meaningless as asking whether the area of a rectangle is more due to its length or its width" (Lykken 1998, quoted in Sesardic 2005, 55), Sesardic writes: "Contrary to what [Lykken] says, however, it makes perfect sense to inquire whether Newton's extraordinary contributions were more due to his above-average intellect or to an above-average stimulating intellectual environment." And in response to what he refers to the "mistaken conclusion" by the Nuffield Council on Bioethics — namely, that "it is vital to understand that neither [the broad or narrow] concept of heritability allows us to conclude anything about the role of heredity in the development of a characteristic in an individual" — he writes: "On the contrary, if the broad heritability of a trait is high, this does tell us that any individual's phenotypic divergence from the mean is probably more caused by a non-standard genetic influence than by a non-typical environment."[6]

What I find curious in all these discussions is that, although the conflation is not restricted to genetics, it does seem both less common and less automatic when talking, say, about drums and drummers. In the latter context, the conclusion that it is possible to decompose a phenotypic difference even when one can say nothing about the relative strength of the causal inputs to the trait itself seems considerably less counterintuitive; indeed, I doubt that anyone would be tempted to infer anything at all about the causal contribution of a drum to the sound that was produced from my hypothetical conclusion that 40 percent of the difference between sounds was due to a difference between drums. So what is it that especially confounds our thinking when we turn our attention to the subject of genetic influences?

Genes as Difference Makers

My claim that an important source of confusion lies deep in the language of genetics requires us to look at the evolution of that science. What precisely is genetics, as a branch of biology, about? How we answer this question of course depends on when, and where, we look. But for classical genetics, especially the paradigmatic school of T. H. Morgan, genetics was about tracking the transmission patterns of units called genes. What was a gene? No one knew, but notwithstanding this ignorance, a gene was assumed to be a unit that could be identified by the appearance of mutants in wild-type populations. That is, a phenotypic difference in some trait (a mutant) was taken to reflect a difference (a mutation) in some underlying gene associated with that trait. But to argue for such an identification between phenotypic difference and underlying gene requires a two-step move. First, change (a mutation) in some underlying entity (the hypothetical gene) is inferred from the appearance of differences in particular phenotypic traits (e.g., blue eye, bent wing, narrow leaf), and second, the existence and identity of the gene itself is inferred from the inference of a mutation. In other words, the classical gene was on the one hand identified *by* the appearance of phenotypic differences (mutants), and on the other hand simultaneously identified *with* the changes (mutations) that were assumed to be responsible for the mutants. Thus, the first map of the mutations thought to be responsible for the observed phenotypic differences in *Drosophila* was called not a map of mutations, but a genetic map, a map of genes.

This is the sense in which the classical gene is often said to be a "difference maker" (see, e.g., Sterelny and Griffiths 1999; Moss 2003). But a gene was taken not only to be a difference maker; it was also assumed to be a trait maker. It was both the entity responsible for the difference observed, and (at least implicitly) the entity responsible for the trait which has undergone a change — i.e., the trait in which a difference has been observed.[7]

We might say, then, that a certain confounding of traits and trait differences was built into the science of genetics from the very beginning;

moreover, we might argue that such a confounding inhered in the logic of their method. The occurrence (and frequency) of trait differences was what geneticists had observational access to: by examining phenotypes, they could detect phenotypic differences—which, in turn, were taken as indicative of changes in some underlying, internal entity. Through breeding, the locus of such changes could be mapped. As Horace Freeland Judson has observed: "In 1913, Alfred Sturtevant, a member of Thomas Hunt Morgan's fly group at Columbia University, drew the first genetic map—'The linear arrangement of six sex-linked factors in *Drosophila*, as shown by their mode of association.' Ever since, the map of the genes has been, in fact, the map of gene defects" (2001, 769). Similarly, much of the same slippage can be identified in the use of the term "gene" in what Lenny Moss (now along with numerous others) calls "Gene-P," or "phenotypic gene," where the unit of interest is clearly a difference maker. The actual referent of Gene-P is thus not a gene but a genetic mutation.[8] Wilhelm Johannsen, the man to whom we owe the word *gene*, was himself clearly worried about this problem when he asked: "Is the whole of Mendelism perhaps nothing but an establishment of very many chromosomal irregularities, disturbances or diseases of enormously practical and theoretical importance but without deeper value for an understanding of the 'normal' constitution of natural biotypes?" (1923, quoted in Moss 2003, 62).

