1. Decomposition and Localization in Perspective

Our focus has been on decomposition and localization and their role in the development of scientific research programs. We have especially emphasized their heuristic role in the construction of explanatory models within problem domains that are ill structured. In the cases we have discussed there was initially no well-delineated space of explanatory models, nor even a clearly defined range of phenomena to be explained. The problem-solving tasks were, then, exploring and constructing the space of explanatory models, and determining the precise range of phenomena to be explained. These tasks are interdependent, their relationship neither simple nor uniform. As Ian Hacking (1983) observes, the relationship between theory and experimental "data" can change with the development of a theory and varies from one science to another. Observations may sometimes be undertaken with no particular theoretical motivation; however, the phenomena one then finds may have no clear significance, for lack of a theory to make sense of them. Alternatively, the observations' significance may become clear only after the fact. In other cases, a program of experimental investigation and observation is undertaken with clear theoretical motivations. In these instances the implications may be clear and immediate. However, on occasion the significance changes and is far from what the initial experimenter thought it would be.

It would be a mistake, from our current vantage point, to try to fit the full range of cases into a single mold. In the cases we have discussed, research programs vary greatly in both motive and method. Some motives were clearly practical: the investigation of fermentation, for example, mattered for the development of wine and beer industries, and research in genetics was not unrelated to agriculture. Other motives were more theoretical: the importance of catalytic reactions to respiration could not be clear in the absence of Lavoisier's view that respiration was a slow combustion. Even the relationship of fermentation and respiration was unclear. At times it was assumed that there should be a single, general theory encompassing both; at other times they were treated in relative isolation. Likewise, it was not always clear what the significance of linguistic or cognitive deficits was for a theory of neural and cognitive functioning. Gall and Spurzheim rejected appeals to deficits on the grounds that work; Jackson acknowledged their importance, but drastically reinterpreted them; and in contemporary research their significance is still a matter of debate. In part, decomposition and localization suggest a form for explanatory models; in part, they serve to impose a structure on the phenomena to be explained. Exactly how they affect the development of explanatory models depends on the theoretical context and the available experimental techniques.

We have described the explanatory models generated by these heuristics as mechanistic in character. This is meant to suggest a number of features common to the cases we have discussed. First of all, the models do feature in what are naturally thought of as *causal explanations*. Models describing the structure of the genetic material are meant to explain development and inheritance in terms of underlying causes. For example, the presence or absence of specific enzymes explains respiration in these terms. However, laws and theories classically construed have little place in the picture. The problems guiding the sort of research we have recounted are (1) explaining how some particular effect is actually produced, here and now, and (2) by what means. The resulting explanations are sometimes what Nancy Cartwright (1983), following John Stuart Mill (1884), calls "explanation by composition of causes." She says this "picture of how nature operates to produce the subtle and complicated effects we see around us is reflected in the explanations that we give: we explain complex phenomena by reducing them to their more simple components" (1983, p. 58). We describe the effect on a body, say, as the consequence of gravitational and electrical forces acting together. Likewise, we describe our linguistic competences in terms of the effect of a multitude of capacities, and we explain phenotypic traits as the consequence of the influence of genetic and environmental contributions. We explain a phenomenon that interests us by identifying a variety of causal factors and showing how they conspire to yield an effect.

Second, the explanatory models appeal to what are naturally thought of as underlying mechanisms. Fermentation is something done by cells; the mechanism is a complex biochemical pathway. Development is something organisms undergo; the mechanism that regulates it is a complex genetic pathway. Language or cognition is something attributable to people; the mechanism is given in neurophysiological—or, perhaps, functional—terms. Thus, the models we have looked at are tied to more than one explanatory level. We shift down from the system to its parts in order to explain how the system does what it does. The models are related to what Darden and Maull (1977) call *interfield theories* (see also Darden 1986). A theory of fermentation must explain fermentation in terms of the biomechanisms, whether psychological or neurophysiological, that mediate behavior. This has important implications for the dynamics of the resulting theories.

