

AD A101491

DTIC FILE COPY

SECURITY CLASS. OF THIS PAGE (When Data Entered)

REPORT DOCUMENTATION PAGE		READ INSTRUCTIONS BEFORE COMPLETING FORM
1. REPORT NUMBER 81-7D, (14) HEST-75-81	2. GOVT ACCESSION NO. A101491	3. RECIPIENT'S CATALOG NUMBER
4. TITLE (and Subtitle) The Logic of Concepts: Case Studies in Engineering and Law		5. TYPE OF REPORT & PERIOD COVERED THESIS/DISSERTATION
7. AUTHOR(s) James H. Parsons		6. PERFORMING ORG. REPORT NUMBER
9. PERFORMING ORGANIZATION NAME AND ADDRESS AFIT STUDENT AT: The Ohio State University		8. CONTRACT OR GRANT NUMBER(s)
11. CONTROLLING OFFICE NAME AND ADDRESS AFIT/NR WPAFB OH 45433		10. PROGRAM ELEMENT, PROJECT, TASK AREA & WORK UNIT NUMBERS (10) 2007
14. MONITORING AGENCY NAME & ADDRESS (if different from Controlling Office)		12. REPORT DATE (14) 1981
		13. NUMBER OF PAGES 201
		15. SECURITY CLASS. (of this report) UNCLASS
		15a. DECLASSIFICATION/DOWNGRADING SCHEDULE
16. DISTRIBUTION STATEMENT (of this Report) APPROVED FOR PUBLIC RELEASE; DISTRIBUTION UNLIMITED		
17. DISTRIBUTION STATEMENT (of the abstract entered in Block 20, if different from Report)		
18. SUPPLEMENTARY NOTES APPROVED FOR PUBLIC RELEASE: IAW AFR 190-17 23 JUN 1981		DTIC ELECTE JUL 17 1981 F <i>Fredric C. Lynch</i> FREDRIC C. LYNCH, Major, USAF Director of Public Affairs Air Force Institute of Technology (ATC) Wright-Patterson AFB, OH 45433
19. KEY WORDS (Continue on reverse side if necessary and identify by block number)		
20. ABSTRACT (Continue on reverse side if necessary and identify by block number) ATTACHED		

DD FORM 1 JAN 73 1473

EDITION OF 1 NOV 65 IS OBSOLETE

UNCLASS

SECURITY CLASSIFICATION OF THIS PAGE (When Data Entered)

S 1-7 D

THE LOGIC OF CONCEPTS:
CASE STUDIES IN ENGINEERING AND LAW

By

James H. Parsons, B.S., M.A.

Military Rank: Major

Number of Pages: 201

Service Branch: USAF


Degree Awarded: Ph.D.

Date: 1981

Institution: The Ohio State University

It has traditionally been supposed that concepts are learned from definitions or analyses; however, many contemporary philosophers and psychologists think that concepts are taught by means of exemplars, particular items which the community accepts as representative exemplifications of a concept. But this view is incomplete without a theory of how the transition is made from viewing group-licensed exemplars to possessing a concept which is shared with one's linguistic community. Further, although there is general agreement that the concepts of a given society may change with time, there is not agreement on how this conceptual evolution occurs and whether or not it is a rational process.

The dissertation contributes to a resolution of these issues by arguing for a particular descriptive theory of the dynamic processes of conceptual activity. According to the theory proposed, a concept is a

particular type of capacity. An idealized model of the items in a given exemplar set bridges the gap between examining the exemplars and possessing the associated capacity. Such models differ from any actual exemplars and, thus, stand in need of justification. 

Case studies of the development of the concept of a feedback control system in engineering and the concept of liability for negligence to noncontracting parties in law reveal some types of arguments which are employed in justifying conceptual models. The same types of arguments are also employed in choosing among competing models and in justifying extension of a concept to nonexemplars. There is a second class of arguments: those which use a concept's record of success or failure to justify alterations in the background beliefs and exemplar set which support the associated model.

Accession For	
ERIC GRA&I	<input checked="checked" type="checkbox"/>
ERIC T'S	<input type="checkbox"/>
Unpublished	<input type="checkbox"/>
Justification	
Distribution/	
Availability Codes	
Avail and/or	
Special	
A	

AFIT RESEARCH ASSESSMENT

The purpose of this questionnaire is to ascertain the value and/or contribution of research accomplished by students or faculty of the Air Force Institute of Technology (AFIT). It would be greatly appreciated if you would complete the following questionnaire and return it to:

AFIT/NR
Wright-Patterson AFB OH 45433

RESEARCH TITLE: The Logic of Concepts: Case Studies in Engineering and Law
(81-70)

AUTHOR: James H. Parsons

RESEARCH ASSESSMENT QUESTIONS:

1. Did this research contribute to a current Air Force project?
☐ a. YES ☐ b. NO
2. Do you believe this research topic is significant enough that it would have been researched (or contracted) by your organization or another agency if AFIT had not?
☐ a. YES ☐ b. NO
3. The benefits of AFIT research can often be expressed by the equivalent value that your agency achieved/received by virtue of AFIT performing the research. Can you estimate what this research would have cost if it had been accomplished under contract or if it had been done in-house in terms of manpower and/or dollars?
☐ a. MAN-YEARS ☐ b. \$
4. Often it is not possible to attach equivalent dollar values to research, although the results of the research may, in fact, be important. Whether or not you were able to establish an equivalent value for this research (3. above), what is your estimate of its significance?
☐ a. HIGHLY SIGNIFICANT ☐ b. SIGNIFICANT ☐ c. SLIGHTLY SIGNIFICANT ☐ d. OF NO SIGNIFICANCE
5. AFIT welcomes any further comments you may have on the above questions, or any additional details concerning the current application, future potential, or other value of this research. Please use the bottom part of this questionnaire for your statement(s).

NAME	GRADE	POSITION
------	-------	----------

ORGANIZATION	LOCATION
--------------	----------

STATEMENT(s):

THE LOGIC OF CONCEPTS:
CASE STUDIES IN ENGINEERING AND LAW

DISSERTATION

Presented in Partial Fulfillment of the Requirements for
the Degree Doctor of Philosophy in the Graduate
School of The Ohio State University

By

James H. Parsons, B.S., M.A.

* * * * *

The Ohio State University

1981

Reading Committee:

Richard T. Garner

Ronald E. Laymon

George S. Pappas

Approved By



Advisor
Department of Philosophy

This dissertation is dedicated to my mother
and father, neither of whom got to see it.

ACKNOWLEDGMENTS

In the course of writing this dissertation, I have incurred intellectual debts to many friends and colleagues at both The Ohio State University and the United States Air Force Academy. In particular, I wish to acknowledge the help and support of my advisor and friend, Ronald E. Laymon. Without his encouragement and advice, this dissertation would not have been written. I shall not list other individuals on either faculty by name, for fear of inadvertently omitting someone important, but I have profited from my associations with each of them.

My wife, Celeste, has shared with me all the hardships and satisfactions of service life, and her support and encouragement has been invaluable to me.

I did the research and writing of this dissertation while on an Air Force Institute of Technology assignment. The United States Air Force has sponsored all of my schooling since I entered the Air Force Academy as a fourth classman in 1962.

The agencies listed below have given their permission to quote from the works indicated:

1. American Society of Naval Engineers: Minorsky, N. "Directional Stability of Automatically Steered Bodies." Journal of the American Society of Naval Engineers 42 (1922):280-309.

2. Journal of the Franklin Institute: Hazen, Harold L. "Theory of Servo-Mechanisms." Journal of the Franklin Institute 218 (1934): 279-331.

3. The M.I.T. Press: Mayr, Otto. The Origins of Feedback Control. English ed. Cambridge: The M.I.T. Press, 1970.

4. The Royal Society of London: Callender, A.; Hartree, D. R.; and Porter, A. "Time Lag in a Control System." Philosophical Transactions of the Royal Society of London A235 (1936):415-44.

5. The University of Chicago Press: Kuhn, Thomas S. The Essential Tension: Selected Studies in Scientific Tradition and Change. Chicago: The University of Chicago Press, 1977.

6. The University of Chicago Press: Kuhn, Thomas S. The Structure of Scientific Revolutions. 2d ed., enl. Chicago: The University of Chicago Press, 1970.

VITA

February 4, 1944	Born - Annapolis, Maryland
1966	B.S., United States Air Force Academy, Colorado
1967-1969.	Copilot, KC-135 "Stratotanker," Strategic Air Command
1969-1971.	Pilot, UH-1F "Huey," Strategic Air Com- mand
1971-1974.	Instructor Pilot, HH-3E "Jolly Green Giant," Pacific Air Forces and Military Airlift Command
1976	M.A., The Ohio State University, Colum- bus, Ohio
1976-1978.	Instructor, Department of Philosophy and Fine Arts, United States Air Force Aca- demy, Colorado
1978	Assistant Professor, Department of Philosophy and Fine Arts, United States Air Force Academy, Colorado

TABLE OF CONTENTS

	Page
DEDICATION	ii
ACKNOWLEDGMENTS.	iii
VITA	v
LIST OF TABLES	viii
LIST OF FIGURES.	ix
INTRODUCTION	1
 Chapter	
I. KUHNIAN PARADIGMS.	5
Kuhn's Reason for Postulating the Existence of Paradigms.	5
The Structure of Scientific Revolutions	9
The Automatic Feedback Control System as a Kuhnian Paradigm.	10
Mayr's Putative Necessary and Sufficient Conditions.	12
II. THE ROLE OF EXPLANATION IN CLASSIFICATION.	29
Putnam's Observations about Natural Kind Terms.	29
Unpacking Putnam's Notion of Hidden Structures.	36
Explanations and Models in Conceptualization.	45
An Illustration: the Acid-Base Research Program	51
III. THE DYNAMICS OF CONCEPTUALIZATION.	56
The Justification of Idealizations in Control Theory.	57
Extending a Concept	75
IV. JUDICIAL CLASSIFICATION.	100
Precedents as Paradigm	100
The Requirement for a Model of <u>L</u>	103

	Page
<u>Langridge v. Levy</u>	107
<u>Winterbottom v. Wright</u>	113
An L- Model is Accepted: <u>Longmeid v.</u> <u>Holliday</u>	118
A Puzzle: Retroactive Reasoning?.	123
Competition among Classificatory Explanations . .	124
Accepting a Precedent	126
The Conservative Nature of the Parke Model. . . .	135
A Methodological Consideration.	138
The Parke Model in America: <u>Thomas v.</u> <u>Winchester</u>	140
Inherently Dangerous Articles	142
An Ambiguity in the Parke Model	145
<u>MacPherson v. Buick Motor Co.</u>	147
Motivating a Model Change	155
The Development of Systems of Classification. . .	157
 V. RATIONAL CONCEPTUAL ACTIVITY.	 168
 BIBLIOGRAPHY.	 196

LIST OF TABLES

Table	Page
1. Comparison of Significant Features of Four <u>L</u> Models. .	136

LIST OF FIGURES

Figure	Page
1. Block Diagram of Classification Process	97
2. Block Diagram of Conceptual Activity.	193

INTRODUCTION

Concepts are capacities for classifying or acting selectively toward items in our ken. More accurately, concepts are capacities for acquiring states which in turn permit the person who is in those states to treat certain things as distinct from certain others.¹ As an analogy, if we provide a computer with the appropriate sensors, then a certain program or arrangement of circuitry will give it the capacity to respond to red stimulae in a way different from its non red stimulae responses. In this dissertation, what I shall call the "concept" of red is not what corresponds to the sensors, the circuit, nor the program, but what corresponds to the ability which these things engender.

As one gains a repertoire of relatively basic concepts, he gains the capacity for ever more complex concepts compounded from the basic ones. For example, having the concepts, red and triangle, along with concepts of various logical connectives, enables us to respond selectively to red triangles, nonred triangles, red nontriangles, things which are either red or triangles, things which are either nonred or triangular, and so forth.

There are two competing accounts of these skills or capacities which we call concepts. The first is that the capacity consists in recognizing whether or not some set of necessary and sufficient conditions for the application of the concept term has been satisfied. When we acquire a concept, we learn the rules for the application of the

concept term. The second account is that the capacity consists in recognizing whether or not something is appropriately similar to certain central cases of application of the concept term. When we acquire a concept on this view, we learn to recognize the appropriate similarity relations, and we learn the central cases.

These two accounts concern standardization, in the sense that they attempt to answer questions about how we can teach concepts already in use in our linguistic community and what it is that we have learned when we join our community in the use of a certain concept. But concepts are tools which we can use well or poorly. It is more valuable for us to be able to make certain discriminations than others, and we want to discover those concepts which will be the most valuable to us. There are parallel accounts to the two theories of standardization already mentioned which attempt to point the way toward optimal concepts. According to the first theory of optimization, our task is one of altering the rules of application in response to our successes and failures, until we achieve concepts which divide the world cleanly into the sets of items which it is the most useful for us to recognize. Perhaps these optimal divisions correspond to "natural" divisions in the world; perhaps there are no such partitions. The important thing is that the activity begins with our tentative formulations of rules, and the end product is a set of optimal rules for classifying things.

In contrast, the second theory would have us begin with a set of things which we suppose to be representative of a useful kind of classification. Optimization consists in learning to spot similarities

to the exemplars which are relevant and "steering" the whole enterprise by modifying the exemplar set. In The Structure of Scientific Revolutions, Thomas Kuhn attempts to apply theories of standardization and optimization of the second type to the scientific enterprise.

Kuhn concerns himself with the manner in which concept standardization and optimization work in the basic sciences. In what follows, I argue that the complex concept of a feedback control system is an example from outside of the basic sciences of a paradigm, i.e., a very important and fruitful concept which is developed from a set of exemplars and which is taught to others by means of exemplars. Using data from the early development of the feedback control system paradigm, I develop an account of the manner in which we attempt to optimize our classificatory capacities as we interact with the world. The account is then extended to account for our use of certain abstract concepts in law.

I begin by summarizing briefly Kuhn's remarks about paradigms and the role they play in the basic sciences. By holding up certain aspects of the feedback control case for comparison with Kuhn's proposals about paradigms, I attempt to show that the feedback control concept may be considered a Kuhnian paradigm. Unfortunately, the discussion also serves to show that Kuhn's comments about paradigms and the role they play are more suggestive than informative. Kuhn has identified an interesting feature of life in the scientific community and given it a name. The problems of exactly how paradigms operate and what might justify their playing the key role which they apparently do

are left unresolved by Kuhn. My work is an effort to contribute to the solution of these problems.

In the second chapter, I use some suggestions by Hilary Putnam to begin developing a descriptive theory of rational conceptual activity. The third and fourth chapters supply important details for the theory, gleaned from case studies of conceptual activity in engineering and law. In the final chapter, the new material from the case studies is included in the theory, and the theory is shown to be of value in filling out Kuhn's comments about the role played by paradigms in our conceptual activity.

CHAPTER I

KUHNIAN PARADIGMS

Kuhn's Reason for Postulating the Existence of Paradigms

While doing the research which culminated in The Structure of Scientific Revolutions, Thomas Kuhn noticed an interesting phenomenon. There are a number of ways of describing this phenomenon; here are a couple of Kuhn's:

. . . the relatively unproblematic character of professional communication and . . . the relative unanimity of professional judgment.²

No conjunction of particular symbolic forms would exhaust what the members of a scientific community can properly be said to know about how to apply symbolic generalizations. Confronted with a new problem, they can often agree on the particular symbolic expression appropriate to it, even though none of them has seen that particular expression before.³

To account for his observations of this phenomenon, Kuhn supposed that there must be rules of application which the scientists learn early in their professional education. Such rules would--

1. give meaning to a theoretical vocabulary by attaching it to a previously meaningful basic vocabulary;
2. show how the terms of a basic vocabulary attach directly to nature by giving a set of criteria which provide definitions of the terms or necessary and sufficient conditions for the proper application

of the terms.⁴

But when one goes into a given scientific community and attempts to discover the rules which the various members of the community are presumed to be applying, the result is "continual and deep frustration."⁵ Although some of the rules can be formulated to everyone's satisfaction, others will defy formulation. A rule that seems right to some members of the community will seem too strong to others. The members of the scientific community themselves cannot produce the rules that they are using. According to Kuhn:

If asked by a philosopher to provide such rules, scientists regularly deny their relevance and thereafter sometimes grow uncommonly inarticulate. When they cooperate at all, the rules they produce may vary from one member of the community to another, and all may be defective.⁶

Of course, Kuhn's inability to isolate any complete sets of rules of application does not prove that there are no such rules. The proposition that for every possible scientific community there is some complete set of such rules is, because of its universal-existential form, impossible either to prove or disprove by experimental evidence. On the other hand, Kuhn's inability to discover even a moderately comprehensive set of rules of application for even one term in one scientific community seems to indicate that the conjecture that rules of application account for the high level of agreement within a scientific community is not a very fruitful hypothesis. In short, Kuhn writes, "one begins to wonder whether more than a few such rules are deployed in community practice, whether there is not some alternate way in which scientists correlate their symbolic expressions with nature."⁷

In probing for an alternative explanation, Kuhn draws upon a well-known section of the Philosophical Investigations in which Ludwig Wittgenstein attacks the notion that an essential part of knowing or having a concept is always knowing exactly what it is that two or more particular items have in common, being able to specify necessary and sufficient conditions for the correct application of the concept. Instead, Wittgenstein says, sometimes what we learn when we acquire a concept is to recognize family resemblances, to see that two particulars share enough of the same properties that they may be paired, i.e., we learn to use a particular scheme for dividing up the world.

Wittgenstein says:

Compare knowing and saying:
 how many feet high Mont Blanc is--
 how the word "game" is used--
 how a clarinet sounds.

If you are surprised that one can know something and not be able to say it, you are perhaps thinking of a case like the first. Certainly not of one like the third.⁸

Kuhn judges that the case he is trying to account for is like the second of Wittgenstein's examples: scientists in the various branches of physical science know how to apply the various formulations of $F = ma$, but they are not able to explain precisely their criteria for applying them.

How do we learn how the word, "game," is used? How do we acquire the concept of a game? By being shown certain applications of the word which are accepted as being correct by all members of the linguistic community: chess, tennis, baseball, monopoly, hide-and-seek. Kuhn's proposal then is that members of a scientific community learn

how the important terms of their discipline are used in just the same way: by being shown certain applications of the terms which are accepted as being correct by all members of the particular scientific community.

Kuhn finds evidence for this conjecture by identifying the mechanism by which this showing of exemplary cases is accomplished:

Students of physics regularly report that they have read through a chapter of their text, understood it perfectly, but nonetheless had difficulty solving the problems at the end of the chapter. Almost invariably their difficulty is in setting up the appropriate equations, in relating the words and examples given in the text to the particular problems they are asked to solve. Ordinarily, also, these difficulties dissolve in the same way. The student discovers a way to see his problem as like a problem he has already encountered. Once that likeness or analogy has been seen, only manipulative difficulties remain.⁹

So Kuhn's hypothesis is that we can account for the extraordinarily high level of agreement among practitioners in a particular scientific community by the fact that they have all learned the use of the difficult terms, concepts, and symbolic generalizations by being shown the same exemplary cases of their application. Notice also that, on Kuhn's account, students are not just given the problems and their solutions in the text of their books; instead, problems are given as puzzles which have been set for the students to solve. In developing the ability to solve these puzzles, the students are not simply learning a set of rules and practicing their application; they are, as Kuhn says, acquiring the "ability to recognize group-licensed resemblances."¹⁰ They are not simply learning information; rather, they are developing a new way of ordering their experiences, a new set of *gestalts*.

The Structure of Scientific Revolutions

We can summarize the main points of The Structure of Scientific Revolutions in the following way. The development of science is not a steady process of accumulating knowledge but a series of abrupt changes. The Copernican and Einsteinian scientific revolutions were not aberrations in the practice of science but business as usual. Such revolutions occur regularly in every branch of science, but they often go unnoticed by all but the specialists whose fields are affected by the changes. The salient feature of these revolutions is that they concern changes in world view. That is, what a particular scientific revolution is about is the proper way of seeing the particular part of the world which is the concern of the field in which the revolution occurs. A major part of the "way of seeing" the world includes similarity relations: what is seen as like, or relevantly similar to, what. The revolution is complete when most of the active practitioners in the accepted discipline have adopted the new way of seeing things. The new world view is spread through the use of exemplars or group-licensed puzzle-solutions. With the new way of seeing, comes an entire disciplinary matrix which includes the types of problems and types of solutions which will be recognized as relevant in the new regime. Following a revolution, there is a period of great productivity in which adherents to the new way of seeing solve many of the puzzles suggested by the new order. Science during this period, Kuhn characterizes as "normal." Occasionally anomalies appear, but, rather than viewing them as defeating counterexamples, much effort is devoted to reconciling

them with the new order. Ultimately, the accepted world view has been patched up so much in order to reconcile anomalies which appeared within it that it becomes suspect. The discipline goes through a period of crisis in which the research activity lacks the direction and purposefulness which it has formerly had. Thus begins a period of "abnormal" or "revolutionary" science. One or more alternative world views are developed; after a period of competition, one of the competitors is accepted, the old order is given up, and the process begins anew.

The Automatic Feedback Control System
as a Kuhnian Paradigm

I want now to consider a particular candidate for a gestalt of the type which Kuhn seems to have in mind when he says, "The practice of normal science depends on the ability, acquired from exemplars, to group objects and situations into similarity sets which are primitive in the sense that the grouping is done without an answer to the question, 'Similar with respect to what?'"¹¹ The candidate which I have in mind is the automatic feedback control system.

A feedback control system is a control system which maintains some variable at a desired level by detecting deviations of the controlled variable from the desired level and makes corrections based on that information. For example, the typical household thermostat maintains a constant preselected temperature during the winter by turning the heater off when the temperature in the house is above the desired temperature and turning it on when the temperature is below that desired. The temperature in the house determines when the heater will run, and the operation of the heater determines the temperature in the house.

The feedback control idea is a powerful one, and engineers have found it useful in such diverse fields as chemical processing, paper making, iron and steel production, machining, assembly, communications, transportation, management, research, and education.¹² In medicine and associated fields, many neural, glandular, and other bodily activities have been successfully modeled as feedback control systems.¹³ The way in which the concept of feedback control has found application in many different fields seems analogous to the way in which $\underline{F} = \underline{ma}$ has been extended to hundreds of different applications in dozens of different fields. Borrowing Kuhn's examples:

For the case of free fall, $\underline{f} = \underline{ma}$ becomes $\underline{mg} = m \frac{d^2 s}{dt^2}$; for the simple pendulum it is transformed to $\underline{mg} \sin \theta = -m l \frac{d^2 \theta}{dt^2}$; for a pair of interacting harmonic oscillators it becomes two equations, the first of which may be written $m \frac{d^2 s_1}{dt^2} + k_1 s_1 = k_2 (s_2 - s_1 + d)$; and for more complex situations, such as the gyroscope, it takes still other forms, the family resemblance of which to $\underline{f} = \underline{ma}$ is still harder to discover.¹⁴

But the feedback control idea has been applied to an even greater variety of situations.

In a typical textbook on feedback control engineering, we find the types of puzzles which Kuhn describes as those which might help a student to come to see certain types of situations as ones in which feedback is present or in which a feedback control system would be appropriate.¹⁵ One of the distinguishing and interesting features of the feedback control paradigm is that it is a paradigm for different specialities; it is one which is used not solely by mechanical

engineers or by electrical engineers, but it is also an important paradigm in fields rather far removed from normal engineering concerns, e.g., management and medicine.

Mayr's Putative Necessary
and Sufficient Conditions

Despite the sort of evidence which we have been considering which seems to support the conjecture that the notion of automatic feedback control is a family concept which is passed along to each new generation of control engineers (or managers or neurophysiologists) in just the manner that Kuhn describes, there may be a reason for thinking that the notion is not one which operates in this way. It is possible to give what seem to be necessary and sufficient conditions for something's being a feedback control system. In a book called The Origins of Feedback Control, Otto Mayr, the book's engineer-historian author, gives just such a set of conditions. He develops his conditions from a definition which was proposed in 1951 by the American Institute of Electrical Engineers:

A Feedback Control System is a control system which tends to maintain a prescribed relationship of one system variable to another by comparing functions of these variables and using the difference as a means of control.¹⁶

In developing his conditions, Mayr says that he wants to define the concept of feedback as rigorously as possible "in order to obtain an instrument with which we can irrefutably identify feedback control systems."¹⁷ Mayr's conditions are the following:

1. The purpose of a feedback control system is to carry out commands; the system maintains the controlled variable equal to the command signal in spite of external disturbances.