It is hard to imagine that the early slippage was entirely accidental. To think of genes simply as difference makers would have been to detract from the power of the gene concept. Mapping difference makers and tracking their assortment through reproduction may have been all that the techniques of classical genetics could do, but the aims of these scientists were larger. What made genes interesting in the first place was their presumed power to mold and form—in a word, their presumed power to act. *Gene action* was the term used for the process by which genes exerted their power in the development of characters or traits. But to illuminate the nature of this process (the developmental process), studies of trait differences alone would not suffice. In fact, neither mapping of the locus of the factors (difference makers) presumed to be responsible for such differences nor analysis of their emergence and intergenerational patterns of transmission

taught us anything about the causal dynamics of the developmental process by which the traits themselves came to be. As John Dupré puts it, "classical genetics was about invisible features that could trigger different developmental outcomes, but not about the causal explanation of these outcomes" (2006, 118). Furthermore, classical geneticists were for the most part well aware of this distinction. Nevertheless, the easy slide between genes as difference makers and genes as trait makers perpetuated the illusion (as widespread among geneticists as it was among their readers) that an increased understanding of the effects of gene differences would enhance our understanding of what it is that the entities called genes actually do.

What Is a Disease?

A similar confounding of the etiology of traits with that of trait differences pervades virtually all of the current literature of medical genetics. Indeed, the very notion of a disease as an individual trait—in the sense that brown eyes, say, is a trait—already incorporates this confusion. We may commonly speak of an individual as "having a disease," much as he or she might have brown eyes, but as many writers have long understood, disease is a state only in relation to another state that has already been established as normal.[9] In his inquiry into the scientific rules for distinguishing the normal from the pathological, written more than a hundred years ago, Émile Durkheim stressed that "a trait can only be characterized as pathological in relation to a given species" ([1894], 4; my translation)—in other words, in relation to a standard of normality or state of health which is, itself, inextricably confounded with the norm of a species. He continued: "One cannot even conceive, without contradiction, a species that could, by itself and in virtue of its fundamental constitution, be irremediably sick. [The species] is the norm par excellence and, accordingly, can harbor within itself nothing of the abnormal."[10]

It is true that the French philosopher of science Georges Canguilhem, following Kurt Goldstein, made valiant efforts to internalize the diagnostic criteria of pathology, locating them within the individual,[11] but for all

his efforts, by far the most common understanding of disease has continued to rely on comparison (or contrast) with a preestablished conception of the normal. English-language dictionaries routinely define disease as a relational state: it is a dis-ease, "an abnormality of the body or mind" (Wikipedia); "a departure from the state of health" (*Oxford English Dictionary*); "a deviation from or interruption of the normal structure or function of a part, organ, or system of the body" (*Dorland's Illustrated Medical Dictionary*); "an interruption, cessation, or disorder of a body, system, or organ structure or function" (MediLexicon). And indeed, French dictionaries do the same. In virtually every dictionary I have consulted, a *maladie* is an *altération de l'état de santé*. In French as in English, an animal can be said to have brown eyes whether or not a comparison with other animals is at hand, but it cannot, without such a comparison, be said to have a disease.

Like medicine, genetics too might be said to be a comparative science. Comparing organisms with differing phenotypes, along with attempting to correlate these phenotypic differences with corresponding genetic differences (mutations), has been the bread and butter of geneticists from the earliest days of that science. But genetics aims beyond comparative judgments, seeking an understanding of the developmental dynamics. And, as I have tried to show, its language invites us to lose sight of the complex moves—first, attributing the cause of a phenotypic difference to a genetic mutation; second, assuming that the presence of a mutation automatically signals the presence of a gene; and third, attributing responsibility for the trait in question to the gene in which the mutation is assumed to have occurred—that are routinely made in effecting this shift from comparative to individual. It seems no accident, therefore, that the adoption of a lexicon of illness referring to disease as an individual attribute comes with the emergence of a medical science grounded in genetics.[12] Medical genetics is a contemporary science aimed at identifying the causes of "inborn errors of metabolism" (i.e., of disease), and it is a direct heir to the more general efforts of classical genetics to associate mutations with particular phenotypic differences. And the main difference between the two kinds of efforts is that, where classical geneticists sought

