Role Functions, Mechanisms, and Hierarchy
Author(s): Carl F. Craver
Source: *Philosophy of Science*, Vol. 68, No. 1 (Mar., 2001), pp. 53-74
Published by: The University of Chicago Press on behalf of the Philosophy of Science Association
Stable URL: http://www.jstor.org/stable/3081024
Accessed: 07/10/2011 12:46

Role Functions, Mechanisms, and Hierarchy*

Carl F. Craver†‡
Department of Philosophy
Florida International University

Many areas of science develop by discovering mechanisms and role functions. Cummins' (1975) analysis of role functions—according to which an item's role function is a capacity of that item that appears in an analytic explanation of the capacity of some containing system—captures one important sense of "function" in the biological sciences and elsewhere. Here I synthesize Cummins' account with recent work on mechanisms and causal/mechanical explanation. The synthesis produces an analysis of specifically mechanistic role functions, one that uses the characteristic active, spatial, temporal, and hierarchical organization of mechanisms to add precision and content to Cummins' original suggestion. This synthesis also shows why the discovery of role functions is a scientific achievement. Discovering a role function (i) contributes to the interlevel integration of multilevel mechanisms, and (ii) provides a unique, contextual variety of causal/mechanical explanation.

*Received June 2000; revised October 2000.

†Send requests for reprints to the author, Department of Philosophy, Florida International University, 3000 Northeast 151 Street, North Miami, FL 33181–3000; craverc@fiu.edu.

‡Thanks to Ron Amundson, Robert Cummins, Lindley Darden, Michael Devitt, Stuart Glennan, Peter Machamer, Karen Neander, Greg Morgan, Pierre Poirier, Gualtiero Piccinini, Wesley Salmon, Ken Schaffner, Steven Small, Wendy Stuart, Nathan Urban, Marcel Weber, William Wimsatt, and Kirsten Wood for comments on earlier drafts of this much reworked paper. Thanks also to the students in Darden's graduate seminar (Fall 1998) in the Philosophy of Biology for useful feedback on multiple drafts. Many of these friends and colleagues will not recognize this as the paper that they so generously read. Any mistakes are mine and mine alone. This work was supported by the National Science Foundation under grant number SBR-9817942. Any opinions, findings and conclusions or recommendations expressed in this material are those of the author and do not necessarily reflect those of the National Science Foundation.

Philosophy of Science, 68 (March 2001) pp. 53–74. 0031-8248/2001/6801-0004$2.00

1. Introduction. Many areas of science (and especially the biological sciences) develop by discovering mechanisms and the role functions of their components. Promises of understanding the mechanisms of, for example, development, disease, and cognition are coupled with claims to have found the roles of various genes, cell types, and brain regions. Yet philosophers have said surprisingly little about how one discovers an item's role, about why such a discovery is a major scientific achievement, or about how these mechanistic and functional descriptions are related.

Most of what has been said about role functions in the philosophy of science has been said by Cummins (1975, 1983), whose position is often repeated and rarely revised. Cummins' account highlights the conceptual interdependence of role functions and what he calls an "analytic explanatory strategy" of understanding the capacities of systems by analyzing them into the capacities of their components. Here I explore the possibilities and consequences of synthesizing Cummins' account of role functions with recent work on the nature of mechanisms and causal/mechanical explanation. The synthesis yields a more detailed analysis of properly *mechanistic* role functions and of the empirical criteria by which mechanistic role ascriptions are evaluated. This synthesis also shows why the discovery of a mechanistic role is a major scientific achievement. It is an achievement, first, because discovering an item's mechanistic role is one way of integrating it into a multilevel mechanism and, second, because integration into a higher-level mechanism constitutes a unique, *contextual* variety of causal/mechanical explanation.

The argumentative structure of the paper is as follows. In Section 2, I review Cummins' (1975) "analytic account" of role functions, emphasizing as he did the connection between functional ascription and analytic explanation. In Section 3, I discuss the character of mechanisms and, most importantly, their active, spatial, and temporal organization. I then use these aspects of mechanistic organization to specify—more precisely than Cummins could—what one asserts of an item in ascribing it a *mechanistic* role function. This understanding of the content of mechanistic role ascriptions highlights the diverse kinds of evidence by which mechanistic role ascriptions are evaluated. In Section 4, I distinguish contextual, isolated, and constitutive descriptions of an item's activity in a multilevel mechanism. It is by elaborating and aligning these different kinds of description, I suggest, that the levels in multilevel hierarchies are integrated together. In Section 5, I introduce the possibility of *contextual mechanistic explanation* as a unique variety of causal/mechanical explanation, and I respond to some criticisms of this suggestion.

2. Cummins on Analytic Explanation and Role Functions. Cummins' (1975) analysis of function-ascribing statements is now the canonical account of role functions. Cummins' "regimented reconstruction" of function-ascribing statements is as follows:

> X functions as a ϕ in S (or the function of X in S is to ϕ) relative to an analytic account A of S's capacity to ψ just in case X is capable of ϕ-ing in S and A appropriately and adequately accounts for S's capacity to ψ by, in part, appealing to the capacity of X to ϕ in S. (1975, 190)

Here, S is a system with the capacity to ψ, and X is a component of S that has the capacity to ϕ. Cummins does not explicate the relationship between S and X, but we can assume that Xs are intended as at least mereological parts of Ss. (Rescher 1955) The relationship between S's ψ-ing and the ϕ-ing of Xs (diagrammed in Figure 1(a)) has traditionally been illustrated with the example of the heart and the circulatory system (diagrammed in Figure 1(b)).

So following tradition[1]: The heart (X) functions as a blood pump (ϕ) in the circulatory system (S) relative to an analytic account (A) of the circulatory system's (S's) capacity to deliver oxygen and calories to body tissues (ψ) just in case the heart (X) is capable of pumping blood (ϕ-ing) in the circulatory system (S), and the analytic account (A) appropriately and adequately accounts for the ability of the circulatory system (S) to deliver oxygen and calories to body tissues (ψ) in part by appeal to the capacity of the heart (X) to pump blood (ϕ) within the circulatory system (S). The clarity of Cummins' regimented reconstruction turns on the clarity of his understanding of systems, of the system/component relationship, and of analytic explanations.

Analytic accounts (A) are explanations. They explain by analyzing the capacities of systems into the capacities of their component parts. System S's capacity to ψ is explained by analyzing S into the parts $\{X_1, X_2, \ldots, X_m\}$ and capacities $\{\phi_1, \phi_2, \ldots, \phi_n\}$ relevant to S's capacity to ψ. Analytic explanations explain by showing the ψ-ing of S to be "reduced to the programmed exercise of the analyzing subcapacities." By "programmed" Cummins means, "organized in a way that could be specified in a program (or flow diagram)." (1983, 100) So the circulatory system's (S's) capacity to deliver goods to body tissues (ψ) is explained by decomposing it into its parts (e.g., hearts (X_1), arteries (X_2), kidneys (X_3), and valves (X_4)) and capacities (e.g., to pump (ϕ_1), to convey (ϕ_2), to filter (ϕ_3), and

1. The central findings of the present essay were originally developed in the context of an example from contemporary neuroscience (specifically, the LTP-Learning hypothesis). The example of the circulatory system is far more familiar and less confusing. This simplified analysis of a familiar case can nonetheless be applied to more complicated cases, like the role functions of genes, pathogens, and brain regions.