We have also focused on what we call *models* rather than formally developed theories. This is, in part, because laws—general and abstract explanatory principles—play a minor role. Traditional views have held that events are explained by showing they are to be expected as the result of natural laws.[1] More recently philosophers have moved away from an emphasis on laws, understood formally. This is sometimes described as a *semantic view of theories*. Models are introduced as intermediate structures, somewhere between laws in a classical sense and data. Bas van Fraassen explains it this way:

> To present a theory is to specify a family of structures, its *models*; and secondly, to specify certain parts of those models (the *empirical substructures*) as candidates for the direct representation of observable phenomena. (1980, p. 64)

Models here are abstract structures; they can be thought of as the (nomologically) possible worlds of which a theory is true. Laws are yet more abstract, leaving specific parameters and parameter values open. In Ronald Giere's (1988) treatment of classical mechanics, the explanatory models are constructed from Newton's laws, though not by any simple derivation; rather, they are developed by fixing some parameter values and introducing a number of simplifying assumptions and approximations. As a consequence the models do not precisely fit any real situation, but to the extent that they approximate real situations, they can be used to understand them.

Some philosophers who have focused on evolutionary biology have defended a similar perspective. In particular, John Beatty (1980, 1981), Elizabeth Lloyd (1984, 1986, 1988, 1989), and Paul Thompson (1983, 1989) have argued that evolutionary biology, population genetics, and evolutionary ecology are better represented in terms of a semantic view. Insofar as we emphasize models, we are not unsympathetic to semantic views. In the cases we have described, however, we find even less of a role for laws. Construing these models as constructed from more abstract laws by fixing parameter values, even with additional simplifying assumptions and approximations, if practicable at all, would be little more than an idle exercise, and one accomplished after the fact. The explanatory task begins and ends with the models.[2] Models in this sense are akin to blueprints; they are partial and abstract representations of the causal mechanisms. A sometimes by appealing to analogies with other more familiar mecha-nisms), and conceptualizing how such parts will interact with each other. The explanatory power of a model stems from its ability to show how some phenomenon or range of phenomena would be the consequence of the proposed mechanism.

By developing a model of mechanistic explanation, our goal has not been to construct a general theory of explanation, or even of mechanistic explanation. Neither has it been to offer any general conception of theories. We do claim that mechanistic models are often the vehicles of explanation in the biological and psychological sciences and that they often constitute what scientists count as theories. But our focus has been on the construction of causal explanations; that is, on the development of mechanistic models. In a suitably broad sense this can be understood as a concern with the development of theories. As we suggested in Chapter 1, it is in the dynamics of theories that we can hope to find an account of scientific discovery.[3] We have so far focused on the kinematics of change. We have described a number of decisions, or choice-points, that define the direction of development. These (illustrated in Figure 8.7) are less chronological than logical choice-points.

Two of the choice-points particularly involve decomposition and localization. The first, which occupied us in Part II, is the identification of a system as the locus of control and the determination of the functional properties of components within the system. Identifying the cell as the locus of control for respiration resolved itself into the question whether respiration was carried out only in specialized loci (for example, the lungs or the blood), or in the tissues as well. As we explained in Chapters 3 and 4, resolving this issue was partly an experimental question and partly a question shaped by more general metaphysical and theoretical commitments. The experimental questions focused on the capacities of specific tissues in isolation. The more general commitments ranged from a commitment to mechanistic explanations—and especially to an antivitalism assumed by most of the major contenders—to more specific commitments shaped by thermodynamic and chemical theories of the time.

The second step, which occupied us in Part III, involves the analysis of the activities of the controlling system into component functions. As with identifying the locus of control, this decomposition must be combined with a localization that resolves the system into appropriate units and assigns relevant functions to them. We have seen that a number of outcomes are possible, ranging from simple localization of a function in a component of the whole system, to models emphasizing interaction of components. If

Once again the specific outcome depends only partly on experimental results; it also incorporates an array of more general commitments. Successful mechanistic explanations require convergence of the decomposition into component functions with successful localization of these components within the system. Developing such convergence often relies on a wide array of experimental and theoretical constraints. These can range from broadly correlational constraints, as in Chapter 6, to more restrictive physical constraints, as in Chapters 7 and 8. At both the point where one chooses to segregate a system from its environment and attribute an activity to it, and the point where one tries to decompose the function and localize it in components of the system, it is possible to seek out alternatives to mechanistic explanation. While we have discussed some of these alternatives, our primary focus has been on the development of mechanistic models.