the causes of phenotypic difference in mutations that were for a long time largely hypothetical, the genetic defects that medical geneticists identify as the causes of a disease state can be clearly specified. The main similarity, on the other hand, is that they are subject to exactly the same sorts of conflation. A disease, we need to remember, is not a trait but a trait difference, and the causal factors involved in producing the former are not the same as those involved in producing the latter.

Like the classical geneticists before them, today's medical geneticists are wont to slide from the identification of a genetic difference (a mutation, a departure from a presumed normal genome) associated with a disease state to assertions of having identified a gene responsible for the disease in question. Where classical geneticists relied on "gene maps" to identify the gene presumed to be involved, today medical geneticists are more likely to rely on analyses of nucleotide sequences. They may continue to talk about *genes*, but it is the DNA itself that has become the focus of investigation. The direct object of much actual research is the identification of any change (or defect) in the nucleotide sequence that is correlated with the expression of the disease. Despite widespread talk of "disease genes," or "disease-causing genes," the information these analyses provide is about the causal implications of mutations, and not about the causal role of genes. Indeed, the notion of a gene "causing" a disease (or even of a particular sequence "causing" a disease) has exactly the same status as the notion of a gene "causing" a phenotypic difference.

Furthermore, and perhaps more importantly, the identification of one or more anomalies in the DNA sequence may or may not point to a defect in a particular gene (however that term is defined). Indeed, for diagnostic purposes (at least for diagnoses based on genetic tests), the attempt to correlate a disease state with an underlying gene may in many if not most cases be largely irrelevant. Contemporary medical genetic diagnoses rely on the identification of aberrant or anomalous sequences, and not on the causal pathways such anomalies may disrupt. Also, the anomalies may be anywhere in the genome—in fact, only rarely are they found in protein-coding sequences (that is, in the segments of DNA usually associated with genes).

For most of us, the crucial questions are, can the identification of such sequences be useful in the treatment or prevention of disease, and if so, how? Most immediately (and perhaps most obviously, even if generally left unstated), such information can be used to promote selective abortion. But if we are interested in therapeutic medicine, we need more than simple correlation between aberrant sequence and aberrant phenotype. It is true that the early days of the Human Genome Project brought the promise that in time we would be able simply to replace defective sequences with normal ones (gene therapy). However, that hope has so far failed to materialize, and at least one of the reasons is that the relation between DNA sequence and phenotype turned out to be far more complicated than originally expected. As to the possibility of other kinds of treatment or prevention in a particular individual carrying the aberrant sequence, this depends on understanding something about the biological function that has been disrupted by the identified change in sequence. Such a quest takes us beyond the analysis of phenotypic differences induced by mutant forms. Indeed, it requires an altogether different kind of analysis, almost always one of a far more difficult nature.

There are, however, some relatively simple exceptions. Phenylketonuria (PKU) is one, and it is probably the most celebrated case of therapeutic intervention in the history of medical genetics. It is everyone's canonical example, mine as well. PKU is a disorder (now recognized as genetic) associated with a range of disabling symptoms, including mental retardation, and it is caused by the inability of the body to properly metabolize the essential amino acid phenylalanine. A major breakthrough in the treatment of this disease came with the recognition that its symptoms can be significantly alleviated if the affected individual adheres to a carefully monitored low-phenylalanine diet for his or her entire life.[13] However, the development of a strategy to treat PKU had nothing to do with either the identification or the mapping of the gene(s) or genetic sequence(s) involved. Today we know that this disorder is caused by one or more mutations in the gene encoding the enzyme that breaks down phenylalanine (to date, as many as 400 such mutations have been identified). In point of historical fact, however, neither our understanding of the (at least proximal) cause of this disease (Jarvis 1937) nor the development of a therapeutic intervention (Bickel et al. 1954) depended on any sort of genetic analysis. It is now believed that the phenotypic expression of PKU (the disabling symptoms) are the direct consequence of the accumulation both of high levels of phenylalanine and of toxic intermediates resulting from its faulty metabolism, but we learned this from direct biochemical analysis; ordinary medical observation showed how dreadful the symptoms can be. Furthermore, we have now acquired the ability to precisely characterize many of the mutations responsible for the absence of the necessary enzyme, and doing so has certainly been instructive. The bottom line, however, is that, thus far, that ability has not significantly added to the possibilities of therapeutic intervention.