2. The system operates as a closed loop with negative feedback.
3. The system includes a sensing element and a comparator, at least one of which can be distinguished as a physically separate element.¹⁸

The first two of Mayr's conditions are merely paraphrases of the American Institute of Electrical Engineers' definition. Mayr says that these two conditions "are necessary to identify feedback systems but they do not suffice."¹⁹ He explains why he thinks the first two conditions are too weak:

Numerous systems exist where input and output are maintained in a "prescribed relationship," and where, either by physical reasoning or by mathematical formalism, a closed loop with negative feedback can be identified. Examples are analog computer programs for differential equations, or simple physical systems with self-regulation, such as the water level upstream of a weir, the R-C circuit, or the weather vane. Indeed, all systems in which the denominator of the transfer function consists of a polynomial containing an absolute member can be represented formally, by means of block diagram algebra, as closed loops with negative feedback.²⁰

With the addition of the third condition, Mayr claims that "we have now obtained a sufficiently complete definition of the concept."²¹ Thus, it is clear that Mayr intends his three conditions to constitute necessary and sufficient conditions for something's being a feedback control system. Finally, he wishes to limit his survey to the history of "automatic feedback control systems, in contrast to manual closed-loop control where the functions of comparison and control action are fulfilled by a human operator."²²

Mayr seems to have given us the makings of one of those rules of application (albeit a fairly high order one) which Kuhn claims to be so difficult to isolate. The rule would look something like the following:

C: It is correct to apply the term "automatic feedback control system" to any and all systems (sets of items)--

1. whose purpose is to carry out commands by maintaining a controlled variable equal to a command signal in spite of external disturbances;
2. which operate as a closed loop with negative feedback;
3. which include a sensing element and a comparator, at least one of which can be distinguished as a physically separate element;
4. whose functions of comparison and control are fulfilled by other than human operators.

These conditions can best be illustrated by considering the most famous application of the automatic feedback control idea, the centrifugal governor introduced in 1788 on the Boulton-Watt steam engine.²³ The purpose of this system is to maintain a preselected (command signal) engine speed (the controlled variable) despite varying loads on the engine (external disturbances). The engine speed is regulated by opening and closing a steam valve. The valve is opened and closed not according to a predetermined schedule, but in response to the speed of the engine (so it employs a closed, rather than open loop control system). If the engine slows below the desired running speed, the valve is opened to speed it up and vice-versa if the engine is running too fast. When there is a deviation from the desired condition, the control system generates a deviation in the opposite direction (negative feedback) of equal or near equal magnitude. To detect the engine speed, the Watt governor uses weighted balls at the ends of

arms on a shaft connected to the output shaft of the engine which are spun out by centrifugal force when the engine is running (the sensing element). The position of the arms is relayed by mechanical linkages to the steam valve. The valve is opened or closed depending on the position of the flyweights as relayed through the linkages, compared with their position at the desired running speed (the comparator). Although steam engines were built and operated which had no governors, the steam valves of these engines had to be opened and closed manually by human operators who monitored the effects of varying loads on the engine speed. The Watt centrifugal governor performed this tedious task automatically.

Let us return to the question of whether or not the notion of automatic feedback control works in the way that Kuhn describes or whether it is fully accounted for by the rule (C) given above. First, Kuhn's main thesis is not threatened, even if C is correct. He notes that at least some rules of application can be discovered.²⁴ His concern is that there don't seem to be enough of these rules to account for the agreement which he has observed. The second thing to be noticed is that, whether C is a correct rule for the use of "automatic feedback control system" or not, it was not generally known in the United States before 1970, the date of the English language edition of Mayr's book. Before that time, the prevailing notion of feedback control systems was something like that expressed by the definition given earlier in this chapter from the American Institute of Electrical Engineers, and that particular definition was not generally available before 1951.²⁵ The point is that, even if we can supply the necessary

and sufficient conditions for the correct application of "feedback control system," the result does not solve Kuhn's problem about accounting for agreement within a scientific community if the conditions we have discovered are not known and applied by members of that community.

Mayr's account of feedback control systems differs from the earlier definition chiefly because of his third condition, the stipulation that only systems which include sensing elements and comparators at least one of which can be distinguished as physically separate elements are to be considered true feedback control systems. His concern is that, without this condition, we would have to consider certain physical systems with self-regulation, such as the water level upstream of a weir or a weather vane, to be feedback control systems.²⁶ Mayr's amended criteria are supposed to rule out these as well as several artifacts which have been thought to be feedback control systems by other writers. These include "carburetors . . . equipped with automatic adjustments for such variables as engine speed, manifold or outside temperature, etc.,"²⁷ a device on an ancient Chinese chariot which keeps a human statue always pointing southward regardless of the orientation of the chariot,²⁸ an ancient Chinese drinking straw which permits wine to flow only at a certain rate,²⁹ a windmill with self-regulating sails,³⁰ and the centrifugal pendulum of Christiaan Huygens.³¹ Each of these has been classified by some authority as a feedback control system, and in most instances, Mayr has indicated the source for that classification. In a 1964 Scientific American article entitled "Control Theory," Richard Bellman, author of a number of books

and articles on feedback control systems, indicates that he views Huygens's centrifugal pendulum as a feedback control system. In Bellman's article, a diagram of Huygens's device is printed on the same page as a sketch of Watt's centrifugal governor; indeed, the diagram of the Huygens pendulum is about twice as big as the diagram of the centrifugal governor.³²

An interesting drama is unfolding here. Let's review. Mayr begins with the American Institute of Electrical Engineers' definition of "feedback control system" and observes that it is too weak. It fails to rule out certain types of mechanisms which most practitioners would not consider feedback control systems. Mayr adds a third condition designed to strengthen the definition so that it rules out these devices. In applying the new definition, Mayr sees that some devices which practitioners consider examples of feedback control systems are ruled out. Instead of considering his amended definition now to be too strong, Mayr rules that the earlier identifications were erroneous. There can be no pretense that Mayr's criteria constitute a rule of application which feedback control specialists routinely employ in recognizing feedback control systems. Not only was it not generally available to them before 1970, even today, there is not complete agreement with Mayr on what is or is not a feedback control system.

Kuhn, perhaps with an excess of optimism, seems to think that, with more time and patience, Mayr could produce a rule of application which accounts for all of the identifications practitioners in the field of feedback control have made, but there would be no guarantee

that the rule would correctly describe future classifications:

If the philosopher wants an adequate body of correspondence rules, he will have to supply most of them for himself.

Almost surely that is a job he can do. Examining the collected examples of past community practice, the philosopher may reasonably expect to construct a set of correspondence rules adequate, in conjunction with known symbolic generalizations, to account for them all. Very likely he would be able to construct several alternate sets. Nevertheless, he ought to be extraordinarily wary about describing any one of them as a reconstruction of the rules held by the community under study. Though each of his sets of rules would be equivalent with respect to the community's past practice, they need not be equivalent when applied to the very next problem faced by the discipline. In that sense they would be reconstructions of slightly different theories, none of which need be the one held by the group. The philosopher might well, by behaving as a scientist, have improved the group's theory, but he would not, as a philosopher, have analyzed it.³³

What Kuhn says here would be the case if the items under consideration bore only a "family resemblance" to one another. This is the kind of case in which the agreement Kuhn notices not only was, but had to be, achieved by means of exemplary cases rather than by rules of application. The idea is Wittgenstein's:

I can think of no better expression to characterize these similarities than "family resemblances"; for the various resemblances between members of a family: build, features, colour of eyes, gait, temperament, etc. etc. overlap and criss-cross in the same way.³⁴

It would be extremely difficult, perhaps impossible, to specify adequate rules for the resemblances we notice among members of a family. Perhaps it might be done for the present members of a particular family, but we could probably not write such rules to cover all future resemblances as well.

It appears then that there are two general strategies for obtaining the ability to classify things correctly. One depends on studying exemplars and "catching on" to family resemblances; the other makes use of rigid criteria contained in rules of application. Kuhn believes that which of these strategies is appropriate is a function of the kind of thing being studied. Of the method which relies upon exemplars, he says:

. . . everything which is special about this method of processing stimuli depends upon the possibility of grouping data in clusters with empty space between them. In the absence of empty space, there is no alternative to the processing strategy that, designed for a world of all possible data, relies upon definitions and rules.³⁵

If Kuhn is right, then with respect to the data concerning feedback control systems, it might be possible to group the data in clusters with empty space between them or it might not. That is, the world might be so constructed that there really is a separate class of feedback control systems such that, even though one might not be able to state firm criteria for class inclusion, the members of the class could be grouped together by recognizing family resemblances to exemplary cases. Or the world might be so constructed that the feedback control systems shade off imperceptibly into something else with respect to every property which seems characteristic of feedback control systems, so that if we are to identify a class of feedback control systems, we must do it by fiat. Or there is yet a third alternative. There may be things which behave like feedback control systems in almost any situation, "bordered" by things which behave like feedback control systems in relatively fewer situations, "bordered" by things

which behave like feedback control systems in only the most *recherché* situations. In this case, having the capacity correctly to classify something as a feedback control system, i.e., having the concept, will depend on our ability to see resemblances between situations as well as systems.

Mayr has bet on the second of these possibilities, i.e., that lines of classification must be established by fiat. There are two reasons for suspecting that he is wrong. The first is that the field of automatic control engineering has matured and had tremendous success without the benefit of necessary and sufficient conditions for something's being a feedback control system. Throughout the field, there is striking agreement on what constitutes a feedback control system, even though the various systems so identified differ widely with respect to their physical appearance, purpose, and the language used to describe them. The situation has all the earmarks of one in which the ability to recognize the appropriate resemblances is passed along from generation to generation of practitioners by means of exemplars. The second piece of evidence against Mayr's assumption is that Mayr's conditions are too strong to let in even the cases which he (and everyone else) would agree are the central examples of feedback control systems. Recall that the first of Mayr's conditions is that in a true feedback control system, "the system maintains the controlled variable equal to the command signal in spite of external disturbances."³⁶ But no actual feedback control system does that, not even the classic Boulton-Watt centrifugal governor!

The problem is that, in an actual feedback control device;

friction, inertia of the sensing and controlling apparatus, vibration, and hundreds of other factors are at work. When a deviation from the commanded conditions is detected there is a slight time delay before the system reacts to correct the undesired condition. During this time, the controlled variable may change, perhaps back to the desired condition, but the control system makes a correction anyway. The correction may thus be too much, too little, or even the wrong type. Of course, control system engineers work to minimize the resulting oscillation, but they cannot eliminate it. In general, their policy is to design mechanisms whose corrections will tend to return the controlled variable back to a point just short of the desired state.³⁷ So, any system constructed according to this strategy is ruled out by Mayr's first condition. Indeed, since every automatic control system has some small amplitude oscillation or "dither," Mayr's conditions seem to indicate that there are no feedback control systems.

We might try to amend the strong wording of Mayr's first condition to avoid this result. Instead of insisting that feedback control systems maintain the controlled variable equal to the command signal in spite of external disturbances, we could say that a feedback control system is one that tends to maintain the equality or that it maintains them roughly equal. Of course, "tends" here would be used in the sense of "successfully tends." But, given the already noted conditions of oscillation for virtually all feedback systems, this means that tending to maintain the controlled variable is not a notion that is any clearer or more primitive than that of a feedback system. Recognizing success in control is also something one must learn.

Furthermore, even if we were to accept the proposed modification of Mayr's conditions, there are still difficulties. For example, Mayr seems inclined to consider the Papin safety valve to be a feedback control system.³⁹ The device is like the weighted safety valve on top of a modern pressure cooker; when the steam pressure exceeds the force exerted by the weight, steam escapes until the steam pressure in the vessel returns to a level commanded by the size of the weight. The device certainly satisfies our weakened first condition. It operates as a closed loop with negative feedback, and Mayr has even supplied a block diagram which shows this loop, so the second condition is satisfied.⁴⁰ The fourth condition is also satisfied, since the functions of comparison and control are performed by other than human operators. That leaves the third condition. Does the Papin safety valve include a sensing element and a comparator, at least one of which can be distinguished as a physically separate element? The answer depends on what is meant by "physically separate."

If the requirement is that there be some item which can be identified as the sensing element or some item which can be identified as the comparator, then the Papin safety valve satisfies the final condition. The weight-loaded valve plug fulfills both of these functions. But if the requirement is that there be some separate item which serves to function solely as the comparator or solely as the sensing element, then Papin's safety valve does not satisfy condition three.

The third condition was introduced by Mayr explicitly in order to rule out "analog computer programs for differential equations, or simple physical systems with self-regulation, such as the water level

upstream of a weir, the R-C circuit, or the weather vane."⁴¹ But, if we consider the action of a weir in maintaining a chosen level of water in a mill pond, it is difficult to see why it should not also be classified as a feedback control system, along with the Papin safety valve. A weir is an adjustable dam. When the depth of water in a mill pond exceeds the depth for which the weir is set, it is permitted to spill over into a runoff channel. As the steam in a pressure cooker presses against the weight-loaded valve plug, the water in the mill pond presses against the weir. When either exceeds a preselected level, it is permitted to spill over. In both cases, there is a single identifiable item which serves as comparator and sensing element, independently of the rest of the system.

It appears that, if either is a feedback control system, so should the other be, but Mayr classifies weirs as not being feedback control systems, despite their having feedback loops, while he considers the Papin safety valve to be a (perhaps borderline) case of a feedback control system. Thus, the conditions given by Mayr are not sufficient to make the discriminations which he intends.

Of course, my remarks do not constitute a fatal objection to Mayr's analysis. A number of moves are open to him. He can rule that his original classifications were wrong or that I have misunderstood him. That is, he could claim that the weir constitutes a feedback control system (it was only the water upstream from the weir which condition three was designed to rule out), thus accepting a weak interpretation of condition three. Alternatively, he could adopt a strong interpretation of condition three, thereby ruling that Papin's safety valve

is not a feedback control system. Finally, he could attempt to devise some formulation of condition three which would permit him to retain the original classifications.

None of these methods of dealing with the problem would affect my main point: in making the classifications of feedback control systems and non-feedback control systems which appear in The Origins of Feedback Control and Feedback Mechanisms, Otto Mayr, a mechanical engineer and an authority on early feedback control devices, did not make conscious use of a set of necessary and sufficient conditions for the identification of feedback control systems. Unless one already knows how to discriminate feedback control systems reliably, one will be unable to get Mayr's results from Mayr's conditions.

NOTES TO CHAPTER ONE

¹ A defense of this view of concepts may be found in David M. Armstrong, Belief, Truth and Knowledge (Cambridge: Cambridge University Press, 1973), pp. 56-76. Of course, it would sound very queer to say that if we punch a lot of one-inch holes in the bottom of a bucket, we thereby give the bucket the concept of things which are less than one inch in diameter. But we do give the bucket the capacity to discriminate between things less than an inch in diameter and those things which are larger. It is in order to get around this sort of counterexample to the sufficiency of his definition that Armstrong explains that a concept is actually a "second-order" capacity (p. 61). The bucket with the holes in the bottom has a first-order capacity: when a solid object A with diameter of less than one inch acts upon it, the bucket must allow it to fall out. But a human being who has the concept of having a diameter smaller than one inch may or may not react towards A as a thing with a diameter smaller than one inch. The person may not desire to undertake any action at all towards items whose diameters are less than one inch. but, by having the concept, he has the (second-order) capacity to do so if he chooses. Buckets may have concepts, but only in those possible worlds in which they may make choices. In this world, they are mere objects, and only first-order capacities are appropriate to them.

² Thomas S. Kuhn, The Essential Tension: Selected Studies in Scientific Tradition and Change (Chicago: The University of Chicago Press, 1977), p. 297.

³ Ibid., p. 301.

⁴ Ibid., p. 302, n. 11.

⁵ Thomas S. Kuhn, The Structure of Scientific Revolutions, 2d ed. enl. (Chicago: The University of Chicago Press, 1970), p. 44.

⁶ Kuhn, Essential Tension, p. 305.

⁷ Ibid.

⁸ Ludwig Wittgenstein, Philosophical Investigations, 3d ed. trans. G. E. M. Anscombe (New York: The Macmillan Company, 1958), p. 36, remark 78.

⁹ Kuhn, Essential Tension, p. 305.

¹⁰ Ibid., p. 306.

¹¹ Kuhn, Scientific Revolutions, p. 200.

¹² Encyclopaedia Britannica, 1979 ed., s.v. "Automation," by Morris Tanenbaum.

¹³ Norbert Wiener, Cybernetics, or Control and Communication in the Animal and the Machine, 2d ed. (Cambridge: The M.I.T. Press, 1961).

¹⁴ Kuhn, Scientific Revolutions, pp. 188-89.

¹⁵ See, for example, Gordon S. Brown and Donald P. Campbell, Principles of Servomechanisms: Dynamics and Synthesis of Closed-Loop Control Systems (New York: John Wiley & Sons, Inc., 1948), pp. 367-87; or John J. D'Azzo and Constantine H. Houpis, Linear Control System Analysis and Design (New York: McGraw-Hill, Inc., 1975), pp. 568-612.

¹⁶ A.I.E.E. Committee Report, "Proposed Symbols and Terms for Feedback Control Systems," Electrical Engineering 70 (1951):909.

¹⁷ Otto Mayr, The Origins of Feedback Control (Munich: R. Oldenbourg Verlag, 1969), English ed. (Cambridge: The M.I.T. Press, 1970), p. 7.

¹⁸ Ibid., p. 8. On the previous page, Mayr explains that a "closed loop" is the closed causal chain which is formed when the results of a system's performance are transmitted from the output side of the system back to the input side. Such an arrangement has "negative feedback" when the sign of the fed back signal is negated so that control activity is based on the difference between the input and output signals.

¹⁹ Ibid.

²⁰ Ibid.

²¹ Ibid.

²² Ibid., p. 7.

²³ Ibid., p. 111.

²⁴ Kuhn, Essential Tension, p. 305.

²⁵ A.I.E.E. Committee Report, "Proposed Symbols," p. 909.

²⁶ Mayr, Feedback Control, p. 8.

²⁷ Otto Mayr, Feedback Mechanisms in the Historical Collections of the National Museum of History and Technology (Washington, D.C.: Smithsonian Institution Press, 1971), p. 90.

²⁸ Mayr, Feedback Control, p. 49-50.

²⁹ Ibid., pp. 50-51.

- 30 Ibid., pp. 95-99.
- 31 Ibid., p. 102.
- 32 Richard Bellman, "Control Theory," Scientific American (September 1964):189.
- 33 Kuhn, Essential Tension, p. 303.
- 34 Wittgenstein, Investigations, p. 32, remark 67.
- 35 Kuhn, Essential Tension, p. 312, n. 20.
- 36 Mayr, Feedback Control, p. 8.
- 37 Encyclopaedia Britannica, 1979 ed., s.v. "Control Systems,"
by Theodore J. Williams.
- 38 Mayr, Feedback Control, p. 8.
- 39 Ibid., pp. 82-3; and Mayr, Feedback Mechanisms, p. 70.
- 40 Mayr, Feedback Control, p. 84.
- 41 Ibid., p. 8.

CHAPTER II

THE ROLE OF EXPLANATION IN CLASSIFICATION

In the previous chapter, I tried to show that the notion of a feedback control system may be considered a Kuhnian paradigm in at least one of the senses in which Kuhn uses the term, "paradigm." It is a concept which is learned through the use of community-approved exemplars (which Kuhn also refers to as paradigms). It is also somehow fixed or delimited by the use of the same exemplars. In this chapter, I begin to sketch a descriptive theory of how a concept is delimited with the aid of exemplars.

Putnam's Observations about Natural Kind Terms

In an article entitled "The Meaning of 'Meaning'," Hilary Putnam proposes a theory of the meaning of terms which we use to name such naturally occurring items as gold, water, and tigers.¹ Putnam's discussion seems relevant here because he appears to believe that our concepts of gold, water, and tigers are delimited by exemplars, which he calls "stereotypes." Thus, I take Putnam to be supplying some suggestions about how we might flesh out Kuhn's remarks about paradigms. As we shall see, Putnam's suggestions are themselves in need of elaboration--at least when they are put to work as elucidations of Kuhn's program.

I shall begin with a brief summary of the claims which Putnam makes about natural kind terms. First, borrowing one of Kripke's terms, Putnam observes that natural kind terms are rigid designators: they refer to the same thing in every possible world in which they designate.² Just as "George Washington" designates George Washington independently of the truth of any particular description of George Washington which we might be able to supply, natural kind terms also designate the same things, even when various accidental features of the items designated are changed. One way of saying this is to say that "George Washington" and "water" designate whatever items have the essential properties of George Washington or water, whether or not we are able to identify those essential properties. Another way of making the same point is by using the metaphor of possible worlds. If we imagine the set of all possible ways the universe might be, including the way it actually is, then whatever "water" designates in one such world will be the same thing which "water" designates in any other possible world in which the term designates, if "water" actually is, as Putnam claims, a rigid designator. As Putnam notes, the notion of sameness here is a theoretical relation.³ Our identification of two items as being the same is always tentative, subject to defeat by further evidence or the overthrow of our theory of what it is for them to be the same. We shall return to this point later.

Putnam makes the further point that, when we decide to designate something by means of a natural kind term ("water"), what we have decided is that the item is the same (in theory) as what we call by that name in our actual world (water). Putnam's way of expressing this

second observation about natural kind terms is to say that they are indexical.⁴ They are indexed to the natural kinds of things we find in the actual world in which we live. Putnam further states that the doctrine that natural kind terms are rigid designators and the doctrine that they are indexical "are but two ways of making the same point."⁵ If by this he means that the set of all rigid designators is identical to the set of all indexicals, surely he is wrong. Instead of indexically tying our natural kind terms to the actual world, we could base them on paradigms from some other possible world and still employ them as rigid designators, so rigid designators do not have to be indexical. Further, it is possible to construct a system of terms whose meaning is initially fixed indexically but which systematically change their meanings from one alternative world to the next. It is hard to see what practical value that sort of a language would have, but the point is that rigid designators are not really identical to indexicals.

Putnam's comments, of course, concern our use of referring terms in language, but the ideas carry over nicely into our discussion of concepts. Our concepts are "indexed" to the world of our experience. Our capacities to react selectively toward different things are developed in our interactions with the furniture of the actual world. We learn actual world concepts, not some others.

These first two points are of relatively minor importance for the purposes of this dissertation. The remaining items, however, are very important for our purposes.

Putnam's third point about natural kind terms is that they may

have their meanings fixed socially, i.e., that members of a linguistic society may be competent users of a term even though they do not know all the criteria for its application but defer to the opinion of experts in the society for resolution of the difficult cases.⁶ Most of us are unable to tell a biscuit from a bun, a muffin, a roll, a scone, or a brioche. Still, "biscuit" has these differences as a part of its social meaning so long as we have some people who understand them and who can make the appropriate discriminations. In every language, Putnam thinks, there is such a "division of linguistic labor."

We must distinguish the capacities for classification which we can exercise as individuals in isolation from the rest of society from those to which we have access as members of our linguistic community. Most of us would be unable to classify a piece of metal as titanium, or at least we would not do it with a great deal of confidence, but we surely have the concept of titanium. Our capacity to react selectively toward samples of titanium depends on the cooperation and assistance of other members of our linguistic community--those skilled in metallurgy.

The fourth point is that items which we classify by natural kind terms usually qualify for those terms in virtue of their "hidden structure," essential characteristics not visible to the naked eye. Alternatively, if the item under consideration is not classifiable (even theoretically) by its hidden structure--perhaps because it has too many hidden structures--then its classification is determined by its superficial characteristics. Thus, an item's classification is rarely determined by its having some cluster of appropriate superficial

characteristics, but neither is it always necessarily determined by its hidden structure.⁷

Putnam's fifth point is that the observations he has made about natural kind terms seem to "apply to the great majority of all nouns, and to other parts of speech as well."⁸ In borderline cases, our linguistic judgments concerning the artifact term "pencil," to use Putnam's example, would probably reveal that we use it rigidly to designate whatever has the same hidden structure as the pencils with which we are familiar. Alternatively, there are a few words which we use as "one-criterion" words.⁹ These words are tied to some particular feature of an item; if it has the characteristic, then we use the word, otherwise not. Putnam notes that even words introduced as one-criterion terms tend after a while to develop a rigid, indexical natural kind sense.¹⁰ For artifact terms, Putnam thinks, this natural kind sense is primary. Thus, if we agree that the light bulbs found in American homes and offices might really be Russian eavesdropping devices, we reveal that "light bulb" and "incandescent electrical device with a glass housing" are not really synonymous for us. Although "light bulb" was almost certainly introduced as a one-criterion term, it has developed a more complex meaning for us. In the light of this, we may feel confident that Putnam would endorse our extension of his observations about our use of natural kind terms to "feedback control system" as it is applied to artifacts as well as to certain naturally occurring systems.

The sixth and final point about natural kind terms which Putnam

wants to make concerns the manner in which we learn them. We learn these terms, he claims, by having pointed out to us the important features of stereotypic examples of the term's application.¹¹ For example, the stereotypic idea of a car includes the notion that a car has wheels that roll on the ground or road surface. Putnam claims that there is a sense in which, if one acquires "car," he or she is obligated to acquire the information that stereotypical cars have wheels that roll on the ground or road surface.¹² Putnam leaves this notion of linguistic obligation unexplicated, but he admits that stereotypes do not have to be perfectly accurate. It might still be appropriate to label a vehicle which rides on a cushion of air a car, but "we could hardly communicate successfully if most of our stereotypes weren't pretty accurate as far as they go."¹³

This sixth point is an observation about concept standardization. We learn the same concepts as those shared by the rest of our linguistic community by sharing their exemplar set. If we examine any toddler's picture book, we shall find pictures of plump red apples, dogs with four legs, and green pine trees; even though withered green apples would still be apples, three-legged dogs would still be dogs, and brown blighted pine trees would still be pine trees.