While we have partially succeeded if the choice-points we have identified provide an accurate account of the kinematics of theory development, our ultimate goal is to go further and develop a dynamic analysis of theory development. This requires us to identify the constraints that play the role of forces governing theory development. The kinematics constitute the phenomena that a dynamic model should explain.

2. Four Constraints on Development

The cases we have discussed can provide a rough taxonomy of the array of constraints that affect the development of mechanistic models; all can be encapsulated in a dynamic model of theory development. This model is speculative and provisional. The constraints we identify, though, are not only consistent with, but are suggested by, the more descriptive, kinematic picture. They fall into four basic groups, summarized in Figure 10.1. The first are *psychological* constraints, including the heuristic strategies employed. If we are right, localization and decomposition are among them. The second are what we will describe as *phenomenological* regularities. These consist of, primarily, information concerning the behavioral capacities of the system, including the effects of perturbations. The third are *operational* constraints, which encompass the experimental procedures and models that limit what can be asked about the behavior of components. The fourth are *physical* constraints, which limit the range of allowed component functions. In what follows we will say a bit more about these four constraints and their role in a more general model of theory change.

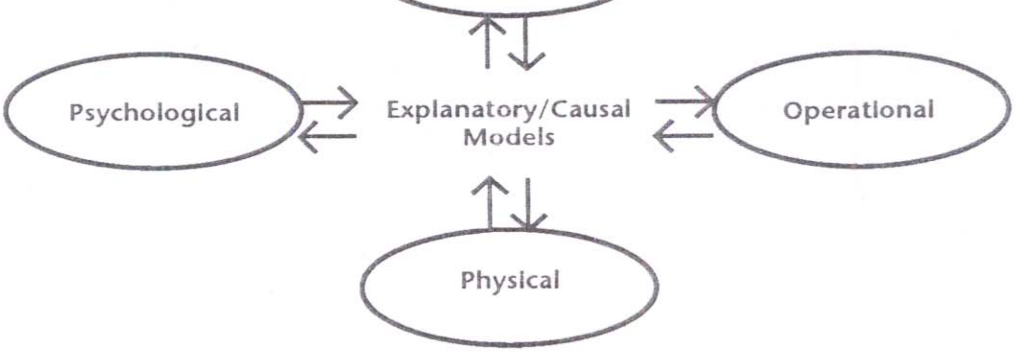

Figure 10.1. Four Constraints on Explanatory Models. There are at least four constraints influencing the development of causal and explanatory models. This provides at most a rough indication of the general categories included and exemplified in one or another of the cases discussed in previous chapters.

Psychological Constraints

Psychological constraints are largely heuristic, in the sense sketched in Chapter 1. Decomposition and localization are heuristics that we have identified as applicable to the development of mechanistic explanations of complex systems. They serve, at least, to simplify the explanatory problem. They are domain specific, and while they are not unconditionally conducive to realistic models, they sometimes lead to realistic causal explanations. One reason decomposition and localization may fail is that the assumptions they impose—that the system is decomposable or nearly decomposable—may be false. Some systems are decomposable, and some are nearly so; others are less so, though assuming that they *are* decomposable is one way to discover that they are not. Insofar as these heuristics are domain specific, they are also liable to error. We saw in Chapter 9 that there is at least some reason to think that decomposition and localization do *not* reveal the actual organization of some systems, neurophysiological, psychological, and genetic. Decomposition and localization are relatively restrictive and so quite powerful. That is, they significantly limit the search space of possible explanatory models to ones with reasonably simple organization. Accordingly, they provide a relatively strong guide on the development of explanatory models. When they fail, it is necessary to expand the range of explanatory models under consideration.

The critical point is that these heuristics, like others, reflect cognitive strategies through which humans approach a complex problem. They are necessary in part because human information-processing capacities are is assumed to be a well-defined space of possible solutions, and the task is to search this space to find the equation or law that solves the problem. If we assume a well-defined and adequate search space, then all discovery requires is a means to search through the space of allowed solutions. Accordingly, BACON searches a limited space of mathematical relations, and DENDRAL assumes a determinate set of organic molecules. Even in a search of a well-defined space, though, the search time might exceed the information-processing capacities of humans or machines, and heuristics are used to focus the search on the most promising parts of that space.