Of course, it need not have happened that way. Identification, location, and characterization of the guilty mutation(s) could have provided the starting point for a research program that, over time, led to an understanding of the disease and possibilities for treatment. But whatever the sequence of events, there is no way in which the genetics of difference alone could have produced that understanding. What was required for a causal analysis of the disease—and, hence, for the possibility of therapeutic intervention—was a biochemical analysis of the metabolic pathway that had gone astray. That analysis *might* have begun with the identification of a particular gene, but if it had, what would have been required is an understanding of the gene's downstream effects, of the particular role actually played in development by the gene in which the mutation had occurred.

What Genes Do

The challenge of attributing causal function to genes (rather than to gene differences)—of understanding what genes do—has plagued genetics from the onset, and it was not until the advent of molecular biology that it seemed possible to meet this challenge. What does a gene do? Defined as a discrete stretch of nucleotides, a gene was said to "make," or code for, a protein. But even with that insight, it was one thing to be able to track the

causal effect of a genetic difference (mutation, or change in nucleotide sequence) on a particular trait, and quite another to track the causal influence of the gene itself (or indeed of the protein it was said to "make") on the development of that trait.

To be sure, enormous progress has been made since the early days of molecular biology, and we now know a great deal more about the ways in which the process of development makes use of an organism's DNA. But what we have learned has not so much answered earlier questions as it has transformed them. We have learned, for instance, that the causal interactions between DNA, proteins, and trait development are so entangled, so dynamic, and so dependent on context that the very question of what genes *do* no longer makes much sense. Indeed, biologists are no longer confident that it is possible to provide an unambiguous answer to the question of what a gene *is*. The particulate gene is a concept that has become increasingly ambiguous and unstable, and some scientists have begun to argue that the concept has outlived its productive prime.

DNA, by contrast, is a concrete molecule—an entity that can be isolated and analyzed, identified by its distinctive physical and chemical properties, and shown to consist of particular sequences of nucleotides. We not only know what DNA is, but with every passing day, we learn more about the exceedingly complex and the multifaceted role it plays in the cellular economy. It is true that many authors continue to refer to genes, but I suspect that this is largely due to the lack of a better terminology. In any case, continuing reference to "genes" does not obscure the fact that the early notion of clearly identifiable, particulate units of inheritance—which not only can be associated with particular traits, but also serve as agents whose actions produce those traits—has become hopelessly confounded by what we have learned about the intricacies of genetic processes. Furthermore, recent experimental focus has shifted away from the structural composition of DNA to the variety of sequences on DNA that can be made available for (or blocked from) transcription—in other words, the focus is now on gene expression. Finally, and relatedly, it has become evident that nucleotide sequences are used not only to provide transcripts for protein synthesis, but also for multilevel systems of regulation at the level

of transcription, translation, and posttranslational dynamics. None of this need impede our ability to correlate differences in sequence with phenotypic differences, but it does give us a picture of such an immensely complex causal dynamic between DNA, RNA, and protein molecules as to definitely put to rest all hopes of a simple parsing of causal factors. Because of this, today's biologists are far less likely than their predecessors were to attribute causal agency either to genes or to DNA itself—recognizing that, however crucial the role of DNA in development and evolution, by itself, DNA doesn't *do* anything. It does not *make* a trait; it does not even encode a program for development. Rather, it is more accurate to think of DNA as a standing resource on which a cell can draw for survival and reproduction, a resource it can deploy in many different ways, a resource so rich as to enable the cell to respond to its changing environment with immense subtlety and variety. As a resource, DNA is indispensable; it can even be said to be a primary resource. But a cell's DNA is always and necessarily embedded in an immensely complex and entangled system of interacting resources that are, collectively, what give rise to the development of traits. Not surprisingly, the causal dynamics of the process by which development unfolds are also complex and entangled, involving causal influences that extend upward, downward, and sideways.