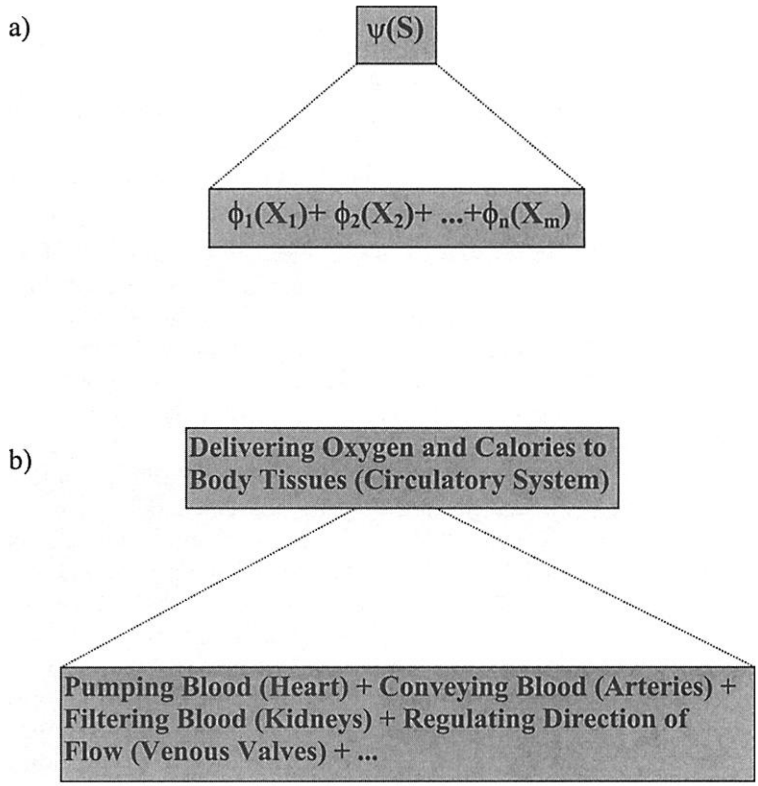

Figure 1. Relation between system S's ψ-ing and the ϕ-ing of Xs, represented abstractly in (a) and applied in the case of the circulatory system in (b).

to regulate the direction of blood flow (ϕ_4)) and linking those parts together in the programmed ψ-ing of the circulatory system.

Cummins emphasizes that not all analytic explanations are interesting. An analytic account of S's ψ-ing is explanatorily interesting only to the extent that:

(C1) The analyzing capacities $\{\phi_1, \phi_2, \ldots, \phi_n\}$ are "less sophisticated" than the analyzed capacity (ψ);

(C2) The analyzing capacities $\{\phi_1, \phi_2, \ldots, \phi_n\}$ are "different in type" from the analyzed capacity (ψ); and

(C3) The analyzing components $\{X_1, X_2, \ldots, X_m\}$ and $\{\phi_1, \phi_2, \ldots, \phi_n\}$ exhibit a complex organization such that together they ψ (see 1975, 191).

The analytic explanation of the circulatory system's working arguably satisfies these criteria, although Cummins does not show that this is so or say what it means for two items to be "different in type" or "less sophisticated." The parts are (perhaps simply by virtue of being parts) less sophisticated and different in type than the whole system; and the activities of hearts, arteries, kidneys, and valves have to be organized properly if each is to play its role. So an analytic account of the circulatory system would arguably be explanatorily interesting. But what does it mean to be "different in type," "less sophisticated," and "organized" in this context?

Here one begins to sense that Cummins hasn't said enough about systems, about analytic accounts, or about how components in analytic accounts are "organized" to distinguish interesting from uninteresting analytic explanations or to distinguish cases that warrant the ascription of role functions from those that do not. Criteria C1–C3 do point to important *symptoms* of interesting analytic explanations, but they do not in any way characterize what makes those explanations interesting. And without such detail about the character of causal/mechanical explanations, one is left with an underspecified notion of a role function. By filling in these abstract placeholders in Cummins' account one can produce a richer and more precise image of properly mechanistic role functions. The characteristic active, spatial, temporal, and hierarchical organization of mechanisms can be used to specify—more precisely than Cummins has—what is asserted by an analytic account, and hence by the ascription of a mechanistic role.

3. Mechanisms and Their Organization. There are many different kinds of systems (e.g., formal systems, procedural systems, and representational systems), and Cummins does not specify which he intends. In this section, I suggest that one restrict the idea of a "system" in Cummins' analysis to mechanisms (described in Section 3.1). I then use the characteristic active, spatial, and temporal organization of mechanisms (described in Section 3.2) to add necessary precision to the idea of a mechanistic role function and, further, to make sense of the empirical criteria by which mechanistic role ascriptions are evaluated (discussed in Section 3.3). There may be other interesting kinds of organization than mechanistic organization, but the content of a functional attribution will only be as precise and meaningfully specified as is the accompanying notion of "organization."

3.1 Mechanisms. Mechanisms are collections of entities and activities organized in the production of regular changes from start or set up conditions to finish or termination conditions (Machamer, Darden, Craver 2000; cf. Glennan 1996, 52; cf. Bechtel and Richardson 1993, 17). The entities in mechanisms can be taken to correspond to Cummins' $\{X_1, X_2, \ldots, X_m\}$; they are the physical parts of the mechanism (e.g., the hearts, kidneys, and veins). Activities are the things that these entities do, either by themselves or in concert with other entities. The activities in mechanisms can be represented as the $\{\phi_1, \phi_2, \ldots, \phi_n\}$ in Cummins' account.[2]

3.2 Mechanistic Organization. Descriptions of mechanisms characterize how entities (Xs) and activities (ϕs) are *organized* to do something (ψ). Cummins' account of analytic explanations appeals, without explication, to the "organized" or "programmatic" exercise of capacities. And criterion C3 explicitly uses "organization" to distinguish interesting from uninteresting analytic explanations. What is the characteristic organization of mechanisms?

It is perhaps too often claimed that the whole is greater than the sum of its parts. However, Wimsatt (1986, 1997) has revitalized this cliché by focusing on conditions under which the whole is *nothing but* the sum of its parts—conditions of *mere aggregativity.* (For a related treatment, see Haugeland's 1998 discussion of morphological and systematic explanation in Ch. 1 and of decomposition in Ch. 9).[3] Suppose that a property ψ of the whole S is a function of the properties $\{\phi_1, \phi_2, \ldots, \phi_n\}$ of the parts $\{X_1, X_2, \ldots, X_m\}$.[4] Then a ψ property of S is an aggregate of the ϕ properties of Xs when:

2. Calling them "activities" rather than "capacities" (as is Cummins' preference) emphasizes the way that mechanisms work as opposed to the way that a set of parts has the capacity to work or is disposed to work. This shift of gestalt is recommended and defended by Machamer, Darden, and Craver (2000). (See Footnote 4 below for a discussion of my shifting ontology for ψs and ϕs).