This much we have already found in Kuhn. What is new here is the notion that, when we study such exemplars, what we derive from them is a "stereotype," an idea of the characteristics of a normal member of the kind of thing exemplified. We learn about apples, dogs, and pine trees by looking at what our society considers to be central cases of them and noting the features of those central cases. If

Putnam is correct, a key element in our progress from observing a set of exemplars to having the capacity to classify things in a certain way is a stereotype, or, I shall say, a model. Such a model is not a complete mental representation of anything in the sense that the truth value of every predicate which might be found to apply to the exemplars is determined in the model. Instead, an exemplar suggests to us the thesis that there is a useful concept to be had, consisting of the capacity to make discriminations according to whether or not something is relevantly similar to the exemplar. The question of which properties of the exemplar are relevant must be answered by a theory of what it is to be the kind of thing exemplified. Such a theory, then, supplies us with an exaggerated portrait, a caricature, of the exemplar. It provides a model of the exemplar or exemplars which is idealized in the sense that it is incomplete, being determinate only with respect to the properties which our theory identifies as important.

Although Putnam's remarks are made in the context of a theory of concept standardization, I think that the idea of an idealized model of exemplars plays a central role in the corresponding account of concept optimization as well. We form our classificatory theories in the faith that there is some theory, although we may not have guessed it, which will result in a maximally useful version of the exemplified concept.

In summary, Putnam's observations about natural kind terms suggest the following hypotheses about concepts:

1. They are capacities for classifying which we can employ
rigidly

2. They are indexed to the actual world
3. They are both individual and social capacities
4. They are supposed to reflect the hidden structures of the things we classify
5. Our concepts of artifacts share their characteristics with our concepts of natural kinds of things
6. They are acquired by learning the important features of what our society takes to be their correct application

It seems unlikely that Putnam would claim that the items listed above are either necessary or sufficient conditions for a capacity's being a concept. Instead, he might claim that they are the important features of the stereotypical application of the term "concept."

Unpacking Putnam's Notion of Hidden Structures

I want now to try to see some of the consequences of Putnam's fourth observation in which he suggests that our concept of items as being of one or another natural kind amounts generally to our capacity to classify them according to their "hidden structure" as opposed to their superficial characteristics.¹⁴ The virtue of this method of classification is that, if we succeed, we can presume to have accomplished the classification of things as they actually are, rather than as they merely appear to be. Let us begin to develop this theory of concept optimization.

Suppose that we decide tentatively to classify some item A as being of kind N on the basis of certain of A's observable properties, x, y, and z. I think Putnam would say that it is our conjecture that

A's hidden structure is the same as the hidden structure of the other members of N. This amounts to saying that we are making the following sort of hypothesis:

H: There is some item C which causes the members of N to have certain of their observable properties, and C also causes A to have x, y, and z.

At this point, it must be left ambiguous exactly which of the various notions of causation correctly describes the sort of cause which C is. It may be that the cause of A's having the particular properties that it does and of the other members of N's having the particular observable properties that they do is that they have the same form, are made of the same material, or are acted upon by the same agent. I suspect that each of these types of causation is appealed to at some time or other in our classifications. The important thing is that our tentative judgment that A is a member of N rests on the conjecture that the two sets of phenomena have the same cause.

Let us assume that to believe of some particular item A that it should be considered a member of a particular kind N just is to believe the appropriate corresponding instance of H. If this is true, then we can discover some of the things which are relevant to our belief that some particular A is an N, by seeing what would falsify a particular instance of H and what would not. We could be wrong in believing that A is of kind N in a number of ways:

1. C might not exist: there simply might be no common cause of the important properties of the members of N

2. C might not be a single item: there might, for example, be a common cause for some of the important observable properties of the members of N and some other cause for the others, or there might be a common cause for the important observable properties of some of the members of N and another cause for the properties of the remaining members, i.e., there might really be two (or more) kinds where we thought there was only one
3. The cause of the important observable properties of A might not be C, the cause of the important observable properties of the members of N

The "important" observable properties are just those which we suspect of having a cause in common with properties of other members of N. For Putnam, these important observable properties are the features of stereotypical assignments of "N." For Kuhn, they are the features of the exemplars which he sometimes calls paradigms.

If the belief that some particular A is an N is equivalent to the corresponding belief that H, then a consideration of what it would take to falsify an instance of H seems to show that items 1 through 3 above describe the only ways in which the belief that A is an N could be false. This has a couple of rather surprising results. First, we would not necessarily have to reject a given instance of H (or a classification of a given A as of a certain kind N) even if it turned out that all of the observable properties of A were different from any of the observable properties of the other members of N, so long as we judged it reasonable to suppose that certain of the observable

properties of A and certain of the observable properties of the other members of N (even though they might be different properties) had a common cause. Second, we would not always have to reject a particular instance of H if we had some theory about what sort of thing C is, and that theory turned out to be wrong. So long as the important observable properties of A and the members of N have some common cause, we are not wrong in saying that A is an N, although our justification for believing that any A is an N is stronger if we know what that cause is.

Given our understanding about the different ways in which our belief that some A is an N could go wrong and the range of situations in which the belief might still be true, what can we say about the minimum level of information which we must have in order to have a particular concept? Suppose that there is a particular capacity for classifying things which the members of a given linguistic community use or might use. What is the minimum information which a person would need in order to have that concept?

If we return to H and examine it carefully,

H: There is some item C which causes the members of N to have certain of their observable properties, and C also causes A to have x, y, and z.

we can discover three items which might be relevant to our having the concept of N:

1. Some known members of N, i.e., an exemplar set for N
2. The important observable properties of the known members of N

3. The common cause C of the important observable properties of the members of N

It seems to me that, when we are unable to use a particular classification schema, it must be because we lack specific information about one or more of the items named in 1 through 3 above. Let us pursue this suggestion. Suppose we let R stand for item 1 above, the set of known Ns; let P stand for item 2, the important properties; and C stand for item 3, the common cause; then we can name seven possible situations, depending on which of R, P, or C we know: (1) R-P-C, (2) R-P, (3) P-C, (4) C-R, (5) R, (6) C, and (7) P. To simplify the discussion, I shall ignore cases involving partially correct beliefs about P, C, or R. In which of these situations can the information which we know yield a usable concept? I think we must say that, if we know R, P, and C for a particular kind, N, then we stand on the firmest conceptual ground available to us. The interesting question is: How much of this information can we lose and still classify things under a particular concept?

In the second kind of situation, in which we know R and P, but we don't know what C is, I think we can still classify things as Ns or non-Ns, i.e., we can have the concept of an N. There were several periods in the history of the concept of an acid in which we could produce samples of acids and name their important properties, but we had no theory about the underlying cause of those properties.¹⁵ During those periods, we were clearly able to classify things as being or not being of a particular natural kind, acid. The third kind of situation, in which we know P and C, but we don't know any particular items which are Ns,

again seems to be one in which we are in possession of a usable concept. We seem to be in this situation today with respect to black holes. We know what the important properties of black holes would be, and we have some ideas about what would cause those properties, but we are not able to point to even one item which we know to be a black hole. Still, it seems possible to classify things as black holes or not on the basis of F and C alone, and the classification is of some value.

What about the fourth case in which we know some examples of specific items properly classified as Ns, and we have a theory about the underlying cause of the important properties of the members of R, but we have no idea what the important properties are? Could these clues furnish the grounds for a usable concept? The answer is a qualified yes. Suppose that a gifted biochemist, Hans Root, dies and leaves his notes, specimens, and equipment to us. In reviewing his notes, we learn that he spend the final years of his life experimenting with Q-radiation. Every biological specimen that he owned was labeled "affected by Q-radiation" or "not affected by Q-radiation." We might conclude that all the members of the first set are of a particular kind, even though we don't know what the particular effects of Q-radiation on these items were. Let us call this new kind the "quaff." We can produce the stereotypic examples of quaffs, and we know that Q-radiation is the cause of the observable properties relevant to something's being a quaff, but we don't know what the relevant properties are. I think that we could classify things as quaffs or nonquaffs, depending on whether or not they exhibit some change or other when irradiated with Q-radiation.

Just as we can't give any certain examples of black holes, but we know what a black hole would be like if we came across one, and just as in certain periods we couldn't explain why certain substances were acids, but we knew what properties needed to be accounted for by such an explanation; here we don't know what the important properties of quafts are, but we know what would count as such a property: some change which normally occurs when the quaft is irradiated with Q-rays.

Someone might object that our ability to classify things as quafts is simply an example of Putnam's third point which I mentioned earlier in this chapter, the "linguistic division of labor."¹⁶ The objector would claim that our use of "quaft" is parasitic on Dr. Root's having known what the important properties of quafts are. I think this may have been true while Dr. Root was alive, but now there is no one in our linguistic community who can discriminate the quaft-properties from any other properties, although trained biologists familiar with his work will be in a better position than the rest of us to guess what the important properties are. The expertise is gone, but still, I think, we have gotten hold of a usable concept, given only our knowledge of specific examples of quafts and the common cause of their quaft-properties.

Putnam is right about there being a linguistic division of labor, and when I say that "we" can classify things as being of a specific kind on the basis of "our" knowing certain things, I am talking about our linguistic community's achievements. I am not, for example, claiming that the average user of "gold" can state precisely the important properties of gold, explain why gold has these properties, or unerringly discriminate gold from other materials. But so long as there is someone

in our linguistic community who can do a sufficient number of these things, then we have a very useful concept. On the basis of such examples as acids, black holes, and quaffs, I conclude that knowing any two of the three relevant features (P, N, or C) of a particular concept permits us to classify according to that particular concept.

But suppose we knew only one of the features. Would we then have a usable concept? Suppose we knew only that, for a given kind, N, certain particular properties were N-properties, but we couldn't identify anything as having N-properties, and we didn't have any explanation of why something might have an N-property. Perhaps Dr. Root, in research notes clearly unrelated to Q-radiation research, had written a detailed description of the N-properties but had never explained what has these properties or why. We could classify things as Ns attributively (to borrow Keith Donnellan's term for a way of referring),¹⁷ I think, as "whatever kind of thing has N-properties," but only if all and only Ns have the N-properties. We could claim that any item A which has N-properties should be filed with whatever else has N-properties, but this would be merely an application of the concept of an N-property, not the concept of an N, unless we also knew that all and only Ns have the N-properties.

Or suppose all we knew about N was R, some paradigmatic Ns. Then we could classify things according to a manufactured relational property; Ns would be "whatever has the same cause of its important properties in common with whatever caused whatever properties these items have that were caused by the same thing." But this is really much too vague to permit us to classify some particular item A as an N.

Finally, suppose we know that some item C causes some things to have some of their properties, but we don't know which things nor which properties. Again, we could try to establish the class of "whatever has some properties caused by C." However, the assertion that A is an N would amount to the assertion that item A should be filed with whatever has some properties caused by C. It would only be reasonable to assert this if we knew that A had some properties caused by C, but, by hypothesis, we do not know this.

In conclusion, it does not seem possible to have a particular concept if we know only some items of the kind or the important properties of such items or the cause of those properties. If we don't already have the particular discriminatory capacity, we could never acquire it from a single feature. But when we know two of the features, we are in much better shape; much of the activity of normal science is devoted to the attempt to strengthen these two-feature concepts by discovering information about the third feature. Indeed, it may be the case that being in possession of two of the features puts us in possession of a concept only indirectly, by providing what we need in order to discover the missing feature.

In this section, I have made some suggestions about how to interpret Putnam's remarks about hidden structure. Two things have the same hidden structure and are thus the same kind of thing when, according to a theory which we hold, certain of their properties are caused by the same thing. The theory involved here is one which explains why things have the properties which they do. In later sections, I shall give up talk about having a common cause of certain observable properties.

Instead, I shall concentrate on the classificatory explanation which tells us which of an item's observable properties are important and why it should have the particular important properties which it has. I intend to make this change for two reasons. First, in this section, I have been intentionally vague about exactly what sort of cause is relevant to classificatory schemata. Shifting the emphasis to explanations seems to me to avoid this embarrassment. Second, whereas it seems to me that the notion of the exact sort of cause involved has to remain unspecified, I can say some quite specific things about the nature and parts of classificatory explanations. The result is a more informative and useful account of the processes of conceptualization and classification than would have been possible if I had kept attention focused on the vague notion of the hidden structure being some common cause.

Explanations and Models in Conceptualization

Let us now reconsider the idea which we developed earlier from Putnam's suggestions that an essential ingredient of a concept is a theory-based model of the exemplars for the concept. Let us begin by considering a formal aspect of our conceptual capacities. When we classify something, what is the nature of the criteria we use in making the classification?

We might suppose that all classes are invented by us and specified by a stipulative definition like that which defines the class of bachelors. We know that, formally, we can construct classes out of anything. We are interested, however, in developing those concepts which will be of greatest value to us in our efforts to understand and gain

some control over our surroundings. Although there are a few well-known classes whose contents are fixed by stipulation, the other classes of things which we do in fact discriminate are determined by some alternative method which permits us to improve our concepts as we gain experience. On the basis of this observation, we might distinguish two types of concepts: (1) artificial concepts (whose rules of classification are fixed by stipulation), and (2) natural concepts (which derive from exemplars). In the case of artificial concepts, the rule of classification has the form:

Anything which satisfies criteria $\underline{C}_1, \underline{C}_2, \underline{C}_3, \dots, \underline{C}_n$
is a member of this class.

The criteria might be disjunctive (\underline{C}_1 or \underline{C}_2 or \underline{C}_3 or \dots \underline{C}_n) or conjunctive (\underline{C}_1 and \underline{C}_2 and \underline{C}_3 and \dots \underline{C}_n) or some combination.

For natural concepts, the classificatory rule has the appearance of being as inflexible as the one stated above, but the situation is actually much more complex. It appears that the idea of a natural classification schema begins with a particular reference sample or exemplar set, e.g., the juice that comes out of a lemon, the fire left by a lightning strike, the lava from a volcano. The rule associated with a natural classification schema is then of the form:

Anything which is or is the same as the reference sample is a
member of this class.

There is, thus, an indexical aspect to natural classes, but, as we shall see, the sort of rule of classification characterized above is

incomplete; it needs to be supplemented by a theory of what it is to be the same as the reference sample.

In a strict formal sense, the only thing which is the same as the reference sample--whatever it is--is the reference sample itself. But as a classificatory scheme, this would generate utterly useless concepts, since everything would be the sole member of its own class. The power that comes from classifying things results from the fact that things in the same natural class behave similarly. If we learn a fact about one particular member of a class, then chances are we have discovered something which is true of every member of the class. Thus it is in our interests to work with large classes of things. Unfortunately, this also can be overdone. There are some properties which almost everything shares with almost everything else, e.g., being self-identical, being conceivable. What is needed is a scheme for constructing natural classes which are not rendered useless for purposes of learning about the world by consisting of only one thing or of almost everything. It is concerning this point that the rule needs to be augmented by a theory of what it is to be the "same" as the reference sample.

The sort of theory which I think is needed has two parts. The first part states which of the reference sample's observable properties are the properties which must be shared by any other sample if it is to be considered the same as the reference sample. This is necessary if the classification is to be of any use to us. Even though we can imagine a theory of classification on the basis of supposed nonempirical properties, such a theory would be completely worthless to human beings. Thus, a part of the theory which supplements the classificatory rule

for membership in a natural class must tell us the empirical criteria for class membership. This part of the theory would have the form:

Anything which has (a sufficient number of) observable properties $P_1, P_2, P_3, \dots, P_n$ is the same as the reference sample.

The second part of such a theory would justify the first part by explaining why we should believe that the particular observable properties named in the first part should be just the ones which will identify for us a class of things which will behave in interesting respects like the reference sample. Perhaps the second part of the theory would be expressible in the form of a covering law explanation:

1. Statement of applicable laws of nature
2. Description of the antecedent conditions, e.g., that the members of the reference class have hidden structure Z

3. Prediction to the effect that anything which has (a sufficient number of) observable properties $P_1, P_2, P_3, \dots, P_n$ is the same as the reference sample in most other interesting aspects

I shall bypass the issues of the nature of lawlike generalizations and the methods of discovering and justifying them. Instead, I shall focus on some problems concerning the description of the antecedent conditions which determines which laws will be appealed to in deriving the prediction which justifies the first half of the theory. For

example, we almost never apply Newton's $\underline{F} = \underline{ma}$ directly. We must first decide what kind of an $\underline{F} = \underline{ma}$ -type situation we are analyzing and then apply the formulation of $\underline{F} = \underline{ma}$ which describes that particular sort of situation. Thus, if we decide that the case we are assessing is best thought of as a falling mass, then we apply the $\underline{mg} = m \frac{d^2 s}{dt^2}$ version of $\underline{F} = \underline{ma}$. On the other hand, if we decide that the case is best thought of as a pendulum, then we apply the formulation of $\underline{F} = \underline{ma}$ which has the form $\underline{mg} \sin \theta = -ml \frac{d^2 \theta}{dt^2}$. If problems arose in using these derivative laws, we would not have to abandon the law; we would simply assume that we do not know how to apply it in certain kinds of situations.

Unfortunately, as we saw in the preceding paragraph, the most important part of a statement of antecedent conditions in a classificatory explanation seems to be a statement about what sort of thing the reference sample is (a falling body, a pendulum). But if we know what it is, we can already classify it. So it appears that a prerequisite for being able to decide how to classify something is knowing how to classify it!

But the situation is not so bad as that. What is required is that we propose a hypothetical model of the sort of thing that the reference sample is. Such a model would be a hypothesis about what the parts of the reference sample items are and how they are related. The derivative laws would then be invoked to predict what would be the observable properties of something with parts of that sort, related in that way. The model would then be tested along with the proposed classification scheme. The hypothetical model would gain justificatory support if the proposed way of classifying things permitted us to make a lot of very useful generalizations concerning members of the proposed class.

There is an interesting problem here which results from the fact that the set of observable properties in question may be derivable from several different models. Does the justifying theory therefore have to show which models work and then show that (and why) there are no others? I am not sure how to respond to this. I think that, historically, it has not happened: the choice of models always has meant a choice between what sets of observable properties are taken to be the important ones. That is, a successful model can override previously-made decisions about the important empirical (Putnam's stereotypic) properties. Examples of this will be given below and in the next two chapters. But it is certainly possible that two different models could justify our thinking that the same sets of empirical properties were the important ones in determining what kind of thing an item is. In this sort of case, I think we would be forced to be instrumentalists about our model making. The model would be viewed solely as a useful fiction and not as a real picture of the item being modeled.

To summarize the points which I have made thus far in this section, natural classifications are made by following the rule which tells us to include in the class whatever is the same as the items in a particular reference sample that we have in mind. A theory is needed to tell us which properties of the reference sample are relevant (or necessary as opposed to accidental) for being the same kind of thing as the items in the reference sample. The goal is to focus our attention on a class which will permit generalizations of discoveries about its members. The set of empirical properties we select is justified by an explanation of why those properties are the relevant ones. The explanation is itself

based on a hypothetical model of what the parts of the reference sample items are and how those parts are related. (Please note that, although my remarks have concerned covering law explanations, the covering law account of explanation is not an essential part of what I have been saying.) The hypothetical model is condemned or justified by the usefulness or uselessness of the resulting method of classification.

An Illustration:
the Acid-Base Research Program

Perhaps a specific example from the history of science will help to illustrate the points I have made. In a dissertation entitled Patterns of Scientific Change: the Acid-Base Research Tradition, Robin Fleming traces the history of the acid-base research program. He suggests that, in the case of acids, what I have been calling the reference sample was originally a collection of juices and by-products of fermentation which tasted sour.¹⁸ Over the past three hundred years, the following properties were thought relevant to being the same as the reference sample:

1. sour taste (in dilute aqueous solution)
2. corrosive (as concentrated)
3. soluble in water
4. easily neutralized (followed by indicators)
5. take place in displacement reactions
6. conduct electricity
7. catalyze reactions¹⁹

It should be noted that, although there have been a number of explanatory theories which attempt to justify a certain empirical characterization of acids, no theory has claimed that all seven of the properties listed above were relevant.²⁰ One set of empirical properties or another is held to be the important set of properties, depending on what

the explanatory portion of a particular theory is like.

The Cartesian explanatory model supposed that acids were composed of particles with tiny points. Acid strength was taken to be a function of the configuration of the particles. Arrhenius proposed that acids were composed of molecules which gave off hydrogen ions. The greater the concentration of hydrogen ions in solution, the greater the strength of the acid. According to the Brønsted model, acid molecules are proton donors, while on the Lewis model, acid molecules are electron-pair acceptors.

It appears that a natural concept is not always, and perhaps not ever, static. If a conception of acids based on the Brønsted model is more scientifically useful than a conception based on the Cartesian model, so much the worse for the Cartesian model and the set of empirical properties which it holds to be the important properties of acids. The dynamic nature of natural classification schemes seems to extend even to the reference sample itself. For example, the Lewis theory has been extremely powerful in studies of general reaction mechanisms in chemistry, and it is thus quite valuable to the overall program of chemical research.²¹ Unfortunately, the Lewis theory would exclude many of the substances in the original reference sample, e.g., acetic acid, hydrochloric acid, and nitric acid. The reaction of chemists has been to propose ad hoc compatibilist accounts which attempt to hold together Lewis's model of acids with the old extension of the term "acid." As Fleming points out, these hybrid theories are uniformly self-contradictory.²² Lewis proposed his model of acids in order to account for the existence of certain nonprotonated acids which were left unexplained by Brønsted's

model. Lewis's project was to save all the Brønsted acids and add a couple of others. Instead, his theory eliminates from the class of acids the most important of the Brønsted acids. Thus, Fleming concludes, Lewis's theory "outran" the justification for it and became merely a stipulated definition of an acid--an "artificial class" in my terminology.²³ But it seems to me that Fleming is wrong about this, since the Lewis model is defeasible by any acid model whose system of classification would be more scientifically useful. The proper moral to draw is that, as a concept matures, even its original reference sample may be superseded.

NOTES TO CHAPTER TWO

- ¹ Hilary Putnam, Philosophical Papers, vol. 2: Mind, Language and Reality (Cambridge: Cambridge University Press, 1975), pp. 215-71.
- ² Ibid., p. 231.
- ³ Ibid., p. 225.
- ⁴ Ibid., pp. 233-34.
- ⁵ Ibid., p. 234.
- ⁶ Ibid., pp. 228-29.
- ⁷ Ibid., p. 241.
- ⁸ Ibid., p. 242.
- ⁹ Ibid., p. 244.
- ¹⁰ Ibid.
- ¹¹ Ibid., p. 250.
- ¹² Ibid., p. 251.
- ¹³ Ibid.
- ¹⁴ Ibid., p. 241.
- ¹⁵ Robin Sherwood Fleming, "Patterns of Scientific Change: The Acid-Base Research Tradition" (Ph.D. dissertation, University of Virginia, 1975), pp. 114-15.
- ¹⁶ Putnam, Papers, pp. 227-29.
- ¹⁷ Keith S. Donnellan, "Reference and Definite Descriptions," The Philosophical Review 75 (1966):281-304.

- 18 Fleming, "Patterns," p. 32.
- 19 Ibid., p. 33.
- 20 Ibid., p. 53.
- 21 Ibid., p. 69.
- 22 Ibid., pp. 64-67.
- 23 Ibid., p. 81.

CHAPTER III

THE DYNAMICS OF CONCEPTUALIZATION

We have seen that classificatory judgments are supported by a theory. The theory contains an explanation which justifies our selection of certain empirical properties as being the characteristic properties of members of the class. The classificatory explanation itself contains a hypothetical model of the items in some original reference class. We have seen, e.g., in the acid-base case, that a particular classificatory scheme may remain vital and useful, despite changes in the set of observable properties taken to be characteristic of the members of the reference class, changes in the hypothetical model, changes in the composition of the reference class itself, and even changes in the classificatory explanation. In this chapter, I use specific examples from the literature of automatic control to construct a descriptive theory of the conceptualization process which accounts for its dynamic nature.

The chapter has two subsections. In the first, I focus on the fact that the hypothetical model used in a particular classificatory explanation invariably differs significantly from any actually existing thing which is modeled. I examine the kinds of arguments which are given to justify the idealizations made by such models, and I construct a preliminary taxonomy of these arguments. In the second section, I use

the results of the first section in constructing a program which describes the process by which we decide how to classify a particular item. The program gives us an insight into the process of conceptualization itself.

The Justification of Idealizations in Control Theory

One element of a complex empirical concept is a set of empirical properties which are thought to be relevant to the correct classification of any item satisfying the concept. The selection of a particular set of properties is justified by appeal to an explanation which attempts to show that the particular properties are the predictable result of the workings of known laws, given an idealized model representing the item being classified. Thus, when we decide to classify something as a feedback control system, we justify that decision by pointing out that the item has the characteristic properties. Our belief that we are paying attention to the right properties is justified by appeal to our explanation of why those should be the important properties. The explanation, in turn, is justified by our beliefs that the laws which it employs actually hold and that the model it uses is really like the thing which we're attempting to understand--in this case, a feedback control system. In this section, I want to investigate the final stage of justification: the justification of the belief that we are basing our classification scheme on the correct model.