tend to be more focused on individual traits than are those of the farmer. The obvious question is, how can scientific research help us? To the extent that the choices we make depend on the social, political, and even moral values we hold, more—or better—scientific data will not be of much use. But on some issues, scientific research can clearly help us set priorities. In particular, it can (at least in many cases) shed light on how effective any proposed intervention is likely to be. The question then becomes, what kind of research can give us this kind of information?

Because I believe that, even without the many problems I have discussed in this book, heritability measures are an extremely indirect way to go about providing such information, my preference would be to reformulate the question altogether, changing it to one that directly focuses on what I think people want most immediately to know, and also one that the realities of biological development permit scientists to actually answer. Let us ask not how much of any given difference between groups is due to genetics and how much to environment, but rather how malleable individual human development is, and at what developmental age. As I have already said, there is no reason to privilege birth as a cutoff point—development is lifelong, and so is its plasticity. We may not share the interests of breeders in artificial selection, but as both scientists and citizens, we are surely committed to trying to maximize the development of individual human potential. And for this, we need a better understanding of what resources can contribute to such development, and of how they can best be deployed. What kinds of research can provide us with such information? I would put my money on the new studies of phenotypic plasticity that we are beginning to see not only in developmental biology, but also in neuroscience, physiology, and ecology. There is no shortage of scientific work that can productively inform us about the things we want to know, but we do need to pose our questions in ways that researchers can answer.

NOTES

Introduction

1 For a particularly valuable review of such alternative "systems of inheritance," see Jablonka and Lamb, 2005.

2 Especially with recent discoveries of the many different roles played by non-coding DNA in a cell, the definitions of *gene* have proliferated wildly (see, e.g., Keller and Harel 2007).

ONE Nature and Nurture as Alternatives

1 For the Confucian literature, see, e.g., Graham 1986; Ames 1991; Ng 1994; Bloom 1994 and 1997; and Xunwu 2000. For a more general and comparative discussion of ancient Greek and Chinese views of the matter, see Lloyd 2007. Particularly striking is Graham's claim that "analogies for human nature in Mencius are always dynamic, trees growing on a denuded mountain, ripening grain, water finding its natural channels" (Graham 1986, 43). Lloyd's discussion is framed in terms of nature and culture, and his conclusion is to the point: "The nature/culture division turns out to be a far more problematic framework than is sometimes realized for the discussion of the views in question, namely about what there is and about the different categories into which that falls" (2007, 132–33).

2 I take the phrase "nurture nature" from Ng's translation of *yang-hsing* in his discussion of the concepts of nature and nurture in Confucian philosophy (Ng 1994); see also Chen, who argues that the notion of culture as "the art of cultivating nature" (2000, 109) lies at the core of Confucianism.

3 John Locke, March 7, 1692 [i.e., 1692/93], "Some Thoughts Concerning Education," http://www.fordham.edu/halsall/mod/1692locke-education.html.

the *Oxford English Dictionary*), the particular kind of distinction that Galton has in mind ought, given his implication of separability (and hence mutual exclusivity), more accurately be referred to as a *disjunction*.

16 I should make clear that in my discussion of "particulate" inheritance, there is no implication of the transmission of "particulate" traits (for instance, there may, as both Darwin and Galton assumed, be many particles that combine to form each trait), and hence there is no implication of an opposition to what is referred to as "blending inheritance." As I am using the term, *particulate inheritance* refers merely to the transmission of discrete material units assumed to provide the causal basis of the development of traits.