3. Thanks to Cummins and Poirier (personal communication) for reminding me of the discussion in Haugeland 1998, Ch. 1.

4. I use the terms "capacity," "activity," "property," and "event" to characterize ϕ and ψ in different parts of this essay. This is mostly an effort to stress points of contact across these terminological differences. The following remarks will perhaps assuage worries that I am ignoring important differences. I use "capacity" when explicating Cummins. I prefer talk of activities, and I take the notion of a "capacity" to be a substantivalist way of expressing the fact that there are certain properties of entities which allow them to engage in activities (see Footnote 2). I use the term "event" to describe ϕ and ψ in Section 5 in an effort to highlight connections with Salmon's (1984) image of causal/mechanical explanation; all activities are events, but not all events are activities. In the present section, I use the term "property," like "event," to stand for either the having of a property at a particular time in a particular place or the occurrence

(W1) ψ is invariant under the *rearrangement and intersubstitution* of Xs;
(W2) ψ remains *qualitatively similar* (if quantitative, differing only in value) with the addition or subtraction of Xs;
(W3) ψ remains invariant under the *disaggregation and reaggregation* of Xs; and
(W4) There are no *cooperative or inhibitory* interactions among Xs that are relevant to ψ.

Wimsatt's criteria (W1–W4) gesture at the sense of "organization" required for explicating the conceptual interdependence of mechanisms and mechanistic role functions, for getting beyond the symptomatic criteria expressed in Cummins' C1–C3, and so for developing a mechanistic account of role functions.

Compare the circulatory system (S_1) to a neat glass of gin (S_2) and, likewise, hearts, kidneys, and arteries (the Xs in S_1) to unit volumes of gin (the Xs in S_2). The circulatory system (S_1) delivers goods to tissues (ψ_1), and the glass of gin (S_2) has a certain volume (ψ_2). The parts (Xs) of the circulatory system, such as the heart and kidney, cannot be intersubstituted for one another (W1). Kidneys do not pump blood and hearts do not filter it. Changing even the spatial relations among the components of the circulatory system would (at least in many cases) completely disrupt the behavior of the whole system (W2). In fact, only judicious removal of parts from (or addition of parts to) the circulatory system is compatible with its continued working as a whole (W3). Finally, there are excitatory and inhibitory interactions between the components of the circulatory system (W4), and this is exactly why you cannot tinker with the entities and activities of the circulatory system too capriciously without breaking it. The volume of gin, in contrast, stays the same volume of gin any way you stir it. The total volume only increases and decreases as you pour and sip, and unit volumes do not work together in any interesting sense to produce a total volume (even if they sometimes seem to).

Each of Wimsatt's criteria (W1–W4) points to the absence of cooperative activity among the parts in mere aggregates. The components of mechanisms, in contrast to those of mere aggregates, have an *active organization*; they act and interact with one another in such a way that the ψ-ing of S is more than just the sum of ϕ properties. In fact, the ϕ properties of mechanisms are not really mere properties at all; they are the activities of and among the entities in the mechanism. Typically, mechanisms are composed of different kinds of entities (like hearts, kidneys, and

of some sort of temporally extended happening or change. The ψ properties of mere aggregates are often, but not always, events in the first sense, while the ψ properties of organized mechanisms are typically, but not exclusively, events in the second sense.

valves) engaging in different kinds of activities (like pumping, filtering, and directing) and acting in cooperation or competition with specific other entities in the mechanism (W4). It matters which Xs ϕ with which others, and it matters where, when, how much, and how often. This is why the parts of mechanisms often cannot be reorganized randomly (W1), added or subtracted at will (W2), or taken apart and put back together again (W3) without disturbing their ability to ψ. As Wimsatt notes, it is often by doing these unhealthy things to a mechanism that scientists are able to tease apart the relevant features of its organization. Not coincidentally, these are also the kinds of experiments done to discover and specify an item's mechanistic role.

The active organization of mechanisms is sustained by their characteristic spatial and temporal organization. The same entities and activities, strung together in different spatial and temporal relations to one another, can yield very different mechanisms. One understands a mechanism by discovering its component entities and activities, and by learning how their activities are spatially and temporally organized (Craver and Darden, forthcoming); aspects of this organization are also among the evidential criteria for evaluating ascriptions of mechanistic role functions.

Starting with the *spatial organization* of mechanisms, the entities in mechanisms often have crucial sizes, shapes, orientations, and locations that allow them to engage in certain activities (and hence certain roles) and not in others. The heart's size is appropriately related to that of the vena cava and that of the aorta such that it delivers blood in regular volumes to distal regions of the body within a specific range of rate and pressure. It has a shape that allows it to act as a bellows, its auricles and ventricles alternately receiving and expelling blood. The heart is situated between the vena cava and the pulmonary artery (on one side) and between the pulmonary vein and the aorta (on the other). Part of understanding the organization of this mechanism, as Harvey ([1628] 1963) exhibits in the diverse observations and experiments of his *Movement of the Heart and Blood in Animals: An Anatomical Essay*, is understanding the spatial arrangement of its component parts. The other part of understanding a mechanism's organization, also featured in Harvey's treatment, is understanding how the activities of these component entities are temporally organized.

The activities of mechanisms exhibit a *temporal organization* in the activity ψ of the mechanism as a whole S. The order, rate, and duration of successive component activities (ϕs) is crucial for S's ψ-ing. Blood entering the heart through the vena cava is then pumped through the pulmonary artery to the capillaries of the lungs, where oxygen and carbon dioxide are exchanged. Blood returns to the heart via the pulmonary vein before being shipped off to the rest of the body. There is a sequence of stages here from beginning to end, and it would not be possible to change their order without gumming up the works (or making it a different mechanism entirely). The different stages also have characteristic rates and durations that are crucial to the working of the mechanism (as the diagnostic value of heart rate and blood pressure attests). Timing is everything for most mechanisms, and learning a mechanism's timing provides important clues to how it works and to how it does not work.

Understanding how a mechanism works is just understanding how one activity leads to the next through the spatial layout of the components and through their participation in a stereotyped temporal pattern of activities from beginning to end. This is what it means to say that a mechanism is "organized" or, as Cummins sometimes says, "programmed"; it means that the components have active, spatial, and temporal relations to one another such that together they ψ. An analytic account for a mechanism is not just a list of entities and activities; it is a description of a mechanism. And that description involves, in addition to a list of entities and activities, a description of how they are organized together actively, spatially, and temporally in S's ψ-ing. Specifying the mechanistic role of some component X, accordingly, involves describing how X is organized with the other entities in S such that it contributes to S's ψ-ing.

3.3 Implications: Mechanistic Role Functions. Attributions of mechanistic role functions describe an item in terms of the properties or activities by virtue of which it contributes to the working of a containing mechanism, and in terms of the mechanistic organization by which it makes that contribution.