Three kinds of models are commonly used in studying feedback control systems: (1) differential equations and other mathematical relations, (2) block diagrams, and (3) signal flow graphy. In this section,

I shall concentrate on the first of these. A widely-used modern textbook, Feedback and Control Systems, by Di Stefano, Stubberud, and Williams, states:

Mathematical models, in the form of system equations, are employed when detailed relationships are required. Every control system may theoretically be characterized by mathematical equations. The solution of these equations represents the system's behavior.¹

Feedback control systems are those which have as one of their important properties a certain type of behavior, discoverable by solving their system equations. The problem is that for any any actual control system, there is not a unique set of equations which is its system of equations. The problem is similar to the translation problem in logic: for any particular argument which we might find in its wild (natural language) state, there is not a unique formal model which represents its logical form.

In deciding which mathematical representation is appropriate, one must already have some idea which properties are important. That is, in modeling a particular control system, one must decide whether to include terms for the amount of heat which it generates, how much noise it makes, how big it is, what it smells like, how much light it absorbs or reflects, and how it tastes. Historically, none of these items has been considered especially important in modeling feedback control systems; however, it would certainly be possible to construct feedback control systems which have any of these items as the reference input or the controlled variable. It happens that the first feedback control systems which were studied mathematically were mechanical speed regulators, and

they were studied in the context of Newtonian mechanics, according to which everything is to be modeled in terms of masses in motion, exerting forces on other masses in motion.

Obviously, there is a gap between the Newtonian description of a speed governor as a collection of point masses moving in a certain way and what we hear, feel, and smell or taste in the presence of the device. But the scientists and mathematicians who applied the Newtonian techniques apparently saw no reason to attempt to justify this gap. The commercial value of the type of speed governor they were studying was in its moving in a certain way and causing the device which it governed to move in a certain way. According to the Newtonian view, an item's initial state, its mass, and the forces exerted on it are the only factors causally relevant to its motion. Along with Newton's laws come a set of invariance principles or what Ron Harré has called "principles of indifference."² Change in location has, of itself, no physical effects. There are no physical consequences of existing at one time rather than another. Changing an item's color, smell, or taste does not, by itself, alter its motion. Although these principles are open to disconfirmation by empirical counterevidence, by the time the Newtonian program was applied to speed regulators, they were well entrenched. The principles had not been disconfirmed, and Newtonian mechanics had proven to be an approach of enormous practical worth.

But even if one restricts the property types of interest and has a science dealing with such properties, there is still the problem of applying this science to the situation at hand. How is the world to be usefully described from this point of view? Typically, there is not a

unique answer; any number of paradoxically false but "approximate" descriptions are the most "useful." Consider, for example, the assumption of linearity which was one of the important elements of control theory until the late 1950's and is still used today except in the analysis of extremely complex systems. Although the British Astronomer Royal, G. B. Airy had suggested it earlier, James Clerk Maxwell was the first to work out the consequences of this important assumption for governors.³ The purpose of Airy's and Maxwell's investigations was to discover how to construct governors which would be stable in operation. In some installations, governors failed to work as desired by their inventors. When a disturbance was introduced, instead of returning the governed device to the commanded speed, sometimes the governor would cause the speed to differ from the desired by an ever-increasing amount, until the limits of the mechanism were reached. That is, if the governed system dropped below the desired speed, the governor would, in certain installations, operate to slow it still further; if the system oversped, the governor would act to increase its speed even more. In other installations, the governor would react to a disturbance by going into oscillations of ever-increasing magnitude. If the initial disturbance caused an overspeed, the governor would overcorrect, then overcorrect for the original error, and so on until the limits of the system were reached. Maxwell hypothesized that the operation would be stable only if the solution of the equation describing the motion converges, i.e., if all the real roots and all the real parts of the imaginary roots of the characteristic equation were negative.⁴ Nonlinear equations were extremely difficult, sometimes impossible, to solve; linear equations were all

solvable, although Maxwell was at the time able to solve only quadratic, cubic, or biquadratic equations, and he wrote that he hoped mathematicians would give the problem their attention.⁵ Nine years later, in 1877, Edward Routh published his Treatise on the Stability of a Given State of Motion in which a general technique is presented for discovering whether or not the solutions of linear equations converge.⁶

Maxwell and those who came after him modeled feedback control systems by linear constant-coefficient (or time-invariant) ordinary differential equations. A linear constant-coefficient ordinary differential equation is an algebraic equality which contains--

1. one or more dependent variables, one independent variable, and one or more derivatives of the dependent variables with respect to the independent variable;
2. no term in which the independent variable, time t , or any power of t is a factor;
3. only terms which are first degree in the dependent variables and their derivatives.⁷

The general form of such an equation is:

$$A \frac{d^n s}{dt^n} + B \frac{d^{n-1} s}{dt^{n-1}} + \dots + C \frac{ds}{dt} + D = 0$$

In Maxwell's article on governors, the assumption of linearity is warranted by limiting the discussion to infinitesimally small deviations. He says, without further discussion or justification, "we may confine ourselves to the case of small disturbances."⁸ Thus, Maxwell is using the technique of "linearizing at a point." Any curve may be described to any specified degree of accuracy as a concatenation of straight

lines. To make the model curve fit the actual curve more closely, we need only reduce the length of the segments being considered. In the same way, a nonlinear disturbance is modeled as a sequence of linear disturbances.

Maxwell's mathematical model presumes not only that any disturbances are linear, but also that all aspects of the governor's motion are linear and time-invariant. In so far as he considers the effects of friction, Maxwell assumes that friction forces remain constant, and he further assumes that the other characteristics of the governor are fixed and do not change with time. But, as Di Stefano remarks, "In actuality, linear systems do not exist. All physical systems are nonlinear to some extent."⁹ What sort of justification can be given for making the assumption of linearity? Maxwell does not admit that such justification is necessary or possible. However as S. Bennett points out in A History of Control Engineering, 1800-1930, Maxwell was concerned with governors which were specially made to exacting standards for use in laboratory work.¹⁰ Compared to the governors found on steam engines of the day, Maxwell's laboratory governors were nearly frictionless.

In a 1922 landmark paper entitled "Directional Stability of Automatically Steered Bodies, " N. Minorsky uses a linear differential equation to describe an automatic steering system for a ship. There appear to be two ways to approach the sort of idealization involved here and in the case of the speed governors mentioned earlier: (1) linearization might be the result of the assumption that the mechanism under study behaves in an ideal way or is an ideal mechanism, or (2) it might result from limiting the investigation to a regime of its operation in which

the actual mechanism is presumed to behave in the ideal way. Both Maxwell and Minorsky opt for the more conservative strategy of limiting the study. I shall develop this distinction in greater detail below.

Actually, we can find Minorsky using both methodologies, although it is only the second one (item 2 above) which he attempts to justify. In addition to limiting his investigation in an effort to catch the system in a linear part of its operation, there are a number of idealizations which he makes about it, e.g., he assumes that the rolling and pitching motions of the vessel are insignificant, and he assumes that there is no delay in the operation of the control mechanism, and in the final part of the paper, where he takes the delay into account, he assumes that the delay is time-invariant. Minorsky's justificatory argument concerns his assumption that, for headings very close to the desired heading, the control action will be linear: doubling the rudder angle will double the effect which it has on the ship's heading.

The argument proceeds in two steps. First, he wants to show us that the sort of limited investigation he proposes is worth doing. Second, he gives evidence that the actual system really does behave linearly in the regime of operation he is investigating.

He justifies the limited scope of his investigation with the following argument. First, a general analysis of the system would not be possible. He says:

The problem of automatic steering of moving bodies cannot be handled mathematically for the case of unlimited angular motion since there is no analytical expression applicable to the various torques acting on a ship, in general.¹¹

Of course, this is not a reason to think that there could not be such an

expression, but given the fact that Minorsky does not have one available, he cannot accomplish an analysis for the general case, but must limit his study. The second part of this argument for limiting the scope of the investigation has to do with justifying limiting it to the particular operational regime which he has selected:

The stability of angular motion has been considered only for small deviations of the steered body from the desired direction which not only simplifies considerably the analytical solution of the problem but gives the only solution that is of practical interest, since only small deviations are possible, if the steering is to be accurate; if there is no stability for small deviations, it means that there is no stability in general since the ship will be continuously deviated from her course.¹²

So Minorsky's strategy is to show us that we must limit the study in any event, if we are to get any useful results, and that we might as well limit it in the way he suggests, since stability in the regime which he proposes to investigate is a necessary condition for stability in any regime. If a system is shown to be unstable in the area of operation infinitesimally close to the desired operating point, then the system will not serve its intended purpose, no matter how it operates at other operating points.

Unfortunately, Minorsky's claim that "if there is no stability for small deviations, it means that there is no stability in general" is not exactly correct. We can imagine a ship which is permitted to deviate randomly five degrees either side of the desired course, and we would still say that the ship has a functioning automatic steering system if the ship's heading always remains within five degrees of the desired course. Indeed, Minorsky himself recommends this sort of control system for certain applications:

The second kind of disturbances is important only for smaller craft, responding instantaneously to each individual impulse or wave (which may become of importance for larger ships in a heavy sea); this disturbance constitutes what is generally called yawing among the waves; in this case the periodicity of such a forced yawing is obviously the apparent periodicity of the waves, and the rudder can do nothing to prevent it; in practice, in such cases intentionally loose steering is generally admissible in order not to over-regulate the rudder too much, which would only uselessly decrease the speed; this is also the only possible method in the case of automatic steering.¹³

Nevertheless, it is not my aim in this section to criticize Minorsky's use of idealizations in constructing his mathematical model, but merely to try to understand their general form and the arguments he gives for them. So far, we have been looking at his reasons for restricting the investigation to the operation of the system at a heading very close to the desired heading. The final part of Minorsky's argument gives us a reason for thinking that the assumption of linearity at that point will fit the facts. He writes:

Analytically this results from the well known method of approximation by infinitesimal analysis, according to which a small arc of continuous curve may be replaced by its chord without committing an error greater than the second order of small quantities, which therefore may be neglected.¹⁴

Minorsky reports carrying out a series of experiments in which he determined that the eye cannot detect rates of yawing less than 3.5×10^{-3} radians/second or angles of yaw less than 4.3×10^{-3} radians. In his article, he assigns these as the first order quantities, the quantities determining the length of the arc to be linearized. He says that these quantities may be "fixed arbitrarily" at the values mentioned above.¹⁵ The choice is one which would guarantee that the system under study, if

stable, would produce results better than any human helmsman could produce. But the choice is arbitrary to the extent that assigning smaller segments of arc would produce increasingly accurate results for increasingly smaller operating regimes. The method of approximation by infinitesimal analysis always gives us the option of improving the fit between theory and reality to any given tolerance.

Summarizing to this point, we find in the Minorsky paper two types of idealizations: (1) the idealization that the various parts of the systems under investigation have time-invariant characteristics, and (2) the idealization that, within a specified range of admittedly non-linear operation, the system operates linearly. The first of these is never argued for, but a justification is offered for the second. The justification finally reduces to the claim that, although the idealization is actually false, it can be made as close to the truth as we desire--if we are willing to pay the price of limiting the scope of the investigation even further.

Notice how the idealizations are used. Minorsky takes Newton's laws (which he presumes to be true) and applies them to an ideal model (which he knows to be a false, but improveable, description of the system under study). The result is a range of predictions of the properties which an ideal control system would have. He then restricts our attention to a restricted subset of the set of predictions and claims that the actual control system has the properties in this subset. This is subjected to empirical test, and, if "successful," we are to describe the relevant features of the actual system in terms of the salient features of the model, even though the model is, strictly speaking, false

or ideal. Although the method of investigation employed by Minorsky appears to be instrumentalist (i.e., the model is taken to be of interest only for some of its predictions), the features of the model are taken to be a realistic description of the artifact being studied.

Let us now consider some other types of arguments offered by various engineers in the field of automatic control to justify the idealizations which they employ. In a 1934 paper entitled "Theory of Servo-Mechanisms," H. L. Hazen offers justificatory arguments for assumptions of the type which Maxwell and Minorsky made without justification.

The argument which Hazen uses more than any other is simply the claim that the idealizations which he makes do not cause any big gaps between theory and reality. Where Minorsky's assumptions of linearity was warranted by his limitation of his theory to a restricted regime of operation, Hazen attempts to justify his idealizations by restricting the type of artifact to which the theory is supposed to apply as well as the range of operation of those artifacts. The restriction comes in several forms:

This representation is very close to the truth for many servos. . . . test and calculation agreed within a small experimental error.¹⁶

Actually there is always some inactive period which may be so small, however, in many instances that its effect is insignificant.¹⁷

Especially in high-speed servos the forces obeying the Coulomb friction law are likely to be quite negligible when compared with the forces having the effect of viscous friction.¹⁸

For many continuous-control servos these assumptions correctly represent the facts.¹⁹

Tests of the servo thus designed show very fast response and an excellent agreement between test and calculation.²⁰

Such arguments do not actually justify the assumptions in question; rather, they reduce the demand for such justification by showing that there is not really so much that stands in need of justification. The closer the ideal model comes to what we know of reality, the smaller the justificatory burden.

The second class of arguments which Hazen uses is related to the first. Instead of claiming that there isn't much that needs to be justified, arguments of the second type promise that there is a way to reduce the gap which needs to be justified. The following are examples:

In most specific cases in which a general analytical solution is too cumbersome to be useful, a restricted analytical or a numerical solution can readily be made, taking into account departures from more idealized conditions. The methods of treating these more involved cases will be outlined briefly at suitable points in the analysis.²¹

A comprehensive treatment of all cases is evidently beyond the scope of this paper, but a few significant cases will be analyzed and from the results obtained deductions can be made covering other cases.²²

This sort of argument has the same point as the one I discussed earlier in connection with the Minorsky paper: the gap between theory and reality is acceptable because we know how to reduce the gap if we should ever want to. If the results of our calculations concerning a general idealized model are not close enough to the facts in any given case, we know how to adjust the idealizations to make the model closer to the specific kind of case we are considering.

The first two types of justificatory argument given by Hazen treat the gap between theory and reality as a bad thing. One tries to

convince us that there's really not much of a gap, and the other claims that whatever gap there is is no problem, since we can reduce it, if we should need to. The third and final type of justificatory argument used by Hazen treats the gap as something desirable. There is no need to apologize for our idealizations or to attempt to justify them. The idealized models are valuable to us because they are not descriptions of any particular actual systems. Consider these examples of this type of argument, picked from Hazen's article:

Although the ideal servo is never realized in practice, its operation furnishes the standard by which the operation of actual servos is judged.²³

It is quite evident that such operation represents a limiting case which at best can only be roughly approached with actual physical apparatus. Practically the amplitude and frequency of oscillation are finite, and in most cases lag error is present. The limiting case is of interest however from the point of view of analysis and from its significance as an ideal.²⁴

In other cases, the idealizations made in order to make a usefully simple analytical treatment possible may depart somewhat more from the facts. Nevertheless, the analysis of an idealized case gives a real insight into the characteristics of a given servo.²⁵

Physically this result is absurd for any actual servo, hence the necessary conclusion is that this case is too greatly idealized to represent the facts. It is of interest, however, as a limiting case.²⁶

Although as a first condition, this is an idealized case that cannot be realized physically, the result has interest and significance. It shows that an ideal relay-type servo is capable of following a constant input velocity with precision, i.e., without finite deviation in the nature of lag or oscillation. That this condition could be approached even under ideal conditions is interesting.²⁷

Although this condition is of interest as a limiting case, the effect of time lag is always present practically.²⁸

Limiting cases give us a glimpse of what the systems being modeled would be like if they were changed in certain ways (changes that would make them more like the idealized model). If we know what properties a frictionless control system has, then if we reduce the effect of friction on an actual control system, we know, at least qualitatively, what sort of changes to expect. Our study of ideal limiting cases gives us knowledge about the actual cases. This argument may be related to the other two by viewing it as another type of "gap closing" argument; we know what it would be like to close the gap between the idealizations and reality by altering the actual system under study. Interestingly, in the context of a conceptualization scheme, this sort of argument seems to license us to describe the artifact being considered in terms of our idealization of it, despite the fact that predictions we derive from the idealization are not usefully accurate.

In summary, Hazen uses three types of justification for the idealizations which he makes:

1. The Negligible Discrepancy Argument: whatever dissimilarity obtains between the actual systems being studied and the idealizations of them results in no significant difference between prediction and experiment
2. The Possibility of Improvement Argument: we know how to complicate the idealized models in ways which will reduce the discrepancy between prediction and experiment to any specifiable limit
3. The Limiting Case Argument: the idealizations are important sources of knowledge about the systems we are studying, even

when there are significant differences between them and any actual systems we may be studying

Not all writers are as clear about the idealizations which they are using nor the arguments which are supposed to justify them as Hazen; nevertheless, in the landmark articles in the development of classic linear automatic control theory, each of the arguments given is reducible to one of the types of argument adumbrated above. In his famous 1932 article, "Regeneration Theory," H. Nyquist uses a version of the limiting case argument:

Now, this fact as to equality of gain and loss appears to be an accident connected with the non-linearity of the circuit and far from throwing light on the conditions for stability actually diverts attention from the essential facts. In the present discussion this difficulty will be avoided by the use of a strictly linear amplifier, which implies an amplifier of unlimited power carrying capacity. The attention will then be centered on whether an initial impulse dies out or results in a runaway condition. If a runaway condition takes place in such an amplifier, it follows that a non-linear amplifier having the same gain for small current and decreasing gain with increasing current will be unstable as well.²⁹

If an ideal amplifier has the property of instability, then an actual one which differs from the ideal one in being more susceptible to instability will be guaranteed to be unstable. By studying a limiting case, we can discover whether or not a particular sufficient condition is fulfilled.

In 1934, H. S. Black wrote in "Stabilized Feedback Amplifiers":

To determine the effect of feedback action upon modulation produced in the amplifier circuit, it is convenient to assume that the output of undistorted signal is made the same with and without feedback and that a comparison is then made of the difference in modulation with and without feedback.³⁰

Here Black seems to be using Mill's methods to study the effect of feedback in a set of limiting cases. I suggest that this use of such methods in strictly theoretical investigations is a widespread and quite useful practice.

In a 1934 paper on process control, "Theoretical Foundations of the Automatic Regulation of Temperature," A. Ivanoff uses the limiting case argument quite consciously:

In such cases a number of conclusions and numerical results will differ from the data given in the present paper. The treatment of the problem proposed is, however, regarded not as a universal and inflexible theory, but rather as a standard by which one can judge the quality of a plant from the point of view of exact regulation, and which one can use to compare the various methods and systems of control. Such being the purpose of the treatment, the author feels that no apologies need be made for a number of such assumptions as the one limiting the variation in the potential temperature to moderate amounts, which, though prejudicing somewhat the general nature of the calculations, simplify to a considerable extent the final application and evaluation of the results.³¹

Quite a few authors talk about making idealizations because doing so "simplifies considerably the analytical solution of the problem."³²

But I contend that simplification alone is not sufficient justification to make simplifying assumptions. To say that doing something makes your work a bit easier is not to say that you ought to do it. Virtually any idealization of an actual thing simplifies the mathematical model describing it and thus simplifies manipulations involving that model. In every case in which I have found a control engineer talking about simplifying assumptions, the simplifications are actually justified by one of the three types of argument given earlier. The simplification is a by-product.

A slightly different form of the limiting case argument is given by A. Callender, D. R. Hartree, and A. Porter in their 1936 paper, "Time-Lag in a Control System":

A systematic theoretical investigation of the factors involved in the process of control is interesting in itself, and is desirable, if not absolutely necessary, as a basis for a detailed design of control apparatus, as its results are much more precise and definite, as well as more general, than any which could be reached by a purely empirical investigation with an actual control system, and also because the field is too extensive to be covered adequately by such an empirical study.³³

The idea is that, since there are so many different types of control systems available, it is impossible to produce an exhaustive analytical study of each of them. A better strategy is to divide the field and study the two or three limiting cases on the basis of which the division was been made. The fact that Callender, Hartree, and Porter expect to achieve "more precise and definite" results from their investigation as a result of this strategy than would be possible from empirical investigation with actual control systems seems less convincing. Theoretical results worked out to a million decimal places are uninteresting, unless it can be shown that they have something to say about actual things.

Later in the paper, Callender, Hartree, and Porter give a version of the negligible discrepancy argument which owes much to the earlier papers of Maxwell and Minorsky:

The equations (1) and (2), or (6) and (7) are linear, as is desirable both for practical and for analytical reasons in order that the superposition principle should apply to their solutions, so that the effects of disturbances occurring at different times should be additive. In practice, the variations of θ are likely to be small, so that linear equations should provide an adequate representation of the behavior of a real system.³⁴

Finally, I wish to point out a couple of possibility of improvement arguments in H. W. Bode's 1940 paper, "Relations between Attenuation and Phase in Feedback Amplifier Design." One of Bode's projects in the paper is to demonstrate a shorthand method of determining the attenuation and phase characteristics of a circuit. The method consists of determining the asymptotes of the characteristic curves without actually making a detailed plot of the curves. Bode notes:

By proceeding sufficiently far in this way, an approximate computation of the phase characteristic associated with almost any attenuation characteristic can be made, without the labor of actually performing the integration in [equation] (2).³⁵

But this shorthand design technique is justified by the detailed equations which Bode has already produced earlier in the paper. If more accurate information is needed, we need only go back to the equations:

Departures from the other assumptions are easily treated. For example, if a varying feedback in the useful band is desired, as it may be in occasional amplifiers, an appropriate cut-off characteristic can be constructed by returning to the general formula (4), performing the integrations graphically, if necessary.³⁶

In this section, I have identified the kinds of justificatory arguments which were used in justifying the use of idealizations in classificatory explanations in several important papers in the development of feedback control theory. The arguments have noticed reduce to three types; in the next section, this taxonomy is expanded.

Extending a Concept

The concept of a feedback control system passed through a number of stages on its way to maturity. I think that the life history of the feedback control system concept is typical of the development of complex concepts, and I intend to test this suggestion in the next chapter. This section is about the way in which the notion of a feedback control system has swelled to include types of systems not obviously the same as the items in the original reference sample. An outline of the history of the concept follows:

STAGE I: PRECONCEPTUAL STAGE. Surrounded by natural feedback control systems (homeostatic mechanisms of the body, weather systems) which they do not recognize as embodiments of a general concept, isolated inventors construct feedback control systems (float-valve regulators, temperature regulators, pressure regulators) employing the feedback principle intuitively, without understanding it as a general concept.³⁷

STAGE II: REFERENCE SAMPLE ACQUISITION. Watt's centrifugal governor appears in 1788 and proves to be of great economic value. Dozens of variations on the Watt design are constructed. The hypothesis is made that these devices are representatives of a type or kind: the speed governors become the reference sample. Things of the same kind (other governors) are recognized by their structural properties (appearance, at this stage) and their performance properties (they keep prime movers from underspeeding or overspeeding).

STAGE III: DEVELOPMENT OF THE CLASSIFICATORY EXPLANATION. Attempts are made to explain why something with the structural properties of the reference sample governors should exhibit their performance properties, given currently accepted laws. This is one of the activities of what Kuhn calls "normal science."

To be useful, laws must be general; at least in principle, the regularity stated by a law must be a regularity which is characteristic of a class of items. The larger the class to which the law applies, the more work we can get out of it. The statement of a law designates the class of items about which it says something true by abstracting from the many properties of actual items to which the law applies certain properties thought to be indicative of the presence of the regularity expressed by the law. Thus, laws, in order to be as useful as possible, contain idealized descriptions of members of the classes to which they apply. An explanation which includes such a law must also employ idealizations, so the classificatory explanations concerning speed governors are not explanations of why some particular governor has the particular performance properties that it does. Instead, the explanations use the laws to explain why an item with certain idealized structural properties can be expected to have certain idealized performance properties. Justificatory arguments are used to show that the explanations are acceptable, despite the idealizations.

The cases involving mechanical feedback control systems make it tempting to suppose that there will always be a clearly discernable distinction between structural and performance properties, no matter what sort of thing is being classified. The structural properties account

for the material that an item is made of and the way in which that material is arranged, while the performance properties have to do with the way in which the item behaves or operates. In a classificatory explanation, the latter is explained as a result of the former and the operation of various laws of nature. But a quick look at another sort of case will show us that a division into structural and performance properties is not always possible.

Let us consider, for example, the theory dealing with amorphous semiconductors.³⁸ Amorphous semiconductors are devices which act like crystalline semiconductors but which are not constructed of crystalline materials. Most of the time they are not especially good conductors of electricity, but they become good conductors when thermal or electrical energy is added.

The model of solid conductors, insulators, and semiconductors is an extension of the model for individual molecules. Solids are like giant molecules. In perfectly crystalline solids, all the electrons are shared with all the atoms, and there is equal probability of finding a given electron anywhere in the material at a given time. In a giant molecule solid, the analogue of the energy shells found at the atomic level is the energy band. We can plot the density of these bands, i.e., the number of states per unit volume per unit of energy. Because the energy bands are built up around discrete energy states, there are gaps between them and sharp "band edges" marking the gaps. As the electrical properties of atoms result from the electrons in the outer shells of the atoms, the electrical properties of solids result from the position of the highest filled state in the energy bands.

The electrical properties of a solid depend (at least in part) on how much energy is required to make electrons available to be charge carriers. Imagine a line drawn halfway between the highest filled state and the lowest unfilled state in a solid. This line marks what is called the Fermi level. If the Fermi level is in the middle of a band, the solid is a metal--very little energy is needed to make electrons available for conduction. If it falls in a large gap, the solid is an insulator. If it falls in a small gap, the solid is a semiconductor. Doping, the process of adding impurities, turns crystalline solids which would normally be insulators into semiconductors by making available some extra energy states in the gap above the highest filled state, thus making the gap in which the Fermi level falls narrow, rather than broad.