17 Galton's suggestion here is reminiscent of Spencer's (1864) model of development as a kind of crystallization.

18 Galton refers explicitly to advances in physical science to support his conviction of the need of a particulate formulation (1876b, 332). See Stanford (2006) for an especially insightful analysis of Galton's particulate views.

19 My argument is not meant to contradict, but rather to complement, Waller's claim for the importance of Galton's methodology to his thinking about heredity (see Waller 2002b).

20 The issue of the invariance of the causal elements is certainly crucial to the debate over the question of the inheritance of acquired characters that later loomed so large, and here too Galton's views departed from those of his cousin. Indeed, Galton's insistence on the invariance of the germ material might be said to have anticipated August Weismann's arguments. But for the present discussion, the more relevant point is that such an assumption greatly supported (as it was supported by) Galton's commitment to the separability of hereditary factors from environmental influences.

21 See Waller 2001; other sources include Rosenberg 1974; Olby 1985; and Bowler 1989.

22 It is likely that Galton's general anxiety on this score was affected by the failure of his own marriage to produce offspring, as well as by the sterility of the marriages of two of his brothers.

23 The shift in terminology from "nature versus nurture" to "heredity versus environment" is of interest in itself and certainly needs further investigation. It is clear, however, that that shift was substantially encouraged by the resonances of the latter phrase with a broad range of discussions of the time, in many different countries (e.g., England, France, Germany, and the United States) in developmental psychology on the one hand, and about race and milieu or environment in anthropology and anthropometry on the other hand. By the end of the nineteenth century, virtually all of these discussions had come to be framed in terms of the relative importance of heredity and environment, and references to Galton's methods as a way of resolving disputes on the matter were common. As Franz Boas explained: "If it can be shown that there is a strong tendency on the part of the offspring to resemble the parent, we must assume that the effect of heredity is stronger than that of environment. The method of this investigation has been developed by Francis Galton and Karl Pearson . . . and we look forward to a definite solution of the problem of the effect of heredity and of environment through the application of this method" (1899, 101).

24 Quoted in Pinker 2004, 8.

TWO Changing the Question

1 This is more an empirical (and practical) issue than a logical one, for it depends on whether gene-environment interactions make a significant difference over the range of environmental parameters deemed relevant. I will discuss this issue at greater length in chapters 3 and 4, but for now, it is worth mentioning that the question of the magnitude of this statistical effect has historically been deeply intertwined with the interdependence of genes and environment in development. For an illuminating discussion of just how this intertwining played out in the debates between Fisher and Hogben, see Tabery 2007 and 2008.

2 Such expectations may of course continue to lurk in the background; indeed, the suspicion that they do accounts for at least some of the political charge attached to these questions.

3 http://en.wikipedia.org/wiki/Nature_versus_nurture, accessed February 8, 2009.

4 To be fair to Woodward, it should be pointed out that he is concerned with only a very limited meaning of "*x* causes *y*," namely, that "*x* is causally relevant to *y*," and if this is all one wants to infer, I have no argument with the replacement of such a claim by the claim that "a difference in *x* causes a difference in *y*" (Woodward 2003, 40). The problem arises when one takes "*x* causes *y*" in the stronger sense required for a measure of the relative importance of *x* as one of two or more causal factors collectively responsible for *y*.

5 In an earlier work (Keller 2000, chapter 4), I refer to the tacit reasoning behind such inferences as a "logic of subtraction"—likening such an effort to the attempt to infer the operation of a car from all the ways in which the car can be caused to malfunction.

6 See chapter 3 for a discussion of the terms *broad* and *narrow heritability*.

7 Judson proposes as a solution to this problem that we "revive and put into public use the term 'allele.' Thus, 'the gene for breast cancer' is rather the allele, the gene defect — one of several — that increases the odds that a woman will get breast cancer" (2001, 769). It seems to me, however, that the equation of *allele* with *gene defect* risks perpetuating precisely the same confusion: it is not the allele itself that is responsible for the phenotypic difference, but the difference between alleles.

8 Moss writes that "Gene-P" is a "phenotype predictor" and cannot be defined by its nucleic acid sequence. But in fact, "Gene-P" is neither a gene nor a predictor of phenotype: it is a phenotype difference predictor. Indeed, as he himself acknowledges, the reason that "Gene-P" cannot be defined by a specific sequence is that "invariably there are many ways to lack or deviate from a norm" (Moss 2003, 60).