This insight about the connection between role ascriptions and mechanistic organization is lost if one abstracts role functions away from the details of how functions are instantiated in mechanisms. Cummins skirts this difficulty when he allows that analytic accounts may be specified by a "flow diagram," and so specified "independent of whatever theory is relevant to describing the details of the realization." (1983, 100) In one version, Cummins (1983) goes so far as to drop explicit reference to the parts $\{X_1, X_2, \ldots, X_m\}$ in analytic accounts, and speaks only of the analysis of one capacity into the programmed exercise of sub capacities $\{\phi_1, \phi_2, \ldots, \phi_n\}$.

But ascriptions of mechanistic role functions are detailed and precise to the extent that they can be explicated in terms of specific details of how an item fits into the active, spatial, and temporal organization of a mechanism that we seek to understand. Flow diagrams (or "box-and-arrow" diagrams) are often useful in this service but they are also often so sketchy that it is impossible to say (at least on the basis of the terse descriptions suspended over the arrows or crammed into the boxes) just how these roles might be filled in that mechanism. Imprecise and abstract role descriptions are often used in discovery as rough-draft stand-ins for more detailed accounts of how an item fits into this mechanistic organization. But the meaningfulness and precision of an item's role ascription should be evaluated with reference to the precision with which one can detail the organization of the system containing the item. It is by detailing how an item fits into the spatial, temporal, and active organization of a mechanism (showing exactly how it contributes to S's ψ-ing) that one specifies its mechanistic role. There may be other kinds of role functions that are not mechanistic role functions, but the point remains that spelling out those alternatives will require an analysis of the intended sense of organization in a manner comparable to that being explored here.

The importance of mechanistic organization in the analysis of mechanistic role functions is reflected in the kinds of evidence by which mechanistic role functions are attributed and evaluated. The active, spatial, and temporal varieties of mechanistic organization are often used to determine whether or not an item can play a given role. The item has to be in the right place at the right time, it cannot be spatially isolated from other components, and it has to have the right size, shape, orientation (and other relevant properties) to interact with the other components of the mechanism. An activity that happens at the wrong time, that takes too long, or that unfolds too slowly for a given role cannot fill that role. Role attributions are also tested by adding and removing parts, by moving parts around, and by substituting them one for the other, as suggested by Wimsatt's criteria (W1–W4).

The analysis of mechanistic role functions produced by melding Cummins' account with work on the organization of mechanisms clarifies what is involved in characterizing an item's role in a mechanism and clarifies what would count as a complete description of an item's mechanistic role. This more precise analysis of mechanistic role functions also helps to show why the discovery of an item's role is a major scientific achievement. The first reason (discussed in Section 4) is that describing an item's role is one way to integrate that item into a multilevel hierarchy of mechanisms. The second (discussed in Section 5) is that role ascriptions provide a contextual variety of causal/mechanical explanation.

4. Interlevel Integration: Contextual, Isolated, and Constitutive Perspectives. Many of the theories in the contemporary biological sciences (and elsewhere) have the mechanistic organizational structure that I've been explicating, and most that do also exhibit a *hierarchical* mechanistic organization. That is, they describe nested networks of mechanisms within mechanisms, in which higher-level activities (ψs) of mechanisms (Ss) are instantiated by the organized activities (ϕs) of lower-level components

(Xs), and these are, in turn, instantiated by the activities (σs) of still lower level components (Ps). The relationship between lower and higher mechanistic levels is a mereological part/whole relationship with the additional restriction that the lower-level parts are components of (and hence organized within) the higher-level mechanism. Lower-level entities (e.g., the Xs) are proper parts of higher-level entities (S), and so the Xs are no larger than, typically smaller than, and always within the spatial boundaries of S. The activities of the lower-level parts are steps or stages in the higher-level activities. Exactly how many levels there are and how they are to be individuated are empirical questions that are often answered differently for different phenomena.

The circulatory system has a hierarchical mechanistic organization. The activities ψ of the circulatory system S are instantiated by the heart's pumping of blood, the kidney's filtration of the blood, and the venous valves' regulation of the direction of blood flow (the ϕ-ing of the component Xs). And these component activities can themselves be described in terms of their underlying mechanisms. The heart's pumping can be explained by reference to the contractions (σ) of component heart muscles (Ps), and the kidney's filtration can be explained by the organized activities of its component glomeruli, tubules, pores, and ionic gradients. This description could potentially continue on to the activities of entities appearing in still lower-level descriptions of mechanisms.

One goal in describing hierarchically organized mechanisms is to integrate these different levels together into a description of one coherent mechanism. This *interlevel integration* of mechanistic hierarchies involves elaborating and aligning contextual (+1 level), isolated (0 level), and constitutive (−1 level) descriptions of the ϕ-ing of some X. Each of these descriptions describes X and its ϕ-ing from a different perspective in a multilevel mechanism.[5]

A *contextual description* of some X's ϕ-ing characterizes its mechanistic role; it describes X (and its ϕ-ing) in terms of its contribution to a higher (+1) level mechanism. The description includes reference not just to X (and its ϕ-ing) but also to X's place in the organization of S's ψ-ing.

The amount of context included in contextual descriptions varies considerably from case to case. Consider four ways of describing the heart's mechanistic role in the circulatory system. The heart:

i. distributes oxygen and calories to the body;
ii. pumps blood through the circulatory system;

5. The connection between levels and functions was first suggested to me by Machamer 1977.

iii. expels blood; and
iv. contracts.

Descriptions (i–iii) are contextual in varying degrees; they each describe things that the heart could not do by itself without being organized together with other entities and/or activities. The heart cannot expel blood (iii) without blood, and the expulsion of blood will only circulate it (ii) if the veins and arteries are appropriately organized. Even then, the heart cannot distribute oxygen and calories (i) in the absence of oxygen and calories. A description of the X's mechanistic role function is contextual to the extent that it makes explicit reference to objects other than (and outside of) X itself.

Reference to objects beyond the boundaries of the heart, notice, is not required in describing the heart's contraction (iv). Contracting is something that the heart does by itself;[6] in describing the heart as contracting, one makes no commitments concerning the mechanistic context in which this activity is embedded. It is an *isolated description*. In forming an isolated description of some item's activity, one draws an idealized dividing line at the spatial boundary of the item and recognizes a limited number of crucial interfaces across that otherwise closed boundary. An interface may here be understood as a reasonably well-defined and regular locus of contact or interaction with objects outside of that boundary (cf. Haugeland 1998, Ch. 9 and also Wimsatt 1974; each ties his discussion to Simon 1969).

Typically there are many such interfaces between an item and its environment, and often only some of the interfaces are relevant to the item's contribution to a contextual mechanism. The heart makes glub-blub noises and generates heat, but these are not relevant to the circulation of the blood, even if they may be relevant to other containing mechanisms (such as diagnostic or thermoregulatory mechanisms). The relevant interfaces between the heart and the other components of the circulatory system lie at the spatial boundaries of the heart with the incoming veins, the outgoing arteries, and the blood coursing through them all. Described in isolation from its context at those interfaces, the contribution of the heart to the circulatory system is to contract; this is what it does (alone) that (in the right context) contributes to the expulsion of blood, the circulation of blood, and the distribution of oxygen and calories.