This is the part of the conventional theory of solids which accounts for the availability of electrical charge carriers. The crystalline structure is one of the properties of the idealized model of solids. The further something deviates from being a crystal, the more difficult it is to make the theory work for a whole solid object, and the object must be considered a collection of segments each of which is a crystalline solid.

The problem with amorphous semiconductors is that the conventional theory says that they ought to have the electrical properties of metals, not semiconductors. Instead of well-defined bands with sharp band edges, amorphous materials have sloppy bands with tails which fill in the gaps. The gaps with sharp band edges found in a crystalline solid result from the periodic structure of the crystal, but this is absent in amorphous materials. In amorphous materials, there are some energy

states available at every energy level. No matter where the Fermi level falls, it seems that amorphous solids ought to behave like metals because the highest filled band always has an unfilled band immediately above it.

A solution to the problem of amorphous semiconductors was proposed in 1971 by Sir Neville Mott of Cambridge University.³⁹ Mott theorized that the states in the band tails are localized within the solid like the phosphorous or arsenic impurities in an ordinary crystalline semiconductor. Hence, the reason that amorphous semiconductors don't exhibit metal conductivity despite the fact that there are no sharp band edges with energy gaps in between is that the carriers which can occupy the states in the band tails can be carriers only in a small region of the solid. This explains why amorphous semiconductors are not good conductors under normal conditions, but it does not explain why they should become good conductors when a certain amount of energy is added.

The reason for this is that there is a particular density of states above which the localized states in the tails become extended throughout the solid. Hence, there are critical energies above which the mobility suddenly jumps, and these establish mobility edges analogous to the band edges of crystalline solids. The mobility edges create a mobility gap between bands; available carriers conduct electricity only if they are free to move.

Now the question I want to ask about all this is: in a given amorphous semiconductor, which of its properties are to be considered structural properties and which are its performance properties? It seems right to say that at least one performance property is its

property of electrical semiconductivity, and at least one structural property is the property of being composed of, e.g., selenium. But now consider the property of having its Fermi level fall in an energy gap with sloppy band edges or the property of having a mobility gap with sharp mobility edges. These seem to be performance properties which result from the structural properties of being composed of selenium or of having a noncrystalline arrangement of the selenium atoms. Unfortunately, they seem also to be structural properties which account for the performance property of electrical semiconductivity. Thus, we are confronted with properties which seem to be both structural properties and performance properties. This bodes ill for the usefulness of the distinction, since it is not difficult to find additional examples of structural-performance properties, once we see the problem. Consider Watt's first centrifugal governor. Earlier, we wanted to say that the performance properties of that governor--the range of stable operation, the rate of correction for deviations, and so forth--resulted from certain structural properties--the length of the arms, the mass of the fly weights, the type of steam valve to which it was connected. But any of the items I have mentioned as structural properties may be seen as performance properties. The length of the arms is what it is while the governor is in operation as a partial result of the resistance to bending of the material of which the arms are composed. The fly weights have the mass they do because of the density of the material from which they have been cut. The steam valve has the properties which it does because of the arrangement and type of the materials from which it is made.

What is evident from this short listing of properties which may be either structural or performance properties is that something is one or the other only in relation to some other property. Nothing seems to be a performance property or a structural property simpliciter--at least it seems quite easy to think of properties which are structural or performance in relation to something else, and I cannot think of one property which is a structural or performance property in vacuo. The mass of the fly weights is a structural property of a governor with respect to the operating characteristics of the governor, and it is a performance property with respect to the density of the brass from which they were cut.

What we have been calling a distinction between structure and performance seems actually to be a distinction between a property which is explained and a property which explains it. Relative to a certain explanation E, some properties Π (pi, for "performance") get explained, while certain other properties Σ (sigma, for "structural") are an essential part of E, the explanation of the Π -properties. The Σ -properties of E may become the Π -properties of some other explanation E', themselves explained in terms of some further set of properties, Σ' .

In a classificatory explanation, the Π -properties and Σ -properties we are dealing with are of two types: those belonging to an idealized model and those belonging to some actual thing we are trying to classify. The arguments which we uncovered earlier in this chapter are arguments which justify our using a model composed of a certain set of idealized Σ -properties and Π -properties to motivate a particular classificatory judgment about some actual item whose Σ -properties and

Π -properties are not identical to those of the model.

In the following, I shall continue to use the terms, "structural" and "performance," but we must keep in mind that these terms are descriptive of properties only relative to a given explanation. The performance properties are the explanandum and the structural properties are a part of the explanans. Now let us return to our discussion of the "life history" of the feedback control concept.

STAGE IV: CONCEPT EXTENSION. The concept is extended on the basis of the explanation(s) developed in stage III. The idealized properties become the stereotypic properties of the class--even though they are strictly false of almost everything, even the reference sample.

STAGE V: CONCEPT COMPLETION. At this stage, the concept is fully-defined. Necessary and sufficient conditions are available for the application of the concept term, although some borderline cases remain because of the difficulty of deciding whether or not the conditions are satisfied in particular situations. The definition is empirically defeasible, but only in a major conceptual revolution.

It is doubtful that stage V is ever finally reached for the most interesting and sophisticated concepts, but the concept of the wheel, the lever, and the inclined plane show us that stage V is possible for some complex concepts. Until stage V is reached, it appears that the processes described in stages III and IV are repeated continually. Just as explanation E_1 is completed and the process of concept extension on the basis of E_1 (a process to be detailed in this section) is begun, someone may be working on a new, improved explanation E_2 , and so on. We shall have to say that a concept is in stage III with respect to

explanation E_2 and stage IV with respect to explanation E_1 .

With regard to stage IV, the concept extension stage of concept development, there are two questions which need to be answered. First, how can it be reasonable to claim that what it is to be a particular type of entity N is to have a certain set of idealized structural and performance properties, when almost nothing in the reference sample for N has the idealized properties? And second, how can it be reasonable to swell the extension of the term " N ," when almost none of the items added to the extension have the idealized properties which we say are stereotypic of N 's? The answers to these questions are related to the kinds of justification which we found being given in the previous section.

Suppose we have a particular article, A , which we suppose, perhaps because of its appearance, to be an item of type N , e.g., a feedback control system. How do we decide whether or not it should be included in the extension of the concept term? Clearly, the answer is that we appeal to the classificatory explanation which is the product of stage III development. Call that explanation E . E consists of certain laws L and a model with properties P , according to which items of kind N have properties P . In classifying some article A , we submit A to the following test:

TEST 1: Does A have all the properties P ?

If the answer is yes, then surely A is an N (a feedback control system or whatever instantiates the concept we are considering). But, as we noted above, the answer is almost never yes. The properties in set P

often include the properties of being massless or frictionless or having coefficients which do not vary with time, and no actual thing has these properties. If the answer to test 1 is no, then must we conclude that A is not an item of the type we are investigating? Actually, no; we submit it to another test:

TEST 2: Is the gap between the actual properties of A
and the properties in the set P too great?

The idea here is that we are interested in the particular type of article because it seems useful. That is why stage II happened; the reference sample is just a collection of things that have the appearance of being the same and which has caught our attention because the things in the collection seem valuable in some way--because of their usefulness or perhaps (in the case of certain types of things) their beauty. If the properties of the particular item we are considering, A, fall very short of the properties which made the items in the reference sample seem valuable, then that seems a good criterion for rejecting it. We treat P as a specification of what the properties of the reference sample items would be if they had their valuable properties in the most ideal way. If A falls just a little short of that ideal, it will (for most purposes) be sufficiently valuable for us to count it as being of the same type as the items in the reference sample.

What I have said above is only a part of the story. In the case of feedback control systems, it is certain of the performance properties we find valuable, but since our classificatory explanation tells us it is the structural properties which are responsible for the valuable

performance properties, the structural properties (the presence of a feedback loop) seem to take precedence in our applications of test 2. I think this is generally true of complex concepts; it is what Putnam's "hidden structure" metaphor points to.

My comments in the previous paragraph may make it sound as though we can dispense with the performance properties entirely in deciding how to classify artifacts. The question of whether or not something is a feedback control system would reduce to the question of whether or not it contains a feedback loop. But the presence of a feedback loop is not a sufficient condition for something's being a feedback control system; we can find feedback loops in almost anything. For example, pick up a pencil by the eraser end and rotate it between your fingers. There is a sense in which the pencil can be considered a type of feedback control system known as a servomechanism. Your twisting motion is the control input, and molecular forces within the materials of which the pencil is composed react to any angular differences between the lead end and the eraser end and operate to eliminate that difference. Of course, the angular difference is never very great, but since the pencil has some mass and is not perfectly rigid, there will be some twisting. Nevertheless, there are very few occasions on which we would think it appropriate to claim that a pencil is (or contains) a feedback control system. A part of the reason for this seems to be that the performance properties of the pencil servomechanism appear on such a small scale that we do not notice them or (unless they are specifically brought to our attention) see them as being valuable in the sense that the performance properties of other servomechanisms are valuable.

If most pencils were completely lax and pliant, so that they behaved like a stack of well-oiled washers, we would find them rather difficult to write with, I think. If someone discovered a way of transmitting information about the angular position of the particular part of the pencil which you hold to the other parts of the pencil and using that information to make the other sections of the pencil follow the position of the manipulated part, then we might be more inclined to take the notion of pencil servos more seriously. Perhaps we need to understand what something would be like if it were uncontrolled, before we can think of it as being controlled--or before we can see that control as being valuable.

I am unable to give clear and complete criteria for applying test 2, but I have pointed out the kinds of considerations which motivate those criteria. It seems likely that the test is always performed using standards of which we are not fully conscious. In any event, if we decide that the gap between the actual properties of A and the properties in the set P is not too great, then we are able to classify A as a member of the class of items we are considering. If the gap is too great, then we must ask whether or not this situation can be corrected. Given that there is a shortfall between the properties which our classificatory explanation says feedback control systems have and the properties which item A actually has, can we see any way to lessen the shortfall and still retain the classificatory explanation? We might express this as:

TEST 3: Can the gap between the actual properties of A and the properties in the set P be reduced to some specifiable tolerance?

I have added "to some specifiable tolerance," as a means of getting back again to the question of the properties of N which are valuable for our purposes. Suppose that the following states one of the properties in P: for a system with an open-loop transfer function of $1/(j\omega+1)^3$, the gain margin will be eight. We can imagine the situation in which the system we are examining, A, has not quite the specified open-loop transfer function and not quite a gain margin of 8, perhaps 7.11. Suppose we are in a situation in which only a control system with a gain margin greater than 7.75 will have a noticeable effect on the controlled variable. Then test 3 seems to ask whether or not we have some way of bringing the gain margin of A to within .25 of the performance predicted by the classificatory explanation, but this is not quite correct. Suppose that we are unable to close the gap to within .25 of the predicted gain margin. Are we to conclude that A is a feedback control system, but not the right one for this application, or that A is not a feedback control system? The answer seems to depend upon the linguistic context in which the classificatory judgment is made. In some contexts, anything that can be said to have an open-loop transfer function is a feedback control system, whether or not its gain margin is so small that its controlling effects would be undetectable. In other contexts, we are more choosy. In the context of this paper, we may be willing to grant that pencils are or contain feedback control systems. In other contexts, we would

deny it. Once again I am unable to specify the subtle criteria which we apply in making such judgments, but I have pointed at the kinds of things which I think are relevant.

There are a number of ways in which an item might pass test 3, and we can list these as subtests:

SUBTEST 3A: Can the gap be reduced by altering the set of structural properties of the model used in the classificatory explanation?

In many cases, we can make the idealization of the structure of the item closer to the actual item by "dirtying" it up with terms for frictions, inertias, and nonlinearities. This in itself will reduce the gap, since the structural properties are a subset of the complete set of properties of the model, but also it will often make the predicted performance properties closer to the actual performance properties.

SUBTEST 3B: Can the gap be reduced by improving the laws used in the classificatory explanation?

The laws which are a part of classificatory explanations are, as I have remarked earlier, derivative laws. A set of these laws is taken to apply in a given situation, depending upon how the item under investigation is modeled. The laws might be improved by, for example, discovering more precisely the values of certain coefficients. Improving them in this way might bring the predicted performance properties closer to the actual performance properties.

SUBTEST 3C: Can the gap be reduced by altering A
itself?

Sometimes the actual performance can be brought closer to the ideal by "improving" the item under consideration, perhaps by reducing the effects of friction by oiling its joints, reducing the effects of inertia by using lighter materials, and so on.

We want to say that an artifact passes test 3 and thus is to be judged an item of type N, e.g., a feedback control system, if it passes any of the subtests. Unfortunately, in passing subtest 3C, the item being assessed is itself changed. We are left wondering whether to judge that item A which has been changed into an N should be considered to have been an N in its original state. This seems too strong; we can turn almost anything into a feedback control system. Again, we seem to be making a subtle judgment, this time having to do with how much we had to alter A to narrow the gap sufficiently to make it a certifiable N.

Subtests 3A through 3C ask whether we can reduce the shortfall between the actual and predicted properties of A by tinkering with certain properties of the idealized model (P), the laws (L), or the actual item under consideration (A). I see no reason to suppose that the gap must be narrowed by only one of these methods at a time. Every time P is changed, a slightly different, although not necessarily "improved" in the sense I mentioned earlier, law is brought into the classificatory explanation.

Notice how these tests and subtests are related to the justificatory arguments which we discovered in the literature of automatic

control theory. The first of these arguments I called the "negligible discrepancy" argument. This argument amounts to the claim that a particular idealization is justified because test 2 would be passed by any item in the reference sample, despite the proposed idealization. The second justificatory argument I called the "possibility of improvement" argument, and it is simply the claim that a particular idealization is justified because test 3 would be passed by any item in the reference sample, despite the proposed idealization. Finally, I discussed what I called the "limiting case" argument, which I believe may be interpreted as the claim that a particular idealization is justified because test 3C would be passed by any item in the reference sample, despite the proposed idealization. We should not be surprised to find in the literature justificatory arguments which correlate to subtests 3A and 3B.

We have seen that if item A passes test 1 and has all the properties in set P, we may classify it as being of the kind, N, and the same is true if it passes test 2 and comes close enough to having the properties of P. Even if it fails tests 1 and 2, it might still be classifiable as an N if it passed test 3 (via subtest 3A, 3B, or 3C or some combination thereof) by having an actual/ideal properties gap which could be reduced (and we have an idea of how it could be done) to a level that would pass test 2. Now what happens if item A fails tests 1, 2, and 3 with respect to a particular classificatory explanation? Then, instead of worrying about whether or not A satisfies the requirements of the classificatory explanation, E, we turn our attention on E itself:

TEST 4: Is there any reason for rejecting E?

If E is completely satisfactory, then we must conclude that A is simply not an N. There are two sorts of things which might be wrong with E which have not already been dealt with in a test. These may be expressed as subtests to test 4:

SUBTEST 4A: Are the laws used by E correct?

From time to time we discover that something we thought to be a law is not one after all. If E employs a derivative of such a "law," then E doesn't explain why the items in the reference sample have the performance properties that they do, given their structural properties. Some other explanation, employing a true law might do that. As I have noted earlier, we are willing to revise our applications of a law if faced with empirical counterevidence. But the laws, themselves, especially if well entrenched, are not overthrown by counterevidence alone, and the process of changing them is a major Kuhnian scientific revolution. If E fails subtest 4A, then, whether the problem is that we have applied the law incorrectly or that the law itself is incorrect, the possibility is left open that A is an N after all, since E is not an acceptable account of N-type items. We had better withhold judgment and seek an improved classificatory explanation.

SUBTEST 4B: Is the reference sample homogeneous?

Perhaps E is being asked to perform an impossible task. Suppose that the items in the reference sample actually have no structural properties in common (no common "hidden structure"), then whatever idealization E makes concerning the structural properties of the reference sample items

must be wrong. It would make no sense to require that E explain why something with the structural properties of the items in the reference sample should have the performance properties that it does, since there would have to be a different explanation for each item. Our stage II hypothesis that the items in the reference sample were all representatives of the same kind of thing was false. The question of whether or not A should be considered to be the same kind of thing becomes otiose.

There is another way in which E could fail subtest 4B. It might be the case that the reference sample is not homogeneous not because each item in it shares no significant structural properties with any others, but because the reference sample contains two or three or n subsets of items which do have common "hidden structures." This was the situation which Maxwell thought he had discovered when he announced in his 1868 paper, "On Governors," that we should make a distinction between regulating devices which are governors (type I automatic control systems in today's terminology) and those which are merely moderators (type 0 systems).⁴⁰ Although both are devices whose purpose it is to regulate the movements of machines and to correct irregularities in the motions, in the case of the moderator, a particular deviation of the actual speed from the desired speed results in a particular position of the throttle valve (i.e., a constant actuating signal results in a constant value for the controlled variable), while, in the case of the governor, a particular deviation of the actual speed from that desired results in a particular rate of change of the throttle valve (i.e., a constant actuating signal results in a constant rate of change of the controlled variable). What Maxwell had noticed is a difference in

performance properties between two subsets of things which had previously all been called governors. Newton's laws could account for the difference in performance properties only if members of the two subsets had different structural properties, so Maxwell hypothesized that their structural properties were different in some important way and devised two different classificatory explanations on the basis of that hypothesis. Perhaps also there were obvious structural differences between the two classes of regulators, and Maxwell was merely pointing out that we could see a difference in performance in performance properties which corresponded to the differences in appearance. But, as Otto Mayr points out in "Maxwell and the Origins of Cybernetics," Maxwell misclassified the regulators built by Thomson and Foucault: the equations Maxwell used to describe them (his idealized descriptions of their properties) describe governors, despite the fact that they are really moderators according to Maxwell's own distinction.⁴¹ Maxwell was apparently fooled by their superficial appearance into thinking that they were governors!

Now suppose that the explanation fails test 4B. First, it is not so much a failure of the explanation as it is a failure of our original hypothesis back in stage II of the development of the concept that the items in the reference sample were the same kind of thing. Second, failure of test 4B may give us grounds for giving up the concept altogether or for supposing that we've actually been dealing with two or more concepts when we thought we had gotten hold of only one, or we might decide that certain of the items in the original reference sample should not have been there to begin with. Finally, failure of test 4B leaves open the possibility that item A is in the extension of the concept term

after all.

In the summary of the classification program displayed in Figure 1, we see the dynamic nature of complex concepts. Every element of the conceptual program is defeasible. Given the appropriate circumstances, we might need to change any of the following:

1. The properties thought to be stereotypical of the concept, including both structural and performance properties
2. The explanation of why something with the stereotypical structural properties should have the stereotypical performance properties
3. The laws used in the classificatory explanation
4. The set of items thought to be representative of the concept


A particular item thought to be stereotypical of feedback control systems at one stage of the development of the concept might later be ruled out as not only not typical of feedback control systems, but not a feedback control system at all.

Finally, consider the role of indexicality in the dynamic theory of conceptualization which I have sketched. The indexical element is supplied by the reference sample. In stage II, we notice a certain set of things as having valuable properties, and we hypothesize that they are members of a homogeneous class of items which have similar structural and performance properties. We then tug and pull at the concept to make it as general as it can be and still retain its usefulness. In this way, we maximize its usefulness to us. Throughout the process of concept expansion, the reference sample serves us an anchor which keeps

the program from drifting away from the project of characterizing the sort of thing which was originally found to have the remarkable or valuable properties back in stage II. In addition, test 2 (Is the gap between the observed and modeled properties too large?) is in part answered by reference to the difference between the ideal model and the paradigm instances. Any difference less than this is not too large; a difference larger than this is (perhaps) prima facie too big.

The indexical element is essential just because we don't know exactly how to characterize the kind of thing in the reference sample; we don't know its "hidden structure." Thus, we operate under conditions of uncertainty or partial knowledge like those which I discussed in the second chapter. We need the reference sample or exemplar set to point to when we refer to the kind or type of thing, N, attributively as "whatever is the same kind of thing as the items in the reference sample."

Despite the importance of the reference sample as an anchor for the conceptual program, the stage II hypothesis which establishes it is defeasible. As Putnam says, "The local water, or whatever, may have two or more hidden structures--or so many that 'hidden structure' becomes irrelevant, and superficial characteristics become the decisive ones."⁴²



E is a classificatory explanation, relative to purpose D, consisting of laws L and a model with properties P, according to which items of kind N have properties P. Classifying, as described above, is simply applying a concept to nonexemplars for that concept.

Fig. 1. Block Diagram of Classification Process

NOTES TO CHAPTER THREE

¹ Joseph J. Di Stefano, Allen R. Stubberud, and Ivan J. Williams, Feedback and Control Systems (New York: McGraw-Hill Book Company, 1967), p. 5.

² [Horace] Rom Harré, The Principles of Scientific Thinking (Chicago: University of Chicago Press, 1970), pp. 234-56.

³ Otto Mayr, "Maxwell and the Origins of Cybernetics," Isis 62 (1971):424-44.

⁴ James Clerk Maxwell, "On Governors," Proceedings of the Royal Society of London 16 (1868):271.

⁵ Ibid., pp. 271-72.

⁶ Edward John Routh, A Treatise on the Stability of a Given State of Motion, Particularly Steady Motion (London, 1877), p. vi.

⁷ Di Stefano, Feedback, pp. 28-30.

⁸ Maxwell, "Governors," p. 278.

⁹ Di Stefano, Feedback, p. 179.

¹⁰ S. Bennett, A History of Control Engineering, 1800-1930 (Stevenage, U.K.: Peter Peregrinus, Ltd., 1979), p. 82.

¹¹ N. Minorsky, "Directional Stability of Automatically Steered Bodies," Journal of the American Society of Naval Engineers 42 (1922): 283.

¹² Ibid., p. 283.

- 13 Ibid., pp. 304-5.
- 14 Ibid., p. 286.
- 15 Ibid., p. 284.
- 16 H. L. Hazen, "Theory of Servo-Mechanisms," Journal of the Franklin Institute 218 (1934):287.
- 17 Ibid., p. 304.
- 18 Ibid., p. 315.
- 19 Ibid., p. 316.
- 20 Ibid., p. 323.
- 21 Ibid., p. 287.
- 22 Ibid., p. 288.
- 23 Ibid., p. 283.
- 24 Ibid., p. 285.
- 25 Ibid., p. 287.
- 26 Ibid., p. 290.
- 27 Ibid., p. 296.
- 28 Ibid., p. 299.
- 29 Harry Nyquist, "Regeneration Theory," Bell System Technical Journal 11 (1932): 126.
- 30 H. S. Black, "Stabilized Feedback Amplifiers," Bell System Technical Journal 13 (1934):9.
- 31 A. Ivanoff, "Theoretical Foundations of the Automatic Regulation of Temperature," Journal of the Institute of Fuel, London 7 (1934): 118.
- 32 Minorsky, "Directional Stability," p. 283.

³³ A. Callender, D. R. Hartree, and A. Porter, "Time Lag in a Control System," Philosophical Transactions of the Royal Society of London A235 (1936):416.

³⁴ Ibid., p. 420.

³⁵ H. W. Bode, "Relations between Attenuation and Phase in Feedback Amplifier Design," Bell System Technical Journal 19 (1940):429.

³⁶ Ibid., p. 446.

³⁷ Otto Mayr, The Origins of Feedback Control (Munich: R. Oldenbourg Verlag, 1969) English ed. (Cambridge: The M.I.T. Press, 1970), p. 131.

³⁸ For a general introduction to the theory of amorphous solids, see: David Adler, "Amorphous-Semiconductor Devices," Scientific American, May 1977, pp. 36-48; and Robert Haavind, "Glassy Semiconductors--a Missed Bet?" High Technology, February 1980, pp. 34-41.

³⁹ Neville F. Mott and E. A. David, Electronic Processes in Non-Crystalline Materials (Oxford: Oxford University Press, 1971).

⁴⁰ Maxwell, "Governors," pp. 274-75.

⁴¹ Mayr, "Maxwell and the Origins of Cybernetics," pp. 432 & 435.

⁴² Hilary Putnam, Philosophical Papers, vol. 2: Mind, Language and Reality (Cambridge: Cambridge University Press, 1975), p. 241.

CHAPTER IV

JUDICIAL CLASSIFICATION

The little case, the ordinary case, is a constant occasion and vehicle for creative choice and creative activity, for the shaping and on-going reshaping of our case law.