9 Of course, the concept of normal is itself fraught with difficulty, subject to its own ambiguities that primarily have to do with persistent confusion between properties of individuals and those of populations. But most commonly, it too is understood as a relational property, pertaining not to comparison between individuals but to the statistical norm of a population. For instance, Durkheim wrote: "The state of health, insofar as it can be defined, never conforms exactly to that of an individual subject, but can only be established in relation to the most common circumstances" [L'état de santé, tel qu'elle le peut définir, ne saurait convenir exactement à aucun sujet individuel, puisqu'il ne peut être établi que par rapport aux circonstances les plus communes ([1894], 62; my translation)]. See also Hacking 1990, 160–64.

10 The full quotation in the original reads as follows: "On voit qu'un fait ne peut être qualifié de pathologique que par rapport à une espèce donnée. Les conditions de la santé et de la maladie ne peuvent être définies in abstracto et d'une manière absolue. La règle n'est pas contestée en biologie; il n'est jamais venu à l'esprit de personne que ce qui est normal pour un mollusque le soit aussi pour un vertébré. Chaque espèce a sa santé, parce qu'elle a son type moyen qui lui est propre, et la santé des espèces les plus basses n'est pas moindre que celle des plus élevées . . . Le type de la santé se confond avec celui de l'espèce. On ne peut même pas, sans contradiction, concevoir une espèce qui, par elle-même et en vertu de sa constitution fondamentale, serait irrémédiablement malade. Elle est la norme par excellence et, par suite, ne saurait rien contenir d'anormal" (http://classiques.uqac.ca/classiques/Durkheim_emile/).

11 E.g., in *Essai sur Quelques Problèmes Concernant le Normal et le Pathologique* (Essay on Several Problems Concerning the Normal and the Pathological; 1943), Canguilhem writes: "We think with Goldstein that the norm concerning pathology is above all an individual norm" (72; my translation). See also his essay, "Le Concept et la Vie" (1966).

12 For an extensive discussion of the relation between the languages of medical genetics and classical genetics, see Childs 1999.

13 This is not to suggest that maintaining such a diet is easy, or that the almost inevitable relapses are not without dire risks. Probably the best discussion of the history and politics of PKU, as well as of risks associated with its treatment, is to be found in the work of Diane Paul (see, e.g., Paul 1998b and 2000, and Paul and Edelson 1997).

THREE From Individuals to Populations

1 In particular, I might mention an article by Scott Stoltenberg (1997) that I came upon after writing this chapter. Stoltenberg's work covers some of the same ground as I do, and is especially valuable for the many examples it offers of the different uses of the terms *heritable* and *heritability*.

2 Variance is the statistical measure of the average deviation of a variable in a distribution from its mean value.

3 See, e.g., Jencks 1992; Block 1996; Sober 2001.

4 Measuring this ratio is hardly a trivial feat. One way that is often used in work with experimental organisms is to assign a range of values to a particular phenotypic trait, breed a population of pure lines (homozygous genotypes), cross them (in pairs) to obtain a set of heterozygous genotypes, and measure the phenotypic variance of the trait within a population of identical heterozygous genotypes growing in different environments. Because there is no genetic variation in these populations, the phenotypic variance is assumed to be the same as the environmental variance. This value (averaged over the different heterozygote populations) can then be subtracted from the phenotypic variance of the original population (of pure lines) to obtain an estimate of the genotypic variance in that population. Alternatively (and virtually the only way available in studies of human populations), estimates are obtained by comparing the correlations of a trait in monozygous twins raised by their biological parents with the correlations observed in dizygotic twins, also raised by their biological parents. Because monozygotic twins share (almost) all of their genes, and dizygotic twins share only half, the difference between the correlations is presumed to give a measure of heritability. But to the extent that behavioral responses to children is determined by cultural expectations associated with such characteristics as skin color, eye color, and height, monozygotic twins also share more of their environment than do dizygotic