The distinction between contextual and isolated descriptions brings out an ambiguity in Cummins' unregimented rendition of his analysis of role functions. He claims that "to ascribe a function to something is to ascribe

6. I am neglecting the numerous interfaces between the heart and the rest of the body that are relevant to its contraction. As B-movies attest, however, the heart can continue to contract in isolation from that context.

a capacity to it that is singled out by its role in an analysis of some capacity of a containing system." (Cummins 1983, 99) But this leaves it ambiguous whether the function is the capacity, described in isolation and simply "picked out" by its contextual role, or, instead, the contextual role by virtue of which the capacity is picked out. A complete description of an item's role would describe each of these. One can know an item's contextual role without knowing the isolated activity by which the item plays that role in a given context; and one can know an item's isolated activity without knowing what that isolated activity contributes to a higher-level mechanism. There is a difference, after all, between knowing that spark plugs produce sparks and knowing how that sparking is situated within the complex mechanisms of an engine. In the former case, we know the spark plugs' isolated role; in the latter we describe their role contextually. (For a related discussion of the contextual aspects of capacity ascription, see Glennan 1997).

Isolated descriptions of an X's ϕ-ing specify the activity for which a lower-level mechanism will be sought and so fix the active, spatial, and temporal boundaries of that mechanism (cf. Kauffman 1971; Glennan 1996). Isolated descriptions, that is, frame *constitutive descriptions* of lower-level mechanisms. Constitutive descriptions characterize the organized activities $\{\sigma_1, \sigma_2, \ldots, \sigma_k\}$ of entities $\{P_1, P_2, \ldots, P_j\}$ that instantiate the ϕ-ing of X.

So it is possible to describe an item's activity in three distinct ways, depending on how one looks at it with respect to a hierarchy of mechanisms. Ignoring its context, one can describe x's ϕ-ing in isolation. Looking down to lower-level mechanisms, the activity is described constitutively. And looking up to higher-level mechanisms, the activity is described contextually.

Bechtel (1986) has sketched a similar picture of these interlevel relations. He criticizes Cummins for focusing "on only a single level in nature" instead of three: one (B1) for "components"; one (B2) for the "system functioning as an organized whole"; and one (B3) for "the environment in which the system is either *adaptive or nonadaptive*" (1986, 39; emphasis added).[7] Bechtel claims that the third level, B3, "allows one to integrate knowledge from several levels and thus provides a needed perspective in developing explanations in the sciences dealing with such systems" (1986, 42).

My picture of these relationships differs from Bechtel's (B1–B3) in that (i) I do not reify these three descriptions as "levels of nature" but as descriptive perspectives on a multilevel hierarchy, and (ii) I do not tie my analysis to the "adaptiveness" of higher-level activities.

7. The way that I read both Cummins and Bechtel, Cummins' account includes both B1 (X's ϕ-ing) and B2 (S's ψ-ing).

Bechtel's image blurs the distinction between descriptive perspectives and mechanistic levels. Contextual, isolated, and constitutive descriptions of an item's activity are three different *perspectives* on that item's activity in a hierarchically organized mechanism; they are not levels *of nature*. My perspectival approach is represented graphically in Figure 2 as a relation-

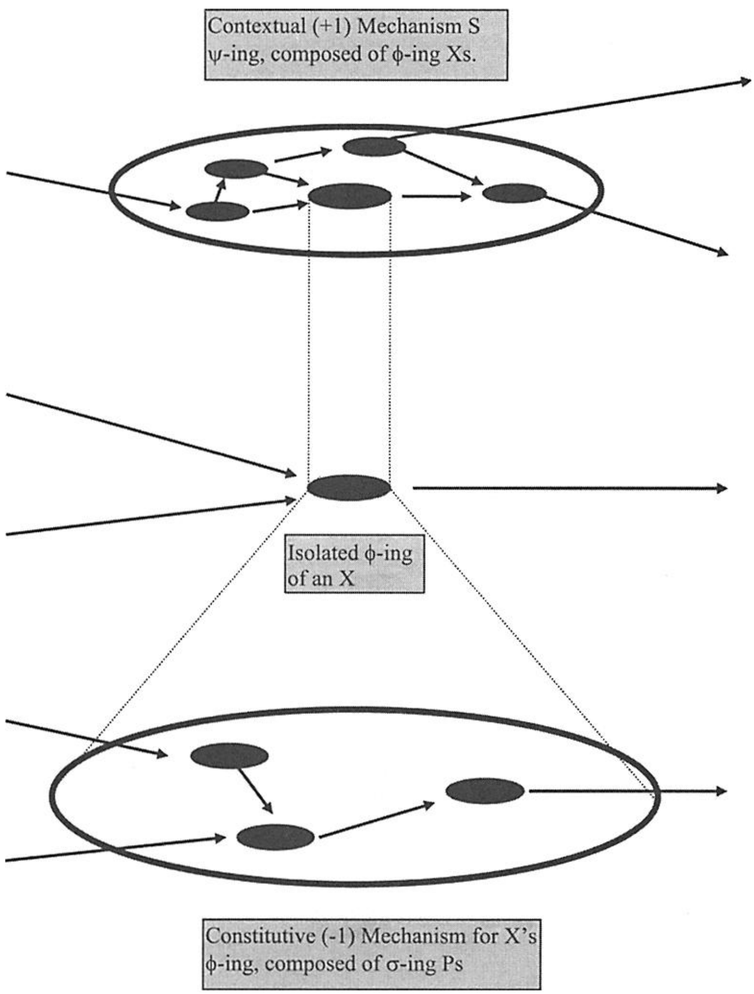

Figure 2. Contextual, isolated, and constitutive perspectives on an X's ϕ-ing. The contextual perspective reveals how X's ϕ-ing fits into a higher level mechanism S's ψ-ing. The isolated perspective describes X's ϕ-ing in terms of input-output relationships across interfaces between X and the other components of S. Finally, the constitutive perspective provides a description of the mechanism for X's ϕ-ing.

ship among three mechanistic levels: one for S's ψ-ing, one for X's isolated ϕ-ing, and one for the organized σ-ing of Ps. Describing X's ϕ-ing contextually involves describing how X's ϕ-ing fits into the organization by which S engages in ψ. This involves looking up (+1) a mechanistic level and detailing X's contextual role in a higher-level mechanism. X's ϕ-ing can also be described in isolation from its context in S—that is, in terms of the commerce across its interfaces with the other components in S. Finally, X's ϕ-ing can be described in terms of its constitutive mechanism. This involves looking down (−1) a mechanistic level and detailing the organized entities and activities that constitute X's ϕ-ing.

Contextual, isolated, and constitutive descriptions should not be taken as divisions in the furniture of the world (as suggested by Bechtel's term, "levels of nature"). Instead, they are distinct perspectives on an activity in a hierarchically organized mechanism. As Lycan puts it, "see Nature as hierarchically organized in this way and the 'function/structure' distinction goes relative: something is a role as opposed to an occupant, a functional state as opposed to a realizer, or vice versa, only modulo a designated level of nature." (1990, 78; cf. Churchland and Sejnowski 1992, 18–27) I would rather put it like this: see the world as a mechanistic hierarchy, and the distinction between a contextual role (+1), an isolated activity (0), and its constitutive mechanism (−1) goes relative to a perspective on an activity at a given level in a mechanistic hierarchy.