In the cases we have considered there is no antecedently specifiable set of mechanisms, and hence no well-defined space to search for possible mechanisms. Indeed, we do not even know whether the space we are searching contains a solution at all. Decomposition and localization guide the search for an adequate model within that space; that is, they guide and constrain the construction of explanatory models. Another, distinct, limitation is critical here: the data vastly underspecifies the nature of the mechanism. Thus, the researcher must develop strategies for finding plausible models given the data. Decomposition and localization simultaneously provide such a strategy. They not only lead researchers to begin with the most manageable models, but also with ones that have high probative value. By assuming decomposition and seeking a direct localization, scientists can be led to models that are easy to manage and which often guide the search for additional factors that must be explained by a more complete model. Hypothesizing that nature is nearly decomposable, and developing linear models, reduce the cognitive demands and can lead to models with areas in which organization is more complex and in which component integrity is undermined. Thus, even failures in decomposition and localization may be guides to fruitful models of integrated component systems.

We emphasized that the need for heuristics stems from the limitations of human information-processing. The actual heuristics employed by a scientist, however, are not determined by these limitations alone. Intellectual and cultural backgrounds also shape the heuristics actually utilized. In particular, the restriction to mechanistic models is culturally conditioned. This affects both the experimental methods that are applied and the theoretical applications that accompany them. What counts as a mechanistic model, moreover, can and does change. Cartesian mechanism excluded action at a distance, and Newton certainly never contemplated the array of mechanical forces that are now commonplace.

In more traditional, empiricist-inspired accounts of scientific explanation, what scientists are thought to explain is a body of data. In the simplest construal of the deductive-nomological model, the data is taken to be a class of observations. The inference from data to theory may be inductive or abductive; either existing or novel observations serve as data against which theories are measured (cf. Bechtel 1988). It is now commonplace to note that "data" is theory laden. Since Hanson (1958) and Kuhn (1962) drew our attention to the phenomenon, philosophers and historians have accepted that data assumes a significance only in light of the theoretical assumptions brought to it. At the very least the theoretical background provides the conceptual apparatus in terms of which the data is understood and described. Yet, it is still often assumed that the goal is to explain the data; that is, to derive the observations from some law or set of laws.

Just as we do not deny the existence and usefulness of laws, we do not deny the existence and usefulness of data. Data in a relatively unproblematic sense is often useful in guiding the search for models and is regularly employed in evaluating proposed explanations. However, data are not the *explananda* of scientific models. We observe data, but we explain phenomena. Phenomena are not data in any useful sense. Phenomena are *repeatable* features of the world which reveal themselves in a variety of experimental arrangements. If one arrangement leads us to think that there is a body of phenomena that cannot be revealed within other paradigms, we suspect that we are confronted with an artifact. We rely on data to identify phenomena to be explained and again to evaluate models. As James Bogen and James Woodward explain in an insightful essay, explanatory models do not predict or explain the data we observe, although that data provides relevant information about the phenomena:

> Data, which play the role of evidence for the existence of phenomena, for the most part can be straightforwardly observed. However, data typically cannot be predicted or systematically explained by theory. By contrast, well-developed scientific theories do predict and explain facts about phenomena. Phenomena are detected through the use of data. (1988, pp. 305–6)

Except in special cases, such as when we try to explain why an experiment did not give us the results expected, we do not, and cannot, explain data. We may have what we think is a reliable means of detecting some phenomenon, and in the process we may rely on data, but we may not be able to explain how we got the data or why it is a reliable indicator of the phenomenon. To explain the data would require knowing, for example, how the instruments are used in the procurement of data work. But we stable, repeatable characteristics which will be detectable by means of a variety of different procedures" (Bogen and Woodward 1988, p. 317). Data are ephemeral. Phenomena are robust.