That is our system of precedent.¹ [*italics deleted*]

Karl N. Llewellyn, The Common Law Tradition

Precedents as Paradigms

In previous chapters, we have developed a picture of the processes by which classification systems based on paradigms are developed and elaborated. The examples used have all concerned physical objects of one sort or another. In the present chapter, I want to show that the picture we have developed can be applied to the classification of abstract objects as well. For this purpose, I shall focus on a particularly well-known legal decision, the judgment in MacPherson v. Buick Motor Co.²

Legalese is sodden with slippery concepts. Until anchored by precedents and carefully drafted legislation, such notions as estoppel, malice, and privity are too vague to be of much practical value. The systems of classification which such notions as these represent are based, in part, on the legal paradigms we call precedents. The job of an appeals court judge is a tricky one, since he must not only render

decisions that are just and fair, but also he must attempt to demonstrate that his judgment is justified, given previous apparently similar cases--all the while operating under the knowledge that his own decision will be a precedent for later cases. The tradition of publishing written "opinions" which give the rationale for judicial decisions makes legal classificatory judgments especially suitable for my purposes. Whereas an engineer may decide to treat a particular item as a feedback control system or as some other sort of system, there is no tradition in the engineering profession of attempting to show that such a decision is justified in the light of the available evidence and of the history of the concept. As we discussed in an earlier chapter, the engineer may include some remarks to justify the more blatant idealizations, but generally, his reasons for thinking that a particular classificatory judgment would be correct are not made public. Of course, a confirmation of the judgment becomes public after the fact, when the engineer finds that he can or cannot get the results that he is seeking, but this sort of after-the-fact confirmation is available to the legal profession also: judges may come to believe that justice was or was not done by a particular decision. Nevertheless, reasons can be given for believing that a particular classificatory judgment would be just or would help achieve some engineering objective. Engineers are not generally in the habit of elaborating their reasons; judges are.

It should not be supposed that the opinion given by a judge is a description of the process by which he arrived at his decision.³ Attempts to describe the judicial "logic of discovery" appear in works on jurisprudence, written out of court. Perhaps the most famous is Judge

Benjamin Cardozo's The Nature of the Judicial Process, and, as Cardozo indicates, the business of describing the creative process is no easier in law than it is in science:

The work of deciding cases goes on every day in hundreds of courts throughout the land. Any judge, one might suppose, would find it easy to describe the process which he had followed a thousand times and more. Nothing could be farther from the truth. Let some intelligent layman ask him to explain: he will not go very far before taking refuge in the excuse that the language of craftsmen is unintelligible to those untutored in the craft.⁴

So, written legal opinions give not an account of the court's thoughts in arriving at a decision, but arguments intended to justify the claim that the decision reached is correct.

It should also not be supposed that the existence of a previous decision on a similar case always provides a binding precedent. As we shall see, this is far from the case. It is always possible to distinguish the case at hand from earlier cases. Indeed, it has been claimed that it is as important for a court to give reasons for following a given precedent as it is to give reasons for not following one.⁵

In this chapter, I want to examine in detail a particular legal decision for two reasons: (1) to demonstrate that the picture we developed earlier of how classification systems are developed applies to a system for classifying abstract objects, viz., an abstract legal concept; and (2) to uncover some of the "logic" of judicial classification, e.g., to see what arguments might be given for altering a given system of classification.

The Requirement for a Model of L⁻

In 1916, Judge Cardozo, then serving on the Court of Appeals of the State of New York, wrote the majority opinion for MacPherson v. Buick Motor Co. Here are the facts of the case as described by Judge Cardozo:

The defendant is a manufacturer of automobiles. It sold an automobile to a retail dealer. The retail dealer resold it to the plaintiff. While the plaintiff was in the car, it suddenly collapsed. He was thrown out and injured. One of the wheels was made of defective wood, and its spokes crumbled into fragments. The wheel was not made by the defendants; it was bought from another manufacturer. There is evidence, however, that its defects could have been discovered by reasonable inspection and that inspection was omitted. There is no claim that the defendants knew of the defect and willfully concealed it. . . . The charge is one, not of fraud, but of negligence. The question to be determined is whether the defendant owed a duty of care and vigilance to any one but the immediate purchaser.⁶

Our intuitions about what justice and good public policy require here conflict. On one hand, we think that consumers should not have to accept being maimed or killed by products rendered dangerously defective by a manufacturer's negligence. They or their survivors should be able to recover damages for injuries resulting from the manufacturer's negligent act. But on the other hand, we think there should be some limit to a manufacturer's liability. In the hands of some consumers, almost anything can be dangerous. Also, a manufactured product may be put to uses not reasonably envisioned by the manufacturer. As a matter of public policy, suppliers of a country's goods and services should be protected from unreasonable prosecution. Excessive product liability can cause manufacturers of some essential goods and services to go out of

business and can result in increased costs for products and services which are available. Further, it inhibits new product development. For example, in 1976, a nationwide epidemic of swine flu was predicted. Congress appropriated \$135 million for a mass immunization program, but the pharmaceutical companies were unwilling to produce the vaccine until Congress passed legislation in which the federal government assumed a part of the risk of the program.⁷ In this case, a socially beneficial product was withheld because of the manufacturers' unwillingness to subject themselves to what they judged to be unacceptable legal risks.

The problem in MacPherson v. Buick Motor Co., then, is to discover what degree of liability would be reasonable. In terms of classification, the problem is to determine the class of cases in which it would be just to hold a manufacturer or supplier of a defective article liable for injuries resulting from the defects in his product. Call this class L. The final decision in MacPherson v. Buick Motor Co. depends on whether or not the case is found to be in L.

We may subdivide L into two subsets. In the first subset will be the cases in which the injured party has purchased directly from the manufacturer or supplier the item which caused the harm. These cases are in L because the contract between buyer and seller gives the seller a prima facie duty to avoid negligence in fulfilling the contract or at least to give the unwary buyer fair warning of the potential danger.⁸ Let us call this subset L⁺, to remind us that: (1) in this class of cases, it would be just to hold the manufacturer liable for his negligence; and (2) there is a contract between the injured party and the manufacturer. Probably not every case in which there is a contract

between the manufacturer and the injured party is an \underline{L}^+ case. Other factors, such as the negligence of the injured party in using the product or causal overdetermination of the injuries might cloud the issue.⁹

Let us use " \underline{L}^- " to designate the remaining subset of \underline{L} which is not included in \underline{L}^+ . Cases in \underline{L}^- would be cases in which: (1) it would be just to hold the manufacturer liable for his negligence, but (2) there is no contract between the injured party and the manufacturer. $\underline{L} = \underline{L}^+ \cup \underline{L}^-$, and $\underline{L}^+ \cap \underline{L}^- = \emptyset$. Since Donald MacPherson bought his new Buick from an automobile retailer and not directly from the manufacturer, we can see that, if MacPherson v. Buick Motor Co. is an \underline{L} case, it must be an \underline{L}^- case. There was no contract between the manufacturer and the injured party.

It is certainly false that every case in which there is no contract between a negligent manufacturer and an injured consumer is an \underline{L}^- case. Indeed, in a standard work on the law of torts, Thomas Cooley notes that "the general rule is that a contractor, manufacturer, vendor or furnisher of an article is not liable to third parties who have no contractual relations with him for negligence in the construction, manufacture or sale of such article."¹⁰ That is, according to the general rule, $\underline{L}^- = \emptyset$. Nevertheless, there are some cases in which it seems just to hold a manufacturer liable for his negligence, regardless of the lack of a contract between the manufacturer and the injured party.

In the next several pages, I shall review the court's attempts to establish and elaborate \underline{L}^- .¹¹ The pattern is a familiar one. The court decides certain fairly clear-cut cases and attempts to construct a classificatory explanation consisting of legal rules and principles

plus a model of the \underline{L}^- cases. The model undergoes continual modification as the \underline{L}^- classification is applied to specific cases.

I use the term "model" here in the same sense in which we use the term in talking about Watson and Crick's model of the DNA molecule or Bohr's model of the atom. In this sense, a model of \underline{L}^- cases is simply an account of what it is to be an \underline{L}^- case, relative to a particular theory. We know, of course, that any plausible \underline{L}^- model must satisfy the following conditions: (1) it would be just to hold a manufacturer liable for injuries caused by his negligence in the case, and (2) there is no contract (regarding the injury-causing article) between the injured person and the manufacturer. But the properties of those cases which we take to satisfy (1) and (2) will depend on our theories of contracts and justice.

Suppose we have on hand a set of cases in which it was ruled that (1) and (2) were satisfied, and we are convinced that, in these cases, justice was done. In order to judge properly in future cases which might not be so clear, we might examine these cases very carefully and try to discover what they have in common. This done, we could make the tentative hypothesis that any case which has all the properties shared by the paradigm cases is itself an \underline{L}^- case. But chances are that the list of shared properties would be extremely long and would contain many shared properties which are simply irrelevant to the question of whether or not a given case is an \underline{L}^- case. What is needed is a theory about which properties are relevant and why. The set of properties identified as relevant by such a theory constitutes the model of \underline{L}^- cases for that theory.

Langridge v. Levy

The history of the L⁻ classification appears to begin in 1837 with Langridge v. Levy.¹² In this case, heard by the English Exchequer of Pleas Court, the court decided in favor of the injured plaintiff, but not on the grounds that the case was an L⁻ case. The facts of the case are these. On the first of June, 1833, George Langridge, father of the plaintiff, visited the shop of Mr. Levy, a gun maker in Bristol. There he admired a double-barreled gun to which was attached a tag which said: "Warranted, this elegant twist gun, by Nock, with case complete, made for his late Majesty George IV." Levy affirmed that the claims made on the ticket were accurate and told Langridge that he had Nock's invoice for the gun. Nock was a gunsmith with a reputation for producing firearms of exceptional quality, and Langridge bought the gun for 24 pounds sterling, mentioning to Levy that he wanted it for himself and his three sons. Unfortunately:

. . . the said gun, at the time of the said warranty and sale, was not made by Nock, nor was it a good, safe, and secure gun, but, on the contrary thereof, was made and constructed by a maker very inferior as a gun-maker to Nock, and was then and at all times a very bad, unsafe, ill-manufactured, and dangerous gun, and wholly unsound and of very inferior materials; of all which premises the defendant, at the time of the making of the said warranty, and of the said sale, had full knowledge and notice.¹³

Later, on December 10, 1835, Langridge's second son, the plaintiff in the case, fired the weapon at some birds in a field near his father's house. The barrel of the gun exploded, mutilating his left hand so severely that it had to be amputated.

In court, Langridge's lawyer argued that since Levy knew that the

plaintiff's father was purchasing the gun for the use of his sons, the court should view the case as an L^+ case, on the basis of an implied contract between Levy and the plaintiff. Barring that, the case should be construed as an L^- case on the grounds that "the law imposes on all persons who deal in dangerous commodities or instruments, an obligation that they should use reasonable care, much more that they should not supply them knowing them to be likely to cause injury."¹⁴ That is, the counsel for Langridge proposed a model for L^- :

LANGRIDGE MODEL: L^- = the set of all cases in which--

1. A person A is injured by some article D;
2. D is supplied by another person B;
3. A is not party to a contract requiring B to supply articles of type D to A;
4. B supplies D, knowing that D is likely to cause injury;
5. D is a dangerous commodity or instrument.

Langridge's lawyer argued for this model by point out the analogy with the 1816 decision of the Court of King's Bench in Dixon v. Bell. In Dixon v. Bell, the defendant sent his young servant girl to fetch a loaded gun. The servant, unaware that the gun was in a dangerous condition, playfully fired it into the face of the plaintiff's nine year old son who lost an eye and two teeth. The court held the defendant liable on the grounds that "the law requires of persons having in their custody instruments of danger, that they should keep them with the utmost care."¹⁵ The Langridge model of L^- would, it was argued, show why

Dixon v. Bell was an L^- case, and it would also provide grounds for deciding for the plaintiff in Langridge v. Levy. It would supply the basis of an explanation of why it was proper to classify Dixon v. Bell as an L^- case and why it would be correct to put Langridge v. Levy in the same class. Simply stated, the explanation would take the form of an argument something like the following:

1. Any case which has properties $P_1, P_2, P_3, \dots, P_n$ is an L^- case (given the Langridge model).
2. Both Dixon v. Bell (which is known to be an L^- case) and Langridge v. Levy have $P_1, P_2, P_3, \dots, P_n$
3. Therefore, both Dixon v. Bell and Langridge v. Levy are L^- cases.

Unfortunately, it is not clear that the Langridge model is an acceptable model of the Dixon v. Bell case. Both Dixon and Bell lodged in the house of a Mr. Leman. Since there had been several robberies in the neighborhood, Bell maintained a gun which was loaded with powder and printer's type, a sort of poor man's buckshot. On October 10, 1816, Bell wanted his gun, and he sent his thirteen or fourteen year old mulatto servant girl to fetch it from Leman. Bell ordered the servant to ask Leman to render the gun safe by removing the priming before delivering it to her. Leman did this and told the girl so before handing the weapon over to her. Thus, Bell attempted to render the gun harmless, and he was unaware that the removal of the priming permitted some grains of gun powder to escape through the touch-hole. Since the Langridge model specifies that the supplier know that the article he supplies is

likely to cause injury, it appears that the Langridge model is not a very good model of Dixon v. Bell. Nevertheless, Langridge's lawyer believed that it was and cited that in its support.

A second argument given by Langridge's lawyer concerned liability for the attacks of dangerous animals: "If a party sold a vicious dog under a representation that he was a quiet one, and being taken home by the buyer, he bit his child; would not the seller be liable for this injury?"¹⁶

In the course of his arguments, Langridge's lawyer cited the following principles:

1. "If any subject sustained a wrong by the unjustifiable act of another, he ought to have a remedy."¹⁷
2. "Whenever by the circumstances of the transaction a duty is imposed upon the defendant, and by a breach of that duty (as distinguished from a contract) an injury happens to the plaintiff, he may sue."¹⁸
3. "The law imposes on all persons who deal in dangerous commodities or instruments, an obligation that they should use reasonable care, much more that they should not supply them knowing them to be likely to cause injury."¹⁹
4. "Prima facie, every man who suffers an injury is entitled to recover against the party who caused it, and who must be taken to have intended the natural consequences of his injurious act."²⁰

Contrast these principles with the principles appealed to by the counsel for Levy, the defendant:

1. "There is no such known right in the English law . . . where-
by the plaintiff is entitled to receive damages from the de-
fendant, with whom he made no contract."²¹
2. "Wherever an instrument is immediately dangerous, and is so
placed as to be likely to do an injury to any of the public,
the party who places it there is liable for such an injury."²²
3. "The damage must be a proximate consequence from the act of
the defendant."²³

At issue here is the question of how many intervening steps in the causal chain there can be from the defendant's act to the plaintiff's injury before the defendant is no longer legally liable for injuries which followed his act. Levy's lawyers tried to show that the cases and principles cited by the plaintiff's attorney had relevance only for cases in which there were few intervening steps from act to injury and that, in Langridge v. Levy, the chain was long and complex:

There are other cases [than Dixon v. Bell and the vicious dog case, both of which were cited by the plaintiff] which may be put, more in analogy with the present. Suppose a chain cable were sold with a warranty of its being secure, when in fact it was imperfect, and the vessel being in a storm, the cable is let go, and breaks; could it be contended that the captain and each of the crew, if injured in consequence, would have a right of action against the seller? So, supposing the owner of an unruly horse, knowing his disposition, sold him with a warranty that he was quiet to drive, and the buyer lent him to a friend, who put other persons into the carriage, and he ran away, and overturned and injured them; would the seller be liable to each of these persons?--Such liabilities would be carried to an extent wholly indefinite.²⁴

In effect, Levy's lawyers were arguing against the sufficiency of the conditions in the Langridge model. The cases cited appear to fit the Langridge model, yet they are cases in which it would not be fair to

hold the supplier of the dangerous article liable for injuries caused by that article. To replace the Langridge model, Levy's lawyers proposed a model which would show that Dixon v. Bell and the vicious dog cases are \underline{L}^- cases, while showing that the anchor chain case, the unruly horse case, and (of course) Langridge v. Levy are not. "The distinction," they said, "is this: is the instrument or other thing immediately dangerous or mischievous by the act of the defendant, or is it such as may become so by some further act to be done to it?"²⁵ So the model proposed by Levy's lawyers was the following:

LEVY MODEL: \underline{L}^- = the set of all cases in which--

1. A person A is injured by some article D;
2. D is supplied by another person B;
3. A is not party to a contract requiring B to supply articles of type D to A;
4. B supplies D, knowing that D is likely to cause injury;
5. D is an instrument or other thing made immediately dangerous by B's act.

Since Langridge's son had to load the gun before it became immediately dangerous, the Levy model would have us conclude that Langridge v. Levy is not an \underline{L}^- case, and Levy should not be held liable for the boy's injury. Dixon v. Bell was an \underline{L}^- case, because the gun had been loaded by the defendant. But, as we noted in the discussion of the Langridge model, Bell was not aware that his gun, with the priming removed, was likely to cause any injury. So, although Levy's lawyers claimed that the

Levy model was a good model of Dixon v. Bell, it was not.

Since the court decided in favor of the plaintiff in this case, we might suppose that the Langridge model was accepted, but this is not what happened. Lord Parke explicitly rejected it, saying:

. . . we should pause before we made a precedent by our decision which would be an authority for an action against the vendors, even of such instruments and articles as are dangerous in themselves, at the suit of any person whomsoever into whose hands they might happen to pass and who should be injured thereby.²⁶

But neither did the court embrace the Levy model; instead, the case was decided on the grounds that Levy's fraudulent warranty was a "falsehood told with an intention that it should be acted upon by the party injured, and that act must produce damage to him."²⁷ The importance of Langridge v. Levy is in the fact that two L⁻ models were explicitly proposed for the court's consideration: the "dangerous articles" model of Langridge's lawyer and the "imminently dangerous articles" model proposed by the counsel for Levy. Both models are attempts to find an acceptable way to delineate those cases in which justice demands that an injured noncontracting party be compensated and those in which a manufacturer or supplier deserves to be protected from excessive liability. The court avoided the adoption of any L⁻ model, although the notion of intentionally injurious fraud on which this case was decided forms an important part of the model which was later accepted by the court.

Winterbottom v. Wright

The matter was reopened in 1842 in Winterbottom v. Wright, again in the English Exchequer of Pleas Court.²⁸ But while the court had rendered a judgment for the plaintiff in Langridge v. Levy, in this case,

it found in favor of the defendant. Wright, the defendant, had contracted with the Postmaster General to provide and maintain a mail coach to carry bags of mail along a certain route between Hartford, in Chester, to Holyhead. Another man, Nathaniel Atkinson, and his associates had contracted with the Postmaster General to convey the mail coach along the appointed route and to supply horses and coachmen for that purpose. The plaintiff, Winterbottom, was one of the drivers hired by Atkinson. On August 8, 1840, Winterbottom was driving the mail coach supplied by Wright along the route from Hartford to Holyhead, when:

the said mail-coach being then in a frail, weak, inform, and dangerous state and condition, to wit, by and through certain latent defects in the state and condition thereof, and unsafe and unfit for the use and purpose aforesaid, and from no other cause, circumstance, matter, or thing whatsoever, gave way and broke down, whereby the plaintiff was thrown from his seat, and, in consequence of injuries then received, had become lamed for life.²⁹

The lawyer for Wright, the defendant, argued that the general rule of nonliability to noncontracting parties should be adhered to without exception: L should be held empty. If exceptions were permitted, "the most alarming consequences would follow the adoption of such a principle."³⁰ Examples of such "alarming consequences" included the following. If a defective axle on a train caused an accident, anyone injured in the accident could sue the manufacturer of the axle. If a gentleman's coachman were injured by the breaking down of his carriage, he could sue the smith or coachmaker, even though his contract is with his master. Also mentioned was the sort of case proposed earlier in which a ship's anchor chain breaks, the ship goes aground, and everyone affected sues the manufacturer of the chain or the suppliers of the iron.

In addition, Wright' lawyer pointed out disanalogies between this case and Langridge v. Levy. First, Levy knew the gun was bought for the use of the plaintiff, but Wright did not know that Winterbottom would be driving the coach. Second, there was fraud in the earlier case but not in the later one. Third, in Langridge v. Levy, the cause of injury was a weapon of a dangerous nature, and a coach is not a weapon at all. So, in addition to his preferred strategy of convincing the court that there should be no L^- cases, Wright's lawyer had a back-up plan. If the court were convinced that Langridge v. Levy is an L^- case, he would supply a model of L^- cases which would include Langridge v. Levy while excluding Winterbottom v. Wright.

WRIGHT MODEL: L^- = the set of all cases in which--

1. A person A is injured by some article D;
2. D is supplied by another person B;
3. A is not party to a contract requiring B to supply articles of type D to A;
4. B supplies D, knowing that D is to be used by A;
5. B fraudulently represents D to A as free of defects;
6. D is a weapon of a dangerous nature.

Notice that, since it includes the element of fraud, the Wright model would clearly exclude Dixon v. Bell, a case that both the Langridge and Levy models were supposed to include (although, as we have noticed, they did not).

Winterbottom's lawyer tried to show that, despite the defendant's claims, the cases are analogous, and the earlier case (Langridge v. Levy)

should stand as a precedent for the later one. The same points were discussed. First, there is no evidence that Levy had ever heard of the particular son who was injured, although he knew that some child of Langridge would use the gun. Here the situation is similar: Wright may not have known that Winterbottom would drive the coach, but he surely knew it would have to be driven by some coachman. Second, there was fraud in this case also:

The defendant represented the coach to be in a proper state for use, and whether he represented that which was false within his knowledge, or a fact as true which he did not know to be so, it was equally a fraud in point of law, for which he is responsible.³¹

Third, it is not the fact that the gun in Langridge v. Levy was a weapon of a dangerous nature that is important. The important point is that it was "an article which, if imperfectly constructed, was necessarily dangerous."³² So yet another model was being proposed for L^- :

WINTERBOTTOM MODEL: L^- = the set of all cases in which--

1. A person A is injured by some article D;
2. D is supplied by another person B;
3. A is not party to a contract requiring B to supply articles of type D to A;
4. B supplies D, knowing that D is to be used by some member or members of a particular class of persons, and A is a member of that class;
5. B fraudulently represents D to A as free of defects;
6. D is an article which, if imperfectly constructed, is necessarily dangerous.

The Winterbottom model also excludes Dixon v. Bell, but it includes both Langridge v. Levy and Winterbottom v. Wright as L^- cases.

Winterbottom's lawyer adduced several imaginary cases to show that the Winterbottom model more closely accords with our intuitions about justice than the Wright model. For example, if a soldier is injured when he fires a defective musket, we think that he should be able to recover damages from the firm that sold the defective rifle to the government. Or, if a coachmaker builds a coach so negligently that one of its wheels flies off, injuring a child of the owner of the coach, the coachmaker should be held liable for the injury. Again, if a contractor repairs a public building so negligently that a person using the building is injured by falling stone, the injured person should have a remedy. According to the Winterbottom model, these are L^- cases, and that seems consistent with our intuitions about what would be just. According to the Wright model, they are non- L^- cases, so the Winterbottom model is held to be superior.

Once again, the court refused to select either model. Instead, the judges followed the suggestion of Wright's lawyer and agreed that the general rule of no liability to noncontracting parties be maintained without exception. Lord Abinger argued:

There is no privity of contract between these parties; and if the plaintiff can sue, every passenger, or even any person passing along the road, who was injured by the upsetting of the coach, might bring a similar action. Unless we confine the operation of such contracts as this to the parties who entered into them, the most absurd and outrageous consequences, to which I can see no limit would ensue.³³

Judge Alderson concurred, with the comment that:

if we were to hold that the plaintiff could sue in such a case, there is no point at which such actions would stop. The only safe rule is to confine the right to recover to those who enter into the contract: if we go one step beyond that, there is no reason why we should not go fifty.³⁴

The importance of Winterbottom v. Wright lies in the fact that two new models of L⁻ were laid out for the court to consider, although the court clung to its old model of L⁻, i.e., the null set.

An L⁻ Model is Accepted:
Longmeid v. Holliday

Nine years later, in 1851, the English Court of Exchequer finally accepted a model for L⁻ other than the null set. The case at issue was Longmeid v. Holliday.³⁵ The defendant, Holliday, kept a shop in London for the sale of a type of oil lamp which he called "The Holliday Lamp." Holliday did not construct the parts for his lamps, nor did he assemble them himself; instead, he had contracted to have both of these tasks performed by others. One day, Frederick Longmeid came into the shop and purchased one of the lamps for ten shillings on Holliday's assurance that the lamp was sound and suitable for lighting Longmeid's shop and rooms. Unfortunately, and unknown to either man, the lamp was "made of weak and insufficient materials, and then was cracked and leaky, dangerous, unsafe, and wholly unfit and improper for use."³⁶ So when the plaintiff, Eliza Longmeid, lit the lamp which her husband had purchased, it exploded, covering her with burning naphtha and causing severe burns.

Longmeid's lawyers argued that the court should find for the plaintiff on the basis of the decision in Langridge v. Levy. They further argued that the decision for the defendant in Winterbottom v.

Wright was inapplicable because, in that case, there was no breach of duty to the plaintiff, whereas, in this case, there was, viz., fraud.

There is a general duty on every shopkeeper who sells articles which are or may become dangerous, to take care that they are proper for use. If he is not himself the manufacturer, and therefore is not aware that the article is unsafe, he should so inform the purchaser; but if he sells it as secure, he is guilty of a breach of duty, which renders him responsible to every one who is in consequence injured.³⁷

Thus, the Longmeid model for L^- is:

LONGMEID MODEL: L^- = the set of all cases in which--

1. A person A is injured by some article D;
2. D is supplied by another person B;
3. A is not party to a contract requiring B to supply articles of type D to A;
4. B represents D as free of defects when B either knows that D is defective or does not know that D is free of defects;
5. D is an article which is or might become dangerous.