The second difference between Bechtel's account of interlevel integration and my own is that my account of mechanistic role functions does not appeal to any sense of adaptiveness in an environment; instead it appeals only to roles in contextual systems. These contextual systems may be adaptive or destructive, and they need not even be the kinds of systems for which talk of adaptation is appropriate. Heart disease, high blood pressure, cardiac arrhythmia, and arterial hardening all have mechanisms that span multiple levels, and this three-tiered perspective is as useful in those contexts as in those that are adaptive. Descriptions of hierarchical mechanisms are always descriptions of the mechanisms for some ψ, where ψ-ing is presumed to be something that one wants to understand (build, control, predict). That ψ in which the hierarchical mechanism "tops off" provides the necessary perspective from which to "integrate knowledge from different mechanistic levels," (Bechtel 1986, 42), without necessarily being adaptive or maladaptive.

These two revisions yield an interestingly different approach to interlevel integration in sciences that describe multilevel mechanisms. The levels to be integrated are perspectival mechanistic levels. The three descriptive perspectives constitute a descriptive goal in the integration of levels in a multilevel hierarchy of mechanisms. This goal might be put in the form of a directive. An activity (ϕ) is fully integrated into a multilevel mechanism when (i) the activity has been fit into the organization of a higher (+1) level mechanism, (ii) the isolated (0-level) activity has been adequately described, and (iii) the activity has been explained in terms of its lower (−1) level mechanism. One understands how the heart or kidney fits into the circulatory system to the extent that one knows how it is organized into the circulatory system, one has characterized its activity in isolation, and one knows the constitutive mechanism of that isolated activity. Integrating an item into a hierarchical mechanism involves elaborating and aligning these three kinds of description. In the process, one shows how the organized σ-ing of Ps could constitute the isolated behavior that is X's ϕ-ing, and shows how X's ϕ-ing is organized together with other components in S such that S can ψ. One might express a quite general *integrative strategy* in the discovery of mechanisms in a slogan: "Up for roles, down for mechanisms" (This strategy might be added to those discussed in Darden 1991, Chs. 2 and 15).

So one reason that it is a scientific achievement to describe the mechanistic role of an item in a mechanism is that doing so is one way of integrating that component into a multilevel mechanistic hierarchy.

5. Contextual Explanation: A Third Aspect of Causal/Mechanical Explanation. A second reason that discovering a role is a scientific achievement is that contextual role descriptions are *explanatory*. Consider Salmon's general sketch of the nature of causal/mechanical explanation (1984). Salmon distinguishes his causal/mechanical account of explanation from Hempel's (1965) covering law model by noting that whereas Hempel's explanations explain an event by showing *that* it fits into a *nomic* nexus, causal mechanical explanations explain by showing *how* an event fits into a *causal* nexus. Salmon's causal nexus is composed of causal processes interacting with one another. On Salmon's causal/mechanical model, explaining an event (or type of event[8]) is a matter of situating that event within this geometrical network of causal processes and interactions (perhaps supplemented by statistical relevance relations linking causal processes and interactions to the explanandum event).[9]

8. Salmon intends this image to work for both singular events and general regularities, except that in the case of regularities, "we apply precisely the same considerations to any volume of space-time that is similar in relevant respects. The relevant similarities are given by the nature of the regularity that we are trying to explain." (1984, 275) A similar move can accommodate general contextual explanations as well.

9. Hitchcock (1995, 1996) has charged that Salmon's rendition of the causal nexus is too austere to provide an adequate picture of causal/mechanical explanation. Salmon (1997) recommends using statistical relevance relations to augment the barren geometrical structure of the causal nexus. This move is explored in Craver (1998, Ch. 4).

Salmon recognizes two ways of situating an event within a causal nexus, and he refers to these as two distinct "aspects" of causal mechanical explanation: an etiological aspect and a constitutive aspect. These are depicted in Figure 3(a), modeled upon Salmon's diagram (1984, 275). The *etiological aspect* of an explanation for some event "fill[s] in the causally relevant processes and interactions that occupy the past light cone" of that event. (Salmon 1984, 275) Etiological explanations are backward looking, revealing the relevant portions of the causal nexus in the event's past. We provide an etiological explanation of why John is a victim of heart disease when we blame his smoking and diet and, perhaps, the mechanisms by

a)

b)

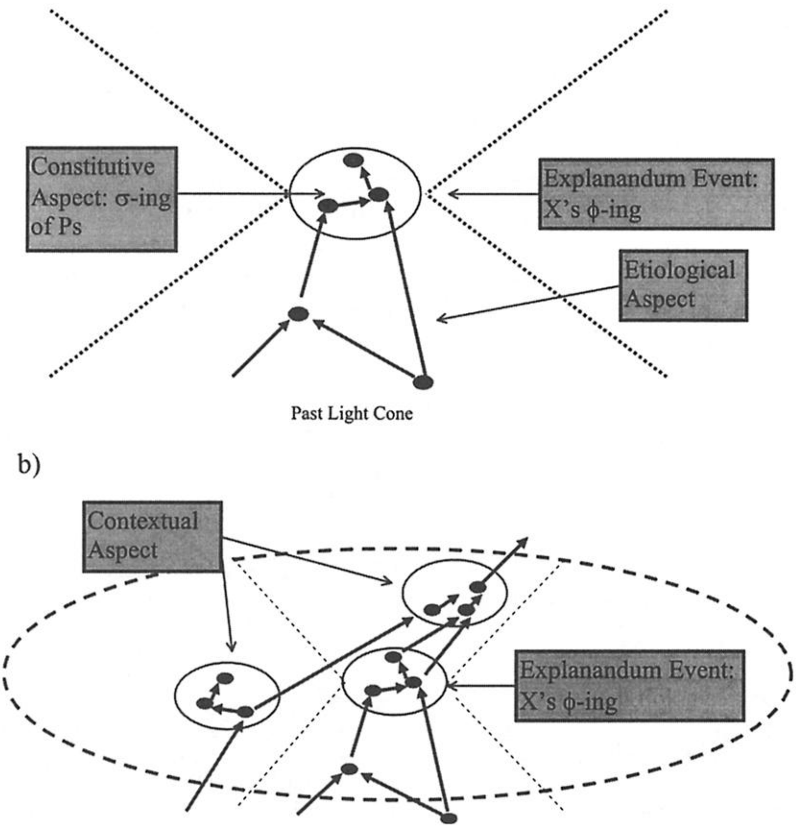

Figure 3. Constitutive, etiological, and contextual explanations. Figure 3(a) is modeled upon Salmon (1984, 275) and exhibits the etiological and constitutive aspects mechanistic explanation. Figure 3(b) superimposes the contextual aspect onto Salmon's diagram.

which smoking and diet produce heart disease. The explanandum is an event or a type of event and the explanans reveals the antecedent mechanisms by which the event occurred/occurs.