Cartwright's following defense of "phenomenological laws" is a defense of phenomena in the sense we intend:

> Phenomenological laws describe what happens. . . . For the physicist, unlike the philosopher, the distinction between theoretical and phenomenological has nothing to do with what is observable and what is unobservable. Instead the terms separate laws which are fundamental and explanatory from those that merely describe. (1983, p. 2)

Pre-Copernican astronomers spoke of "saving the phenomena," meaning both the more regular motions of heavenly bodies and the less regular occurrences such as eclipses, occlusions, and retrogression. *Phenomenological* thermodynamics includes within its domain the regularities that govern the behavior of gases, such as the Boyle-Charles Law and much more. Though some researchers would offer an empiricist rendering of these laws (Duhem 1906; van Fraassen 1980), reducing them to the data points from which they were derived, data in the sense of the products of individual observations and experiments are not what scientists set out to explain. Hacking illustrates the difference between a datum and a phenomenon:

> A phenomenon is *noteworthy*. A phenomenon is *discernible*. A phenomenon is commonly an event or process of a certain type that occurs regularly under definite circumstances. The word can also denote a unique event that we single out as particularly important. When we know the regularity exhibited in a phenomenon we express it in a law-like generalization. (1983, p. 221)

There is feedback between phenomena and explanatory models. Because something must be noteworthy or detectable as a regularity to be a phenomenon, what counts as a phenomenon to be explained depends on what explanatory models are available and alive. In acknowledging this, however, we do not intend to endorse the constructivism found in Hacking's comment that "experiment is the creation of phenomena" (ibid., p. 229). Ours is an ecumenical realism. Minimally, a phenomenon is a regularity in a system's behavior. The fact that yeast cells ferment, and that they do so at a particular rate under specified conditions, is a phenomenon that calls for explanation. So is the fact that a trait is heritable within a particular environment. The fact that tissues can respire in the absence of blood is another phenomenon that must be explained by a theory of respibe explained by a model of spatial-problem solving. All of these are phenomena to be explained. They are real, and so, we think, are their causes.

Phenomena are an important constraint on the attempt to develop an explanation. An explanation is erroneous if it fails to account for them. They speak with a voice, though, that is subject to interpretation. In recent debates over the architecture of cognition, some theorists claim that any explanation of the phenomena surrounding cognition must invoke symbolic structures and rules. Fodor and Pylyshyn (1988), for example, contend that the productivity and systematicity exhibited by cognitive phenomena *require* a system with a certain architecture. While the characteristics of the phenomenon to be explained may be suggestive as to possible structures of the explanatory model, they are not sufficient to necessitate a particular form to the model. As we observed earlier, Justus Liebig attempted to formulate the structure of a general model of animal chemistry on the basis of the phenomenon of animal nutrition: animals consume foodstuff containing more complex chemical structures than are found in their waste product. Liebig concluded that all the reactions in the animal body accordingly had to be catabolic. Subsequently, Bernard demonstrated the occurrence of glycogenesis and argued for the more general phenomenon of synthetic reactions in the animal. This meant not only that the explanatory model needed revision, but that metabolism needed to be understood in different terms. Metabolism was not simple decomposition of foodstuff. Likewise, classical genetics construed its phenomena in terms of the transmission of phenotypic traits. One consequence of Beadle and Tatum's work was the recognition that the phenotype, and the phenomena of genetics, must themselves be understood in biochemical terms. Phenomena, as well as the explanatory models, change. This is what we have called reconstituting the phenomena (see Chapter 8).

Operational Constraints

We remarked in Chapter 2 that in complex systems, and particularly in self-organizing systems, the nature of the components and their contributions cannot generally be inferred from normal behavior of the system alone. Interaction between components in a smoothly functioning system makes any simple inference from the phenomena to the mechanism problematic. What is required are independent means for assessing and understanding the behavior of components. Often a great deal of resourcefulness is needed in the development and utilization of experimental procedures to study a system and its components.

normal system. We would, for example, gain little information about the function of the liver by simply cutting it out of the body, unless we have developed means for interacting with it that help to reveal its natural operation. What we require are techniques that show what the parts of a system are and how they function within the system.