This model could be used in the manner described earlier to account for the decisions in Dixon v. Bell and Langridge v. Levy, but it does not account for the Winterbottom v. Wright decision. According to the Longmeid model, Winterbottom v. Wright should have been decided for the plaintiff, and it was not. Longmeid's lawyers cited Langridge v. Levy and also several cases in which a surgeon was held "liable for injury resulting to a patient from his unskillful treatment, although the patient neither employed nor was to pay him."³⁸

The attempt of Longmeid's lawyers to distinguish the present case from Winterbottom v. Wright is confusing at best. As noted above, their other claims seem to indicate a model which includes both Winterbottom v. Wright and the present case, Longmeid v. Holliday, as L cases. Nevertheless, they said that Winterbottom v. Wright is "distinguishable, inasmuch as there the plaintiff was no party to the contract, neither was there any breach of duty towards him."³⁹ But the plaintiff in Longmeid v. Holliday, Eliza Longmeid, was no party to the contract of sale for the lamp; that was between Holliday and Frederick Longmeid, her husband. Further, if Holliday had a duty to refrain from representing as sound an article which he did not know to be sound, so did Wright. But the court had found that Wright had no such obligation; indeed, the lawyers for Holliday cited Winterbottom v. Wright as an "express authority" that "no duty is imposed on a tradesman to furnish articles fit for the purpose of every individual into whose hands they may come."⁴⁰ If there were such a duty, then, they pointed out, every passenger injured in a steam boat or omnibus accident caused by some hidden defect in construction could sue the builder.

Holliday's lawyers further argued that the lesson to be learned from Langridge v. Levy is that injured noncontracting parties may recover damages for injuries resulting only from fraudulent representations concerning dangerous articles. In other words, they wanted to construe the Longmeid v. Holliday case in terms of the Langridge "dangerous commodity or instrument" model, plus the stipulation that the supplier knowingly committed fraud. In the original hearing of Longmeid v. Holliday, the jury found that no fraud had been committed.

The judgment of the court was in favor of the defendant. Lord Parke rendered the opinion. There was no fraud, so Langridge v. Levy does not apply, since that case was decided on a determination of fraud. Nevertheless, Langridge v. Levy is, the court ruled, a genuine L⁻ case, and further, Lord Parke said, "there are other cases, no doubt, besides those of fraud, in which a third person, though not a party to the contract, may sue for the damage sustained, if it be broken."⁴¹ Examples of such cases occur when a person is injured by the malpractice of a surgeon or apothecary, the carelessness of a coach driver, or the defective construction of a road builder. Additional cases occur "when any one delivers to another without notice an instrument in its nature dangerous . . . or if he places it in a situation easily accessible to a third person who sustains damage from it."⁴² The court cited Dixon v. Bell as "a very strong case" in favor of this model. Cases involving the provision of "a machine not in its nature dangerous,--a carriage for instance,--but which might become so by a latent defect entirely unknown, although discoverable by the exercise of ordinary care," were specifically excluded from the model accepted by the court.⁴³ The result is the following:

PARKE MODEL: L⁻ = the set of all cases in which EITHER--

- 1a. A person A is injured by some article D;
- b. D is supplied or made available to A by another person B;
- c. A is not party to a contract requiring B to supply articles or services of type D to A;

- d. D is a type of article or service which places on its supplier a duty to the public to avoid negligence in supplying it;
- e. B is negligent in supplying D; OR
- 2a. A person B knowingly tells a falsehood with intent to induce some other person A to do an act F which could reasonably be expected to result in A's loss;
- b. A does F because of B's falsehood;
- c. A is injured as a result of doing F;
- d. A and B are not parties to any contract concerning A's doing F or B's inducing A to do F.

An example of the sort of article or service which places on its supplier a duty to the public is anything which is in its nature dangerous. In supplying instruments in their nature dangerous, the supplier incurs a duty to the public to give notice of their dangerous nature.

On the basis of the Parke model, the court decided in favor of the defendant, Holliday. An oil lamp is not an instrument which is in its nature dangerous, and Holliday did not fraudulently misrepresent his product. The model also serves as the basis of classificatory explanations which account for the decisions in the other cases we have discussed. Dixon v. Bell is an L⁻ case because Bell supplied a loaded gun, an instrument which is in its nature dangerous. Langridge v. Levy is an L⁻ case because Levy sold a gun with a defect which made it dangerous, knowingly misrepresenting it as safe. Winterbottom v. Wright is not an L⁻ case because a coach is not an instrument in its nature dangerous

(and condition ld is not satisfied for any other reason), and Wright was unaware of the defect in the coach which injured Winterbottom. In his opinion, Lord Parke explicitly noted these facts, mentioning Dixon v. Bell and Langridge v. Levy by name and Winterbottom v. Wright by implication in listing a carriage as a paradigm example of "a machine not in its nature dangerous, . . . but which might become so by a latent defect."⁴⁴

A Puzzle: Retroactive Reasoning?

At this point we should perhaps stop and take note of an interesting phenomenon. In each of the cases we have examined, the court has stated clearly its reasons for deciding for the plaintiff or for the defendant. In every case prior to Longmeid v. Holliday, these reasons were intentionally and explicitly independent of any model of L^- . But in Longmeid v. Holliday, Lord Parke seems to have attempted to supply an alternative explanation for the decisions. This seems odd. Consider the situation in more general terms. At time t_1 , person S does X and gives Y as the reason. At some later time t_2 , another person S^* gives some other item Y^* as the reason for doing X at t_1 . This seems rather presumptuous on S^* 's part, unless S^* judges that S had been confused at t_1 about his true motives for doing X or that S had simply lied about his reasons for doing X . But this case is different.

Lord Parke, the S^* of our story, did not suspect that his brethren were either confused or lying. Instead, something else is afoot. Recall that a legal opinion is not a record of the process by which a court arrived at a particular judgment. It is also not necessarily a

listing of the judges' motives for deciding as they did. Instead, the opinion is a classificatory explanation intended to justify the judgment reached. The explanation consists of: (1) a statement of applicable legal rules and principles, and (2) a model of the case at issue as well as those cases being cited as precedents. Thus, when Lord Parke offered a reexplanation of the judgments rendered in earlier cases, he was not supposing that the judges who wrote the earlier opinions were saying something false. Instead, he took the other cases as data and offered his explanation as a better account of the data than had been given in the previous individual decisions. Demonstrating that the Longmeid v. Holliday decision is consistent with the explanation served to justify that decision.

Competition among Classificatory Explanations

Before we continue, something needs to be said about what would make one classificatory explanation "better" than another. The explanation offered by Lord Parke seems to be superior to those found in the opinions of previous cases because it accounts for more of the data. The previous opinions were, for the most part, intentionally written narrowly, to cover the data of just the case at hand. Not one of the various competing models which we have discussed can serve as the basis of a classificatory explanation of each of the cases discussed. The Langridge model from the case about the gun that blew up requires that the cause of injury be a dangerous commodity or instrument and that the supplier be aware of the danger, but the defendant in Dixon v. Bell believed that the gun in that case was safe, since he had ordered the priming removed. The more conservative Levy model requires that the

cause of injury be an instrument or other thing made immediately dangerous by the supplier and that the supplier be aware of the danger. But again, the model fails to fit Dixon v. Bell, and it also does not fit Langridge v. Levy, since the defective gun was not made immediately dangerous by Levy. Both cases were ruled L cases by the court.

The Wright model from the mail coach case requires that the injury-causing article be a weapon of a dangerous nature, fraudulently misrepresented by the supplier. But once again, this model fails to fit Dixon v. Bell in which there was no fraud. And the Winterbottom model stipulates that the cause of injury be an article which, if improperly constructed, is necessarily dangerous and that it be fraudulently misrepresented by the supplier. So again, it fails to model Dixon v. Bell, and it also fails to model Winterbottom v. Wright, since the court decided that that was not an L case.

Finally, in the exploding lamp case, Holliday's lawyers favored Langridge's dangerous commodity or instrument model plus fraud, so it also fails to account for Dixon v. Bell. And the Longmeid model stipulates that the cause of injury be an article which is or might become dangerous, either intentionally or unintentionally misrepresented as safe by the supplier. But this model fails to account for any of the decisions except Langridge v. Levy.

So the Parke model forms the basis for a classificatory explanation which accounts for many more decisions than either the written opinions or the various models which were considered but not accepted by the court. In this sense, it is superior as an explanation of the data to anything considered earlier. But there is a problem about what

should be taken as "the data." If Dixon v. Bell were eliminated as a paradigm L case, the Langridge model, the Wright model, and the version of the Langridge model favored by Holliday's lawyers would each provide an adequate account of the remaining decisions. Why should the Dixon v. Bell decision be retained as a part of the data, thus giving the Parke model a victory over its competitors? Perhaps because Langridge v. Levy was cited as a precedent in both of the later cases, and Dixon v. Bell was cited as a precedent by the winning side in Langridge v. Levy. It is interesting to note, however, that, although Lord Parke delivered the opinion of the court in Langridge v. Levy, he did not at that time mention Dixon v. Bell as a precedent. To combat the Parke model, then, one could begin by showing that the counsel for Langridge made an error in giving Dixon v. Bell as a precedent in Langridge v. Levy.

Accepting a Precedent

In The Common Law Tradition, Karl Llewellyn documents sixty-four separate techniques which judges commonly use in handling precedents, including sixteen techniques for escaping the results that a particular precedent might seem to demand. He reckons only two of the sixteen techniques to be flatly illegitimate: (1) knowingly disregarding the precedent case, and (2) misrepresenting or misclassifying the facts in the pending case in order to evade the influence of the earlier case.⁴⁵

If we consider the various elements which make up a court case, I think we can easily grasp the manner in which the remaining sixty-two techniques would work. The following possibly incomplete list names

the significant features of court cases. Consider the argument schema with which we were working earlier:

1. If a case has $p_1, p_2, p_3, \dots, p_n$; it is an N case.
2. Case A has $p_1, p_2, p_3, \dots, p_n$.
3. Therefore, case A is an N case.

For any case which has been decided, there is a decision, and this is what is represented by item 3 above. Second, there is a justification of the decision or, in the terms I have been using, a classificatory explanation. That is what is represented by items 1 and 2. A classificatory explanation consists of two parts: (1) a statement of the facts of the case, represented by item 2; and (2) a statement of the legal rule or principle used to justify the decision, represented by item 1. As we saw earlier, such rules are tied to theory-based models of what it is to be an N case. Third, there is some sort of authority for the rule. In America, this could take the form of statutes, the Constitution, or prior decisions. Fourth, are what are called "dicta," the additional comments which a judge may make concerning issues related to the pending case, but not really essential to deciding it. For example, a judge may talk about how he would have decided the case if it had been different in one way or another. The fifth and final element is the post-decision history of the case. The case may have been cited as a leading precedent in a hundred subsequent cases, or it may have been overturned at a later time. In summary, court cases have decisions supported by classificatory explanations composed of statements of the facts and of legal rules for which there is some authority. There may

be dicta, and there will be a history, although that history could be that, after being decided, the case was ignored and never thought of again.

Now where is the slack in all of this which permits what Llewellyn calls the "leeways of precedent"? First, consider the facts. We would like to say that, if the facts in the pending case are the same as the facts of an earlier case, the earlier case is a precedent for the later one, but not if the facts are different. Unfortunately, any event or condition admits of an extremely wide variety of true descriptions. It appears that, for any two cases, we can give a true description according to which they are different. We notice that descriptions of the former type must be very general for cases which are "really" different, but they can be quite specific for cases which are "really" the same. Still, if the cases are actually two and not one, there comes a level of specificity at which they are revealed to have different facts. By adjusting the level of specificity of our descriptions of the cases, we may "demonstrate" that they are different (and neither can serve as a precedent for the other) or the same (and the earlier case must be taken as a precedent for the later one). If we are not cynically manipulating the facts as in Llewellyn's two "flatly illegitimate" techniques, what level of specificity should we use in deciding whether or not two cases have the same facts? The answer depends on which facts we take to be "material," and that is related to the rule or principle on which the previously decided case was decided. Let us then shift our attention to the flexibility associated with the rule or principle on which a case is decided, the ratio decidendi.

There are two questions to be decided here. First, what was the rule on which the earlier case was decided? Second, should that rule be applied in the pending case? In a classic article entitled "Determining the Ratio Decidendi of a Case," Arthur L. Goodhart argues that the principle of a case may be determined by taking account of which facts were found to be material and immaterial by the judge and taking account of his decision based on them. The ratio decidendi may then be expressed as a conditional sentence of the following form:

If a case has material facts M, then regardless of whether or not it has immaterial facts I, it should be judged an N case.

The principle of a case is not found, Goodhart demonstrates, in the reasons given in the opinion nor in the rule of law set forth in the opinion.⁴⁶ Goodhart lists ten "rules for finding what facts are material and what facts are immaterial as seen by the judge." These "rules" are, Goodhart admits, actually "tentative suggestions."⁴⁷ I won't list the rules here, but I shall try to indicate the general strategy on which the rules are based.

First, Goodhart notes, "there is a presumption against wide principles of law, and the smaller the number of material facts in a case the wider will the principle be."⁴⁸ So our first strategy will be to presume that all of the facts which may be determined from the record of a case are material. This strategy must be modified to account for the presumed generality of the law: "As a rule the law is the same for all persons, at all times, and at all places within the jurisdiction of the

court."⁴⁹ Our second policy, therefore, will be to presume that any facts of person, time, place, kind, or amount are immaterial except for those which the judge has specifically designated material. A final modification is needed to take into account the judge's role in creating law:

The same set of facts may look entirely different to two different persons. The judge founds his conclusions upon a group of facts selected by him as material from among a larger mass of facts, some of which might seem significant to a layman, but which, to a lawyer, are irrelevant. The judge, therefore, reaches a conclusion upon the facts as he sees them. It is on these facts that he bases his judgment, and not on any others. It follows that our task in analyzing a case is not to state the facts and the conclusion, but to state the material facts as seen by the judge and his conclusion based on them. It is by his choice of the material facts that the judge creates law. A congeries of facts is presented to him; he chooses those which he considers material and rejects those which are immaterial, and then bases his conclusion upon the material ones. To ignore his choice is to miss the whole point of the case. Our system of precedent becomes meaningless if we say that we will accept his conclusion but not his view of the facts. His conclusion is based on the material facts as he sees them, and we cannot add or subtract from them by proving that other facts existed in the case. It is, therefore, essential to know what the judge has said about his choice of the facts, for what he does has a meaning for us only when considered in relation to what he has said.⁵⁰

Hence, our third policy will be to presume immaterial any facts which the judge specifically states or impliedly treats as immaterial, e.g., by intentionally omitting from the opinion a fact which appears in the record. I believe that the results of following this tripartite strategy are extensionally equivalent to the results of obeying Goodhart's ten rules. Stated all at once, the strategy is:

Presume that all facts in the record are material except facts of person, time, place, kind, or amount (unless stated to be material) and facts specifically designated or treated as immaterial by the judge.

Notice that this strategy gives us an indication of the level of specificity which we might reasonably use in deciding whether or not two cases have the same facts. We don't want to be so specific in describing the facts of the case that our description designates particular persons, times, places, and so forth.

What we take to be the principle or ratio decidendi of a previous case is very important in our determination of whether or not the rule of that case should be applied in a pending case. Despite Goodhart's arguments, if a judge states a rule in his opinion and says, "This is the ratio decidendi of this decision," it is tempting and not entirely unjustifiable to take his word for it. So it is easy to see how we could have two candidates for the ratio decidendi of a case: one stated by the judge and one generated by Goodhart's proposals. Further, since it is going to be difficult to tell whether or not a judge "implicitly treats" a given fact as immaterial, the Goodhart program itself will be likely to produce a number of plausible candidates for ratio decidendi. Since there is likely to be more than one reasonable interpretation of the ratio decidendi for any case which has been decided, it is easy to see how the principle on which a case is decided could be a source of "slack" which would permit us to accept or reject the case

as a precedent for a later case.

A rule given or implied in the decision for a case is itself supported or justified by some authority. If the authorities used in an older case have been superseded, this would provide a reason for ignoring the older decision. On the other hand, it would be fallacious to reject a true principle or a just decision because we reject the authority offered for it. Here, again, is a locus of maneuver: a poor rule may be rejected on the grounds that the authority for it is void.

Any dicta found in an earlier decision may be cited as support for a present decision or it may be identified as dicta and ignored.

Finally, the subsequent history of a case may provide reasons for accepting the case as a precedent (it has been cited as a precedent for similar cases hundreds of times) or for rejecting it (the earlier decision was later overturned or was completely ignored by generations of judges and lawyers).

In the preceding paragraphs, I have tried to give an indication of the types of things which would count for or against the use of a particular case as a precedent for a later one. The one fact which stands out when we go over this issue is that often, within the system I have described, it will be possible to give good reasons for accepting a certain case as a precedent for another one and to give equally good reasons for rejecting it. Perhaps this is the case with Dixon v. Bell. It can be argued that the facts in the case are similar enough to the later cases that it should be taken as a precedent for them. In both, someone was injured by an article which was supplied by someone with whom the injured party had not contracted. But according to another

true description of the facts, Dixon v. Bell is not similar to, e.g., Langridge v. Levy. The second case concerned the fraudulent sale of a defective gun; in the first case, there was no fraud and no sale, and the gun was not defective.

As we have noted, the post-decision history of the case also provides no unambiguous reason for deciding to include or omit Dixon v. Bell. In favor of including it is the fact that it was cited by the winning counsel in Langridge v. Levy; further, Langridge v. Levy was itself cited as a precedent in each of the later cases we are interested in. But against this evidence is the fact that the judge in Langridge v. Levy did not mention Dixon v. Bell in his opinion, despite the fact that it had been cited as a precedent in the arguments.

No dicta seem to be at issue here, and the ratio decidendi will be viewed as either the same or different, depending on what we take to be the material facts of the cases. The authority for the principle will also depend on what we see as the principle which, in turn, depends on what facts we count as material. Thus, unfortunately, the system which I have described based on the various components of cases provides no unambiguous answer to the question of whether or not Dixon v. Bell should be considered a part of the data to be included in any acceptable model of \underline{L} . It appears that this question cannot be answered within the system which I have been talking about based on the various components of a court case.

On what basis, then, might this decision be made? If we are unable to decide on the basis of rules internal to the system, then the only alternative seems to be to step outside of the system. The legal

system is a means for achieving certain ends which we value. If there is evidence that deciding in one way would make the system better able to achieve those ends and deciding in the other way would interfere with that process, that would be a good reason for choosing the former decision.

What are the ends at which our legal system aims? It seems to me that there are two. The first goal of our legal system is to make life in society more predictable. We want to be able to predict, to some extent, what will happen to us. Will we be able to travel on the public highway without being robbed? Will we be able to rob travelers on the public highway without being punished? Questions of this sort need to be settled before people can reasonably live together, build, and prosper. In a lecture called "The Path of the Law," Oliver Wendell Holmes remarked that the proper object of study for law students is "the prophecies of what the courts will do."⁵¹ We are willing to hire legal counsel in order to obtain reasonably accurate forecasts of whether society will punish or protect us in our endeavors. This is manifestly valuable information, and changes to the legal system which improve our ability to forecast our fate are changes for the better.

The second goal of the legal system is to make life in society more equitable than it would otherwise be. Our lives could be perfectly predictable and still the legal system would be a failure if what we could predict was that, no matter what the circumstances, the worst possible thing would happen to us. At least a legal system should resolve intrasocietal conflicts in a way which is generally perceived to be fair. It would be wrong to assume that fairness is itself a simple

concept; it also seems to be developed from paradigm cases. Nevertheless, the L^- concept seems to be a derivative of it, and so we can talk about using the concept of fairness (at some stage of its development) to help decide whether or not a given case should be taken to be a paradigm for L^- .

Perhaps there are other goals or ends of the legal system which I have not listed. If so, these also should count in a judge's decision. Perhaps predictability is actually a part of fairness. In any event, since Dixon v. Bell apparently cannot be unambiguously demonstrated to be a precedent for L^- within the legal system, it appears that one must transcend that system and ask whether or not including Dixon v. Bell would make a greater contribution to the system's ability to achieve its ends than not doing so. Apparently Lord Parke, in his Longmeid v. Holliday decision, was of the opinion that it would. I shall have more to say about this type of decision in the final chapter.

The Conservative Nature of the Parke Model

Suppose we disregard the Dixon v. Bell case; is there then any reason for choosing the Parke model over the others? Without Dixon v. Bell, we have four models, each of which can account for the three remaining cases. For comparison, let us summarize the significant features of the competitors in the following table:

TABLE 1
COMPARISON OF SIGNIFICANT FEATURES
OF FOUR L MODELS

Model Name	Significant Features
Langridge	Dangerous commodity + knowledge of danger
Wright.	Weapon of a dangerous nature + fraud
Holliday.	Dangerous commodity + knowledge + fraud
Parke	<u>Either</u> article/service giving public duty of nonnegligence + negligence, <u>or</u> fraud with intent to cause harm

From this listing, we can see immediately that the Parke model describes a much broader class of cases than the others. The Wright model is the narrowest, describing only cases involving dangerous weapons fraudulently misrepresented. Could this broadness be a factor in favor of the Parke model? After all, in the earlier cases, the court was determined to limit or eliminate altogether the L cases. Lord Parke himself said in the earlier Langridge v. Levy case:

. . . we should pause before we made a precedent by our decision which would be an authority for an action against the vendors, even of such instruments and articles as are dangerous in themselves, at the suit of any person whomsoever into whose hands they might happen to pass, and who should be injured thereby.⁵²

Hence, we might suppose that Lord Parke would have selected the Wright model instead of the Parke model, if he were forced to choose among the four models being considered here.

But this would be wrong. In Langridge v. Levy, although the court rejected both the Langridge and Levy models of L^- , it did accept a model for a certain subclass of L^- which is identical to part 2 of the Parke model. Cases involving falsehoods told with the intention that someone should be induced by them to do acts likely to be injurious to themselves are cases in which the injured party has a right of legal action regardless of the presence or absence of a contract. Thus, the court was already committed to part 2 of the model. Indeed, as noted in the opinion of Langridge v. Levy, that part of the model derives from several earlier cases of which Pasley v. Freeman (1789) is the leading decision.⁵³

Part 2 of the Parke model identifies as L^- cases any cases picked out by the Wright and Holliday models. Nevertheless, the other models seem to point out an obligation to be especially careful not to harm others when we deal with things which are inherently dangerous. If we supply inherently dangerous articles to people without at least giving notice of their dangerous nature; and someone is injured as a result, he has a right to sue for the damage whether or not he has a contract with us.

The court was again already committed to the proposition that certain types of goods and services create a duty to the public of non-negligence. This duty holds independently of contractual obligations. Pippin v. Sheppard⁵⁴ (1822) and Gladwell v. Steggall⁵⁵ (1839) established that:

if an apothecary administered improper medicines to his patient, or a surgeon unskilfully treated him, and thereby

injured his health, he would be liable to the patient, even where the father or friend of the patient may have been the contracting party with the apothecary or surgeon; for though no such contract had been made, the apothecary, if he gave improper medicines, or the surgeon, if he took him as a patient and unskilfully treated him, would be liable to an action for misfeasance.⁵⁶

Thus, the Parke model is an attempt to gather up all the kinds of cases in which the court had already committed itself to liability independent of contractual obligations. There are two types of such cases: those in which someone injures another (1) through failure to fulfill a public duty of nonnegligence (part 1 of the Parke model models these) and (2) through fraud (part 2 models these cases). The competing models were not as good not only because they failed to account for Dixon v. Bell but also because they modeled only a few of the L⁻ cases which the court had recognized. The Parke model was broader than any of the others, but it was actually more conservative, since the other models required the elimination of a number of previously accepted L⁻ cases.

A Methodological Consideration

Before continuing, I must confess that my statement of the Parke model is patched together from the court report of Langridge v. Levy. Although part 2 of the model follows the original language quite closely, there is no clear statement of part 1. The notion of a public duty to avoid negligence which I used in part 1 is not explicitly included in the court report. Instead, there are four types of cases mentioned as examples of cases in which there is no fraud involved but in which an injured noncontracting party would have a right of action against the person who caused the loss. The four types of cases are those in

which a person is injured by:

1. An apothecary who administers improper medicines or a doctor who renders unskillful treatment
2. A stage coach proprietor or his servant who drives without due care
3. A mason who constructs a defective bridge or other work in a public road
4. Anyone who delivers to (or makes easily accessible to) another without notice an instrument in its nature dangerous⁵⁷

Part 1 of my expression of the Parke model is simply the narrowest description which I could discover which would include all four types of cases without listing them separately. Notice the similarity of what I was trying to do and the judicial process itself. Both involve model building by trying out various descriptions of the data in search of the "tightest" description which is not simply a concatenation of names or definite descriptions of the data.

My inspiration for part 1 comes from a comment by Lord Abinger in Winterbottom v. Wright:

Where a party becomes responsible to the public, by undertaking a public duty, he is liable, though the injury may have arisen from the negligence of his servant or agent. So, in cases of public nuisance, whether the act was done by the party as a servant, or in any other capacity, you are liable to an action at the suit of any person who suffers. Those, however are cases where the real ground of the liability is the public duty, or the commission of the public nuisance.⁵⁸

It seemed to me that this description fit the four types of non-fraud L cases listed by Lord Parke. Perhaps one might want to rename part 1

the "Abinger model" and parts 1 and 2 together the "Abinger-Parke model." For simplicity, I shall use the shorter label, but this tie with the earlier case lends support to my portrayal of the model as being basically conservative and not signaling a major change in the thinking of the court.