The *constitutive aspect* of an explanation for some event, "lays bare the causal structure" of the event by revealing "the internal causal mechanisms" that account for the event's "nature" (Salmon 1984, 275). Constitutive explanations describe, in the language of Section 4, the lower-level mechanisms of X's ϕ-ing. Constitutive explanations are inward looking and downward looking, looking within the boundaries of X to determine the lower level mechanisms by which it can ϕ. The explanandum of a constitutive explanation is the ϕ-ing of an X, and the explanans is a description of the organized σ-ing of Ps. Constitutive explanations explain by showing how the ϕ-ing of X fits into the portion of the causal nexus made up of the organized activities of X's component parts.

If one accepts Salmon's broad outline of causal/mechanical explanations (i.e., that they explain by showing how an item or event fits into a nexus of causes), then there is a clear candidate for a third aspect or variety of causal/mechanical explanation. This is a *contextual* aspect or variety of causal/mechanical explanation, the details of which were discussed in Section 4. A contextual explanation explains an entity or activity by showing what it is for, that is, how it fits into the organization of a higher-level mechanism. Contextual explanations are superimposed on Salmon's two aspects of causal/mechanical explanation in Figure 3(b). As the figure illustrates, contextual explanations explain X's ϕ-ing by tracing out the portion of a causal nexus to which X's ϕ-ing contributes. Contextual explanations are characteristically outward looking and upward looking. They are outward looking because they refer to components outside of X; they are upward looking because they contextualize X within a higher-level mechanism (S). Contextual explanations quite literally show how an entity or activity fits into a mechanism.[10] So there are not two, but three

10. One might object to the fact that contextual explanations for some item include reference to components that are at later stages in the mechanism, and hence that come after, or are the consequences of, that item. Much of the worry about the legitimacy of teleological explanation has traditionally been that (at least some kinds of) teleological explanations posit an occult causal influence from later stages in the mechanism to earlier stages of the mechanism. Causal/mechanical explanations respect the asymmetry of causation. Earlier stages produce, allow, or otherwise influence later stages, but later stages, excepting unobjectionable cases of feedback, cannot produce earlier stages.

There is no occult, later-to-earlier causal influence in contextual mechanistic explanations. Contextual explanation is typically given against the presumed backdrop of a contextual mechanism: either an etiological mechanism by which some particular (or type of) end state comes about, or a constitutive mechanism by which some particular (or type of) higher-level activity is carried out. Examples of contextual explanation in etiological mechanisms are familiar in the medical sciences (e.g., explaining the role of different types of causal mechanical explanations corresponding to three distinct ways of situating an item into a nexus of mechanisms: etiological, constitutive, and contextual.

Since any given X or ϕ may contribute to the diverse ψ-ings of any number of Ss, contextual explanations are ineliminably perspectival. That is, they rely upon shared background assumptions that S can ψ or that the ψ-ing of S is important, significant, or relevant. This perspectival feature of role functions has been taken as a sign of weakness. Learning that the heart pumps blood seems to explain what the heart is for in a way that learning that the heart makes glub-blub noises does not. But this weakness reflects instead a breakdown in the shared background assumptions required for contextual explanation to get hold. Either we find it difficult to conjure the ψ-ing of an S to which the glub-blub noises contribute, or we find it difficult to conjure such an S whose ψ-ing is important, significant, or relevant. The idea that the heart is for making glub-blub noises only seems absurd until one is able to conjure (often with some contortion) a suitable mechanistic context for those glub-blub noises (like a diagnostic mechanism or perhaps the mechanisms of fetal auditory development; see Amundson and Lauder 1994). This perspectival feature should not be surprising; it is not unique to contextual causal/mechanical explanations, but can also be found in constitutive explanations (as emphasized by Kauffman 1971) and etiological explanations (as argued by van Fraassen 1980).

Walsh and Ariew (1996) propose a taxonomy of senses of "function" based on the different "causal," "etiological," and "teleological" explanatory uses to which that notion is put. The first two of these (causal and etiological) correspond, respectively, to constitutive and etiological explanations. The last, teleological explanations, "explain the prevalence and (or) persistence of trait types by citing their causal contribution to average fitness of individuals" (513). Like Bechtel (1986), Walsh and Ariew do not recognize the possibility of a contextual variety of explanation that is broader than (though inclusive of) teleological explanation. What is required is a sense of "function," and of causal/mechanical explanation, that is couched in terms of the contribution of an item to some higher-level mechanism, irrespective of the contribution of that mechanism to fitness, adaptation, or the good life. Walsh and Ariew do not recognize the possibility of contextual explanation, and I think this is because they have

cholesterol in heart attacks or a heart attack) and in historical research (e.g., explaining the role of Vesalius' *De Fabrica* in Harvey's discovery of the circulation of the blood). More importantly for the biological sciences, contextual explanations appear in descriptions of constitutive mechanisms such as the mechanism of the kidney's filtration of the blood or the mechanism of the heart's contraction. These contextual explanations cite factors at later stages of the mechanism, but they do not cite them *as* productive of X's ϕ-ing.

not paid sufficient attention to the importance the multilevel structure of mechanistic theories and causal/mechanical explanations. Once this unilevel bias is corrected, the place for contextual explanation is apparent enough.

Hardcastle (1999) has taken some initial steps towards working out a pragmatics of functional explanation that fits nicely with the view of contextual explanation that I have been developing. She correctly emphasizes that the choice of a contextual mechanism or "goal state," for the highest level mechanism depends on accepted theories, the research community, the background context of the research, and the biases of individual scientists (1999, 39). But Hardcastle criticizes Cummins on the grounds that (i) his analysis is restricted to "how a system actually behaves" and hence "misdescribes the functions of malformed or broken things" and (ii) his analysis "overlooks interspecies variation" since even minor differences in containing systems entail differences in function even if the parts have the same role (described at some degree of abstraction). (Hardcastle 1999, 36–37) Mechanistic role functions may seem to be especially vulnerable to such criticisms, and so I close by suggesting some ways around them that are, in the end, consistent with Hardcastle's pragmatic project.

The first objection can be met by noting that the ascription of a function to a malformed or broken part is derivative upon a description of how that *type* of part (X) fits into a *type* of higher-level mechanism (S). The malformed and broken part can be identified as an X by the typical properties and activities of Xs (within which there may be considerable variation even among properly formed parts). Parts can be identified by, for example, their size, shape, location, composition, and development, and by their diverse properties and activities. Malformed parts will not share all of these properties, and if they share too few they may become unrecognizable as Xs, but this is most often not the case. (More extensive arguments of this sort can be found in Amundson and Lauder 1994). The mechanistic role of the broken part only appears against the fixed backdrop of shared assumptions about a type of mechanism within which parts of this type generally (or preferably) make important contributions. A broken kidney, for example, can still be identified by its position in the torso, its connections with the renal vein and artery, its being composed of nephra, and its characteristic shape, color, and size. The kidney's mechanistic role is then identified against the fixed backdrop of a description of the way the circulatory system generally works, or the way that it preferably works, or the way that it works in whatever (normal or pathological) mechanism that we seek to understand.