Often what are required are appropriate instruments and experimental techniques. To measure the respiration of organisms, Lavoisier and LaPlace had to develop a calorimeter. To measure the respiration promoted by Atmungsferment, Warburg had to develop a special flask. To determine the role of the hippocampus in spatial memory, O'Keefe and Nadel had to carefully lesion the hippocampi in experimental rats. Not only must instruments and techniques be developed, but they must be used in a manner that is truly probative. It was no small flaw in Flourens's technique that he lacked aseptic methods, and it is not small virtue of Pasteur that he did not. At other times what is required is a useful model system. In some cases a model system might be a natural system that is particularly useful because it is easy to work with or thought to provide a good example of the phenomenon. The use of insect models in determining the locus of respiration, and of *Drosophila* as a useful genetic model, both played critical roles in the development of mechanistic models. And sometimes the model systems are artificial systems that are thought to simulate real systems and provide demonstrations of how they might work.

The use of instruments, techniques, and model systems are all prone to the production of artifacts. For experimental scientists this is a continuous worry. Part of the reason this concern looms large is that the use of instruments, techniques, or model systems cannot generally be justified by appeal to well-developed and tested theories about how they work. They are generally widely employed long before researchers understand why they work as they do. Rather, they are often justified by more indirect means. For example, the data they give may seem so clean that researchers are convinced that it must reflect the phenomena under study. In other cases the picture they provide of phenomena is corroborated by evidence produced by other techniques. Finally, sometimes the fact that they provide evidence that makes sense in terms of emerging models is seen as vindicating their use (Bechtel 1990).

Although we have not focused in detail on the introduction of instruments and research techniques, or on the factors used to evaluate potential model systems, we have discussed two general strategies. One termed this an inhibitory method. In the simplest form, this approach involves the observation of deficits in the presence of damage to the system, as when we observe linguistic deficits with trauma to the left frontal lobe. In most instances, however, the case is more indirect. O'Keefe and Nadel inferred the function of the hippocampus, in part, from the behavior of animals suffering hippocampal lesions, assuming that they would compensate for the damage. They sought what Jackson called *positive* symptoms rather than simple deficits. The requirement that we be able to isolate intermediaries noted in Chapters 7 and 8, is a relatively strong operational constraint. Abnormal excretions, as in alkaptoneuria, provide strong evidence about metabolic intermediaries in intact systems. Meyerhof's use of flouride to inhibit muscle glycolysis, with a resulting buildup of hexosediphosphate, was evidence of the latter's importance as a metabolic intermediary.

The second experimental approach we discussed follows the opposite course, providing an unusually strong stimulus in order to elicit evidence of what a component of a system does. We have termed this an *excitatory* method. Artificial stimulation of the cortex, as in the classic work of Fritsch and Hitzig (1870), provides relatively direct information about centers of control in the nervous system. The injection of metabolic intermediaries is similar. An important part of the positive case for the role of phosphates in glycolysis depends on Embden's demonstration that an artificial increase of hexosediphosphate leads to an increase in lactic acid.

As we have noted, the issue of how the results should be interpreted arises in both inhibitory and excitatory studies. The fact that a deficit occurs from an inhibitory study does not prove what the inhibited part contributes to the normally operating system. The deficit may in fact be the result of a variety of interactions in the system that occur in an atypical fashion due to the interruption. This is particularly true in integrated systems. Similarly, the increased performance, or the failure to generate an increased output, from an excitation does not categorically point to the (non)contribution of the excited component. For example, the addition of intermediates to the fermentation system did not necessarily result in an increase in alcohol production, because this activity required the integration of different components, not just the one being excited.

We have focused on only two strategies for the development of mechanistic models. While we do not intend this to be a complete list of research tools used in such investigations, we do want to draw attention to the fact that what can be experimentally accessed is as important as what might be theoretically recognized. Science is, in large measure, a practical art.

dependencies":

> It is not a satisfactory explanation of an outcome merely to assert that it is due to some general mechanism, where the details of the mechanism are left unspecified. Instead, a satisfactory systematic explanation must show how the features of the explanandum-phenomenon systematically depend upon the factors invoked in the explanans of that explanation. (1988, p. 323)

When we are dealing with models developed through the heuristics of decomposition and localization, we embrace strong assumptions about how the system is physically realized. It is an important aspect of the cases discussed in Part III that evidence for the physical realizability of the component functions was required. This entails our having evidence from the lower level for the mechanisms posited in the decomposition: Wernicke defended his explanation of the aphasias in part on the basis of Meynart's work in neuroanatomy. Researchers investigating fermentation required that each step in a proposed model involve an independently known chemical reaction. As we noted, it was one of the primary virtues of Neuberg and Kerb's model that it met this constraint. Garrod's work, as we have also mentioned, faced the same constraint.

The insistence on physical realizability is suggestive of reduction. In any useful sense of the term, we do accept that decomposition and localization constitute a reductionistic strategy. Mechanistic explanations *do* explain the activities of a system in terms of the behavior of component parts. Decomposition and localization are not reductionistic, however, in the sense that is current in the philosophical literature. Traditional literature has focused on laws stated as universally quantified statements and on whether the laws of one theory can be derived from those of another. In addition to the derivation of one set of laws from another, philosophical models of reduction require that *bridge laws* connect the terms of one theory with those of the other. There has, for example, been a great deal of controversy within philosophy over whether Mendelian genetics has actually been reduced to molecular genetics, or whether psychology might thus reduce to neuroscience. Much of this controversy focuses on whether we can establish appropriate bridge laws linking molecular and Mendelian concepts (cf. Hull 1974; Kitcher 1984; Wimsatt 1976). It appears that Mendelian concepts fail to map smoothly onto molecular mechanisms, so that the same Mendelian mechanism is subserved by different molecular mechanisms, and the same molecular mechanism may serve different Mendelian functions in different contexts. Philip Kitcher declares flatly,

> body of statements one is attempting to axiomatize. The history of genetics makes it clear that Morgan, Muller, Sturtevant, Beadle, McClintock, and others have made important contributions to genetic theory. But the statements occurring in the writings of these workers seem to be far too specific to serve as parts of a general theory. (1984, p. 351)

We have no doubt that the cases we have discussed, including the work in biochemical genetics, will fail to fit the philosophical mold that has been cast for reduction. The investigators we have examined were not concerned with developing general laws, and when they did turn to the lower level it was not in order to derive antecedently developed laws at the higher level. Rather, they turned to the lower level to create models of mechanisms that would explain specific processes observed at a higher level. Thus, such scientists are subject to the very objection Kitcher advances. Because we regard the philosophical model as not much more than a philosophical construct, having little to do with scientific practice, this leaves us unconcerned. Whether it counts as reduction or not, we think that the history of genetics makes it transparently clear that the work of Morgan, Muller, Sturtevant, Beadle, McClintock, and others *was* part of a general theory of development and heredity, and that this general theory *was* intended to explain development and heredity in biochemical terms.

3. Conclusion: Looking Forward

Our efforts have been directed at reaching a realistic account of the development of science. If we have been successful, then we have correctly characterized the kinematics of one sort of research program—one directed toward developing mechanistic explanations in the face of complex natural systems. We have, more speculatively, indicated four factors, or types of factors, that feature in the dynamics of theories. More is still to be done in developing a dynamic model and in elaborating the component forces.

The focus on decomposition and localization takes us in the direction of models that span more than one level. As we have seen, decomposition and localization are fallible, as are all heuristics; but even when they fail, they may serve as probative tools for facilitating discovery. The emphasis on the lower-level mechanisms in explaining processes identified and conceptualized at higher levels is what many scientists think of as reductionistic. In any case, lower-level constraints are essential in developing and elaborating explanatory models. Absent such constraints, either for lack of appropriate research tools or because of a failure to pursue experimental elaborating and developing interlevel mechanistic explanations. We have offered a qualitative model of theory development designed to account for the processes by which scientists develop mechanistic models of phenomena of interest. This account is meant to be psychologically and historically realistic. The choices we have identified are intended as points in the development of research programs at which decisions shape the course of research and the formulation of models.

The result may not suit a general philosophical temperament. We offer no universal procedures, no single method for science. We see no invariant pattern. Scientific disciplines evolve under differing constraints with differing histories. Their actual development, as historians have long seen, cannot be understood apart from the changing historical context. This does not mean the process is ungoverned or arbitrary. As Peter Galison (1990) suggests, the development of scientific theories is not so much plastic as it is "immensely constrained." If we are right, the problem is one of reaching a solution that simultaneously satisfies a complex array of changing constraints. An account of theory development that takes such constraints seriously requires a measured historicism. Such is also a realistic historicism.