Hardcastle's second objection similarly points to the need for close attention to the varying degrees of abstraction and generality of biological generalizations (see e.g., Schaffner 1993; Darden 1996) and to the fuzziness of biological kinds (see, e.g., Waters 1998). Even slight differences in mechanistic context entail different mechanistic role functions. But the need to keep track of such differences is not unique to contextual explanations; the need also arises when we ask whether two events have the *same* etiological explanation or when two events have the *same* constitutive explanation. Judgements of "sameness" in these cases depend upon an agreed-upon tolerance of diversity among tokens within types. For this reason, talk of close or distant analogy, both among mechanisms and among mechanistic roles, is more appropriate than talk of "sameness" for discussing interspecies comparisons. Both for roles and for mechanisms, closeness of analogy depends upon the similarity of their components and similarity in their active, spatial, and temporal organization.

6. Conclusion. The discovery of an item's mechanistic role is considered a first rate scientific achievement. This is because role ascriptions help to integrate the levels in multilevel hierarchies and because learning an item's role provides a kind of understanding of that item—a contextual explanation of how that item (or type of item) fits into a nexus of mechanisms.

Describing an item's mechanistic role is a perspectival affair. This perspectival take on functional ascription should be a reminder that what we take as functional descriptions can be tinged in a very direct way by our interests and biases (see e.g., Amundson 2000; Gould 1981). Multilevel mechanisms are framed relative to a shared topping off point. Perhaps grounding functional description in the details of mechanistic organization will provide a set of criteria for assessing the precision and accuracy of functional ascriptions and will perhaps help to guard against empirically inadequate, vague, or overly abstract functional ascriptions. In an age of obesity genes and humor centers, any clarity concerning the process of assigning mechanistic roles should be a welcome contribution. Here, I have tried to sketch a framework within which this clarification can take place.

REFERENCES

Allen, Colin, Mark Bekoff, and George V. Lauder (eds.) (1998), *Nature's Purposes*. Cambridge, MA: MIT Press.

Amundson, Ron (2000), "Against Normal Function", *Studies in the History and Philosophy of the Biological and Biomedical Sciences* 31: 33–53.

——— and George V. Lauder (1994), "Function without Purpose: The Uses of Causal Role Function in Evolutionary Biology", *Biology and Philosophy* 9: 443–469. Reprinted in Allen, Bekoff, and Lauder 1998, 335–370.

Bechtel, William (1986), "Teleological Functional Analyses and the Hierarchical Organization of Nature", in N. Rescher (ed.), *Teleology and Natural Science*. Landham, MD: University Press of America, 26–48.

——— and Robert C. Richardson (1993), *Discovering Complexity: Decomposition and Localization as Strategies in Scientific Research*. Princeton: Princeton University Press.

Churchland, Patricia S. and Terrence Sejnowski (1992), *The Computational Brain*. Cambridge, MA: MIT Press.

Craver, Carl F. (1998), *Neural Mechanisms: On the Structure, Function and Development of Theories in Neuroscience*. Ph.D. Dissertation, Pittsburgh: University of Pittsburgh.

——— and Lindley Darden (forthcoming), "Discovering Mechanisms in Neurobiology: The Case of Spatial Memory", in P. Machamer, R. Grush, and P. McLaughlin (eds.), *Theory and Method in Neuroscience*. Pittsburgh, PA: University of Pittsburgh Press.

Cummins, Robert (1975), "Functional Analysis," *Journal of Philosophy* 72: 741–765. Reprinted in Allen, Bekoff, and Lauder 1998, 169–196. Page references are to the reprint.

——— (1983), "Analysis and Subsumption in the Behaviorism of Hull", *Philosophy of Science* 50: 96–111.

Darden, Lindley (1991), *Theory Change in Science: Strategies from Mendelian Genetics*. New York: Oxford University Press.

——— (1996), "Generalizations in Biology", *Studies in the History and Philosophy of Science* 27: 409–419.

Glennan, Stuart S. (1996), "Mechanisms and the Nature of Causation", *Erkenntnis* 44: 49–71.

——— (1997), "Capacities, Universality, and Singularity", *Philosophy of Science* 64: 605–626.

Gould, Stephen J. (1981), *The Mismeasure of Man*. New York: W.W. Norton and Company.

Hardcastle, Valerie G. (1999), "Understanding Functions", in Valerie G. Hardcastle (ed.), *Where Biology Meets Psychology*. Cambridge, MA: MIT Press, 27–43.

Harvey, William ([1628] 1963), *Movement of the Heart and Blood in Animals: An Anatomical Essay*. Translated by K.J. Franklin. London: Dent.

Haugeland, John (1998), *Having Thought: Essays in the Metaphysics of Mind*. Cambridge, MA: Harvard University Press.

Hempel, Carl G. (1965), *Aspects of Scientific Explanation*. London: Collier Macmillan Publishers.

Hitchcock, Christopher (1995), "Discussion: Salmon on Explanatory Relevance", *Philosophy of Science* 62: 304–20.

——— (1996), "The Mechanist and the Snail", *Philosophical Studies* 84: 91–105.

Kauffman, Stuart A. (1971), "Articulation of Parts Explanation in Biology and the Rational Search for Them", *Boston Studies in the Philosophy of Science* 8: 257–272.

Lycan, William G. (1990), "The Continuity of Levels of Nature", in W. G. Lycan (ed.), *Mind and Cognition: A Reader*. Oxford: Bradford, 77–96.

Machamer, Peter K. (1977), "Teleology and Selective Processes", R. G. Colodny (ed.), *Logic, Laws, and Life*, Vol. 6. Pittburgh, PA: University of Pittsburgh Press.

——— Lindley Darden, and Carl F. Craver (2000), "Thinking about Mechanisms", *Philosophy of Science* 67: 1–25.

Rescher, Nicholas (1955), "Axioms for the Part Relation", *Philosophical Studies* 6: 8–10.

Salmon, Wesley C. (1984), *Scientific Explanation and the Causal Structure of the World*. Princeton: Princeton University Press.

——— (1997), "Causality and Explanation: A Reply to Two Critiques", *Philosophy of Science*. 64: 461–477.

Schaffner, Kenneth (1993), *Discovery and Explanation in Biology and Medicine*. Chicago: University of Chicago Press.

Simon, Herbert A. (1969), *The Sciences of the Artificial*. Cambridge: MIT Press.

van Fraassen, Bas C. (1980), *The Scientific Image*. Oxford: Clarendon Press.

Walsh, Denis M. and Andre Ariew (1996), "A Taxonomy of Functions", *Canadian Journal of Philosophy* 26: 493–514.

Waters, Ken C. (1998), "Causal Regularities in the Biological World of Contingent Generalizations", *Biology and Philosophy* 13: 5–36.

Wimsatt, William (1974), "Complexity and Organization", in Kenneth Schaffner and Robert S. Cohen (eds.), *PSA 1972*, 67–86.

——— (1986), "Forms of Aggregativity", in Alan Donagan, Anthony N. Perovich, Jr., and Michael V. Wedin (eds.), *Human Nature and Natural Knowledge*. Dordrecht: Reidel, 259–291.

——— (1997), "Aggregativity: Reductive Heuristics for Finding Emergence", *Philosophy of Science* 64 (Proceedings): S372-S384.