The Parke Model in America:
Thomas v. Winchester

A year after the Longmeid v. Holliday decision, the New York Court of Appeals applied the Parke model in one of its decisions. The case was Thomas v. Winchester.⁵⁹ On March 27, 1849, the plaintiff, Mary Ann Thomas, consulted her physician about an illness, and the doctor prescribed extract of dandelion, a mild medicine. Her husband, Samuel Thomas, went to the store of Dr. Alvin Foord, a druggist in the town of Cazenovia, and asked for some of the medicine. Foord sold Thomas a preparation which he took from a jar labeled "1 lb. dandelion, prepared by A. Gilbert, No. 108, John-street, N.Y. Jar 8 oz." Thomas returned home and administered a small dose to his wife. The result of the medication was not the expected relief from her illness, but "coldness of the surface and extremities, feebleness of circulation, spasms of the muscles, giddiness of the head, dilation of the pupils of the eyes, and derangement of the mind."⁶⁰ Instead of extract of dandelion, Dr. Foord had sold the Thomases a portion of extract of belladonna, a deadly poison!

The jar of belladonna was sold to Foord as dandelion by James S. Aspinwall, a New York druggist, who in turn had bought it from the defendant, Winchester. Winchester's firm made some of the preparations

which it sold and bought others, but, regardless of the source, everything was put up in jars with the "prepared by A. Gilbert" label. Gilbert began using the labels as an independent supplier of medical preparations, and, when Winchester took over the business and employed Gilbert as an assistant, Winchester decided that it would render the articles more salable to continue using the same labels. The belladonna which eventually was sold in the jar labeled "extract of dandelion" was one of those items purchased by Winchester from another manufacturer, and was apparently mislabeled by him or his assistant.

Winchester appealed an earlier decision against him on the grounds that his liability was limited to the original purchaser, Aspinwall. Winterbottom v. Wright and other cases were cited in support of this view.

The court affirmed the earlier decision for the plaintiff. The judgment was delivered by Chief Judge Ruggles. In Winterbottom v. Wright, he said, Wright's "duty to keep the coach in good condition was a duty to the postmaster general, with whom he had made his contract, and not a duty to the driver employed by the owners of the horses."⁶¹ But a poisonous drug is an article which gives its supplier a duty to the public to avoid negligence in supplying it. This is because a deadly poison which is negligently mislabeled puts human life in imminent danger. Thus, the case is an L case according to part 1 of the Parke model:

The defendant's duty arose out of the nature of his business and the danger to others incident to its mismanagement. Nothing but mischief like that which actually happened could have been expected from sending the poison falsely labeled into the market; and

the defendant is justly responsible for the probable consequences of the act. The duty of exercising caution in this respect did not arise out of the defendant's contract of sale to Aspinwall.⁶²

Inherently Dangerous Articles

In order to use the Parke model, we need to have some way of deciding which goods or services place upon us a duty to the public of nonnegligence when we provide them. Certainly the particular things listed in Longmeid v. Holliday (medicines, surgery, coach driving, public facilities, and naturally dangerous instruments) are supposed to be such items. But what is it about these particular items which creates the duty to the public? In Longmeid v. Holliday, recall that Lord Parke made a distinction between supplying an instrument in its nature dangerous and supplying one which might become dangerous at some future time through some presently unknown latent defect. We have a duty of nonnegligence to the public when supplying the former sort of article but not the latter. It is apparently this distinction to which Judge Ruggles was referring when he said:

In Longmeid v. Holliday (6 Law and Eq. Rep. 562,) the distinction is recognized between an act of negligence imminently dangerous to the lives of others, and one that is not so. In the former case, the party guilty of negligence is liable to the party injured, whether there be a contract between them or not; in the latter, the negligent party is liable only to the party with whom he contracted, and on the ground that negligence is a breach of the contract.⁶³

The act of supplying surgery or an instrument of a dangerous nature is a type of act which, if negligently performed, is imminently dangerous to the lives of others. The source of our public duty of nonnegligence is this imminent danger to the lives of others which our negligence

would cause. The action of Wright in supplying belladonna in a jar labeled "dandelion" was certainly imminently dangerous to the lives of others, so the court held that Wright was liable for Mrs. Thomas's injury despite the absence of a contract between them.

Following later jurists, let us use "inherently dangerous" to describe the goods and services which are such that, if we should be negligent in supplying them, our act is imminently dangerous to the lives of others. It is inherently dangerous articles whose supply gives us a public duty of nonnegligence.

In 1870 and 1873, there were cases in which the New York court found that the defendant was not liable for the plaintiff's injuries--again using the Parke model. The 1870 case, Loop v. Litchfield,⁶⁴ concerned a defective balance wheel for a circular saw, and the 1873 Losee v. Clute,⁶⁵ concerned a defective boiler which blew up. In deciding that these were not L⁻ cases, the court was convinced that the defendants had not acted negligently in supplying the articles in question. So there was no requirement for the court to rule on whether or not boilers or balance wheels are inherently dangerous or whether it is only defective boilers and balance wheels that are inherently dangerous. That is, the court did not have to decide the question of whether the "inherently dangerous" label should be applied to classes of items in their "normal" condition or only to classes of items in some "abnormal" condition which makes it more likely that they will do harm.

During the next three decades, English and American courts used the Parke model frequently and often found cases to be L⁻ cases. In these judgments (so far as part 1 of the model is concerned), the

court could not avoid deciding two things: (1) that the defendant had a duty to avoid acting negligently in supplying some item (because the item was of a type which, if negligently supplied, was imminently dangerous to the lives of others); and (2) that the defendant had, in fact, acted negligently. The list of things judged to be inherently dangerous grew to include some rather surprising articles, including: (1) a 90-foot-tall painter's scaffold,⁶⁶ (2) a boat painter's scaffold,⁶⁷ (3) a catwalk on a threshing machine over a revolving cylinder equipped with steel teeth,⁶⁸ (4) a bottle of aerated water,⁶⁹ and (5) a coffee urn.⁷⁰ These items were classed as being somehow relevantly similar to the defective gun of Langridge v. Levy and the mislabeled dandelion of Thomas v. Winchester; they were all inherently dangerous.

At the same time, the list of things judged not inherently dangerous but liable to become dangerous through some unknown latent defect was also growing. In addition to the mail coach of Winterbottom v. Wright, the oil lamp of Longmeid v. Holliday, the balance wheel of Loop v. Litchfield, and the boiler of Losee v. Clute, quite a few other kinds of items were judged not to be inherently dangerous. In 1903, Judge Sanborn (Eighth Circuit Court of Appeals, Minnesota) listed some of them in his opinion in Huset v. J. I. Case Threshing Machine Co.:

a defective chain furnished one to lead stone, . . .
 an improperly hung chandelier, . . . an attorney's
 certificate of title, . . . a defective valve in an
 oil car, . . . a porch on a hotel, . . . a defec-
 tive side saddle, . . . a defective cylinder in a
 threshing machine, . . . a defective wall which
 fell on a pedestrian, . . . a defective rope on a
 derrick, . . . a defective shelf for a workman to
 stand upon in placing ice in a box, . . . a defec-
 tive hoisting rope of an elevator, . . . a runaway
 horse, . . . a defective hood holding a heavy weight

in a drop press, . . . a defective bridge, . . . shelves in a dry goods store, whose fall injured a customer, . . . a staging erected by a contractor for the use of his employes, . . . defective wheels. 71

An Ambiguity in the Parke Model

Thus, by 1916, when the MacPherson v. Buick Motor Co. case came before the court, the situation was this. A number of L models had been proposed, the Parke model had been accepted, and it had been applied to quite a few specific cases. The cases in which it was applied made it clear that, in order to be judged inherently dangerous, the similarity between a particular item and the inherently dangerous articles in the original paradigms for the model (the loaded gun, the defective gun, the mislabeled belladonna) need not be especially strong. Unfortunately, the same is true of the judgment that a particular type of item is not inherently dangerous. It is not at all easy to see why the bottle of aerated water in Torgeson v. Schultz should be considered inherently dangerous, while the boiler of Losee v. Clute should be considered not inherently dangerous. Although the Parke model had seemed to be an admirably clear and correct model of the earlier cases, its application in certain later cases was less impressive. Although the Parke model was still being invoked in justification of judgments, it was virtually impossible to use the Parke model in advance of a decision to predict what a court's judgment would be. The Parke model could be used to "justify" the decision, regardless of whether the judgment was for the plaintiff or for the defendant. The court need only issue its decision that the article in question was or was not inherently dangerous. If an article had been ruled inherently dangerous, a later

court could rule that a similar article was not inherently dangerous by the simple expedient of pointing out the inevitable differences between the two items, although this sort of maneuver was generally avoided. The court's decisions about what was required by justice had begun to outrun the Parke model, even though lip service was still being paid to it.

The problem lay in an ambiguity in the model. As noted earlier, the model is ambiguous between the decision that guns are inherently dangerous and the decision that guns with a certain defect are inherently dangerous. There is a sense in which it is inherent in any gun to be capable of having or developing a defect in virtue of which it would be dangerous. There is some danger inherent in almost every kind of article, if certain things should happen to them, e.g., they are caused to become radioactive. But what range of counterfactual situations are we required to consider in determining whether or not a thing or type of thing is inherently dangerous according to the Parke model? A good guess seems to be to restrict the permissible counterfactual situations in which an item would be or become dangerous to those which could reasonably be expected to occur in the normal course of events. Belladonna mislabeled as dandelion can reasonably be expected to cause some mischief if it gets out into the marketplace, so it is inherently dangerous, and the act or acts which make it inherently dangerous are imminently dangerous to the lives of others. By the time of MacPherson v. Buick Motor Co., inherently dangerous articles were generally taken to be those things which in their normal operation are dangerous.⁷² On this view, it would not be reasonable to view a

stage coach with an undiscovered defect as inherently dangerous. Still, there is a problem in deciding which possible chains of events it is "reasonable" to expect; it might be equally reasonable to expect one future in which an article would be dangerous and another in which it would not. Further, it is notoriously difficult to give an adequate analysis of "normal" conditions.

MacPherson v. Buick Motor Co.

MacPherson v. Buick Motor Co. has become a famous case, I contend, because, in deciding it, the New York Court of Appeals adopted a major revision to part 1 of the Parke model which significantly reduced the ambiguity of the old model. In his dissent, Chief Judge Willard Bartlett argued the case from the perspective of the old model, so let us first examine Bartlett's dissenting opinion in order to highlight the contrast with the majority opinion based on the new model.

Recall that MacPherson v. Buick Motor Co. concerned a man who had bought a car from a retailer who was not the manufacturer of the car. The automobile had a faulty wheel which collapsed while he was driving at only eight miles per hour, and he was injured. He sued to recover damages from the manufacturer, even though he had no contract with the manufacturer.

Judge Bartlett's argument was true to the Parke model and previous decisions which that model was supposed to describe:

1. First,

. . . a contractor, manufacturer, vendor or furnisher of an article is not liable to third parties who have no contractual relations with him for negligence in the construction, manufacture or sale of such article,⁷³

the only exception (barring fraud) being:

cases in which the article sold was of such a character that danger to life or limb was involved in the ordinary use thereof; in other words, where the article sold was inherently dangerous.⁷⁴

2. In Winterbottom v. Wright, an English court ruled that a stage coach is not an inherently dangerous article.
3. In Thomas v. Winchester, an American court cited Winterbottom v. Wright as an authority for "the absence of any liability for negligence on the part of the original vendor of an ordinary carriage to any one except his immediate vendee."⁷⁵
4. Therefore, the American court is committed to the view that a stage coach or carriage is not an inherently dangerous article.
5. But--

in the case at bar the defective wheel on an automobile moving only eight miles an hour was not any more dangerous to the occupants of the car than a similarly defective wheel would be to the occupants of a carriage drawn by a horse at the same speed.⁷⁶
6. If an automobile with a defective wheel were inherently dangerous, it would be at least as dangerous as a carriage with a defective wheel.
7. Hence, the court is committed to the view that an automobile with a defective wheel is not inherently dangerous.
8. Therefore, given the Parke model of L (and since no fraud is involved), the manufacturer in this case is not liable for his negligence to anyone with whom he has no contract.

Some of the suppressed premises in the above argument are doubtful, but the argument serves as an illustration of the way in which the Parke model was employed. In addition to the above argument, Judge Bartlett mentioned the decision in Cadillac Motor Car Co. v. Johnson⁷⁷ handed down the previous year: "That case, like this, was an action by a subvendee against a manufacturer of automobiles for negligence in failing to discover that one of its wheels was defective, the court holding that such an action could not be maintained."⁷⁸

As Karl Llewellyn comments in The Common Law Tradition, "How anyone can argue that Bartlett is not here presenting 'good law,' I do not see. It is not only technically sound; it is closer to the 'feel' of both the authorities and the case in hand than is its opposite."⁷⁹

What, then, was the majority opinion, written by Judge Cardozo? It consisted of two parts: (1) an argument for a revised model and (2) an argument to show that the facts of the case fit the new model:

1. Part 1 of the Parke model specifies that the L⁻ cases are those involving articles which are inherently dangerous, "things which in their normal operation are implements of destruction,"⁸⁰ e.g., poisons and explosives.
2. But the court has determined that certain cases involving articles which are not implements of destruction in their normal operation are L⁻ cases:

A scaffold (Devlin v. Smith, supra) is not inherently a destructive instrument. It becomes destructive only if imperfectly constructed. A large coffee urn (Statler v. Ray Mfg. Co., supra) may have within itself, if negligently made, the potency of danger, yet no one thinks of it as an implement whose normal function is destruction.

What is true of the coffee urn is equally true of bottles of aerated water (Torgeson v. Schultz, 192 N.Y. 156).⁸¹

3. Therefore, the set of L^- cases is not limited to those cases involving "poisons, explosives, and things of like nature . . . which in their normal operation are implements of destruction,"⁸² i.e., part 1 of the Parke model (as presently interpreted) is inadequate.
4. A model which more adequately characterizes the cases which the court has found to be L^- cases is the following:

CARDOZO MODEL: L^- = the set of all cases in which EITHER:

- 1a. A person A is injured by some article D;
- b. D is supplied or made available to A by some other person B;
- c. A is not party to a contract requiring B to supply articles of type D to A;
- d. D is a thing of danger, i.e., either:
 - (1) D is inherently dangerous (an implement of destruction in its normal operation, or
 - (2) D's nature is "such that it is reasonably certain to place life and limb in peril when negligently made;"⁸³
- e. B knows that D will be used without new tests and that the users will be persons with whom B has no contract to supply D;
- f. B fails to take reasonable precautions to protect

the future users of D from injury (by giving warning of the inherent danger in D or by making D carefully); OR

2. Part 2 of the Parke model applies.

5. Therefore, in the absence of any more adequate model, the Cardozo model should be accepted.

I have drawn the above model from the following paragraph of the Cardozo text which I shall quote in full:

We hold, then, that the principle of Thomas v. Winchester is not limited to poisons, explosives, and things which in their normal operation are implements of destruction. If the nature of a thing is such that it is reasonably certain to place life and limb in peril when negligently made, it is a thing of danger. Its nature gives warning of the consequences to be expected. If to the element of danger there is added knowledge that the thing will be used by persons other than the purchaser, and used without new tests, then, irrespective of contract, the manufacturer of this thing of danger is under a duty to make it carefully. That is as far as we are required to go for the decision of this case. There must be knowledge of a danger, not merely possible, but probable. It is possible to use almost anything in a way that will make it dangerous if defective. That is not enough to charge the manufacturer with a duty independent of his contract. Whether a given thing is dangerous may be sometimes a question for the court and sometimes a question for the jury. There must also be knowledge that in the usual course of events, the danger will be shared by others than the buyer. Such knowledge may often be inferred from the nature of the transaction. But it is possible that even knowledge of the danger and of the use will not always be enough. The proximity or remoteness of the relation is a factor to be considered. We are dealing now with the liability of the manufacturer of the finished product who puts it on the market to be used without inspection by his customers. If he is negligent, where danger is to be foreseen, a liability will follow.⁸⁴

Thus, with the Cardozo model, the court announced its resolution of the

ambiguity in the Parke model. Guns in general as well as automobiles in general give their suppliers certain public duties of nonnegligence and care. Certain of those duties are peculiar to the manufacturers of guns because guns are dangerous in their normal operation. But other duties fall on the manufacturers of both types of articles because it is recognized that both guns and automobiles are liable to become dangerous if carelessly made.

Notice the following sentence from the material quoted above:

"That is as far as we are required to go for the decision of this case."⁸⁵ There are at least three senses in which the court might have gone further. First, the court decided that the duty involved falls upon the final manufacturer of the injury-causing article. It could have judged that the suppliers of component parts or the suppliers of the materials from which those parts were constructed were liable. Second, the court decided that the manufacturer has the duty in those cases in which he has knowledge of a probable danger which arises from his product. But the court could have judged that the duty is binding on a manufacturer when he has knowledge of merely possible danger caused by his product. Third, the court judged that manufacturers have a duty to take minimal actions to protect their customers, viz., by making the dangerous articles carefully. But it could have demanded much more sweeping and general protections for consumers, e.g., by prohibiting certain products or certain designs.

Judge Cardozo indicated his belief that it would not be proper for the court to go further in the second sense:

It is possible to use almost anything in a way that will make it dangerous if defective. That is not enough to charge the manufacturer with a duty independent of his contract.⁸⁶

But with respect to the first and third senses, he left open the possibility that later courts might go further. Stated differently, Judge Cardozo made use of a model which, he asserted, should not be altered to exclude the pending case, but which might be broadened to include a lot of other cases if later judges decided that was needed. Meanwhile, the arguments needed to motivate adoption of the narrow model were weaker than those which would have been needed to support a broader model.

Apparent counterexamples to the sufficiency of the Cardozo model can be outmaneuvered. For example, Loop v. Litchfield, the case involving the defective balance wheel for a circular saw was judged to be a non-L case by the court, but it seems to satisfy the Cardozo model for being an L case. It does not defeat the Cardozo model because the manufacturer discharged his duty by pointing out the defect to the buyer, who was happy to assume the risk, since the balance wheel was being offered at a bargain price. When the buyer leased the saw to a third party five years later, the manufacturer had no liability for the lessee's injuries. In Losee v. Clute, the case of the explosion of a steam boiler, the manufacturer knew that his test of the boiler would not be the final one, and "the finality of the test has a bearing on the measure of diligence owing to persons other than the purchaser."⁸⁷ Although these two cases do not exhaust the list of possible counterexamples, they illustrate the way in which such putative counterexamples

would be dealt with.

Judge Cardozo's opinion concludes by showing that, on the Cardozo model of L, MacPherson v. Buick Motor Co. is an L case:

6. Donald MacPherson was injured by his car which was constructed by the Buick Motor Company.
7. MacPherson did not purchase his car directly from the Buick Motor Company.
8. Further:

beyond all question, the nature of an automobile gives warning of probable danger if its construction is defective. This automobile was designed to go fifty miles an hour. Unless its wheels were sound and strong, injury is almost certain. It was as much a thing of danger as a defective engine for a railroad. The defendant knew the danger.⁸⁸

9. In addition:

it also knew that the car would be used by persons other than the buyer. This was apparent from its size; there were seats for three persons. It was apparent also from the fact that the buyer was a dealer in cars, who bought to resell.⁸⁹

10. Finally, the Buick Motor Company was negligent in that the defect in the wheel "could have been discovered by reasonable inspection, and that inspection was omitted."⁹⁰
11. Therefore, MacPherson v. Buick Motor Co. is an L case, i.e., MacPherson is entitled to recover damages from the Buick Motor Company.

Motivating a Model Change

In defending the decision from Winterbottom v. Wright and other precedents cited by Judge Bartlett, Judge Cardozo noted:

Precedents drawn from the days of travel by stage coach do not fit the conditions of travel to-day. The principle that the danger must be imminent does not change, but the things subject to the principle do change. They are whatever the needs of life in a developing civilization require them to be.⁹¹

But this rhetoric seems unnecessary. Judge Cardozo had no need to deny the validity of the old paradigm for non-L⁻ cases. As he noted in a later paragraph, "the defendant was not the manufacturer [of the defective mail coach]. He had merely made a contract to keep the van in repair."⁹² Further, it was never established that the defect which made the coach dangerous was one which Wright could reasonably have been expected to discover. Both the Parke model and the Cardozo model agree with the English court's judgment that Winterbottom v. Wright is not an L⁻ case. And the Cardozo model shows clearly in what respects MacPherson v. Buick Motor Co. differs so that it is an L⁻ case. So Judge Cardozo's oratory about "the needs of life in a developing civilization" seems to be beside the point.

But this is only true once the Cardozo model of L⁻ has been accepted. Before that, there is enough ambiguity in the old Parke model that Judge Bartlett's "precedents drawn from the days of travel by stage coach" can be decisive. At their lower speeds, automobiles are not any more to be considered implements of destruction in their normal operation than stage coaches. Judge Cardozo's "developing civilization" speech may plausibly be seen as an additional plea for changing the

model. Judge Cardozo's words echo those of Judge Coxe (Second New York Circuit Court of Appeals) who dissented to the majority finding for the defendant in Cadillac Motor Car Co. v. Johnson, a case whose facts are almost identical to those of MacPherson v. Buick Motor Co. Judge Coxe wrote:

The principles of law invoked by the defendant had their origin many years ago, when such a delicately organized machine as the modern automobile was unknown. Rules applicable to stage coaches and farm implements become archaic when applied to a machine which is capable of running with safety at the rate of 50 miles an hour. . . . "New occasions teach new duties;" situations never dreamed of 20 years ago are now of almost daily occurrence.⁹³

The message is that, although the Parke model was once adequate to model those L cases involving articles which generate public duties of nonnegligence, it had become inadequate. Modern society had developed new and unanticipated types of products to which the Parke model, as it had come to be understood by the courts, did not affix these public duties. But the courts were wrong not to associate these duties with such articles because, if negligently constructed, they can do great harm. Again quoting Judge Coxe:

If the law, as stated in the prevailing opinion is sustained, the owner of an automobile entirely free from fault may be injured for life by the collapse of a decayed wheel occurring a few months after its purchase, and be entirely without redress.

It is, I think, doubtful whether, in the circumstances disclosed, an action can be brought to a successful termination against the Pennsylvania company where the wheel was manufactured. If this be so, it follows that an injury may be occasioned by the grossest negligence and no one be legally responsible. Such a situation would, it seems to me, be a reproach to our jurisprudence.⁹⁴

In Judge Coxe's view, not only was the Parke model inadequate, its mechanical application by the courts was occasionally producing unjust decisions.

The Development of Systems of Classification

In the foregoing pages, we have seen the development of the Parke model for L from its beginnings in particular cases, through its adoption by the court in Longmeid v. Holliday, through a period of about sixty-five years during which the courts extended the model by applying it to new types of cases and eventually came to apply the model mechanically, to its final overthrow by the Cardozo model.⁹⁵ I suppose it may be debated whether or not the court's acceptance of the Cardozo model represents the "final overthrow" of the Parke model or merely a purification of it. In any event, it is apparent that the picture of classification systems we developed in earlier chapters is applicable to the abstract legal notion we have been considering in this chapter. The Parke model seems to have had a "life history" similar to that which we noticed earlier in connection with the feedback control concept. First, there must have been a preconceptual stage in which persons suffered injuries under circumstances in which it would have been just to hold another person with whom they had no contract liable for their loss. Next came the stage of reference sample acquisition, the early cases of Dixon v. Bell, Langridge v. Levy, and Winterbottom v. Wright. Third, a classificatory explanation of the judgments in the reference sample was developed. This step is marked by the Longmeid v. Holliday decision; the heart of the classificatory explanation was the

Parke model. Next, the \underline{L}^- concept was extended by deciding, a case at a time, using the Parke model, which types of cases were \underline{L}^- cases. As this process continued, it became possible in more and more cases to apply the model mechanically without "breaking new ground." But, as Roscoe Pound complained in his 1908 article on "Mechanical Jurisprudence," when law is applied mechanically, its "practical function of adjusting every-day relations so as to meet current ideas of fair play" is given up, and it becomes "a pseudo-science of technical rules existing for their own sake and subserving supposed ends of science, while defeating justice."⁹⁶ And this was the situation when the Parke model was superseded by the Cardozo model.

Pound recognized that there are strong motivations driving us to harden our dynamic schemes of legal classification into rigid "artificial" classification systems. To abnegate mechanical jurisprudence entirely is to opt for caprice and, perhaps, corruption. Mechanical jurisprudence is the basis for legal systems which are reasoned, uniform, and certain; but it does not ensure justice in individual cases. Indeed, given the diversity of the world, it practically guarantees injustice in some cases.

In jurisprudence, as in science, we must construct classificatory systems if we are to progress. The world is too full for us to treat everything as sui generis. But our progress toward whatever goals we seek will be limited unless we use our classificatory schemes as defeasible tools, subject to modification or defeat by concrete cases. In the final chapter, I shall draw some conclusions about the way in which this process of modification or defeat by concrete cases proceeds.

NOTES TO CHAPTER FOUR

¹ Karl N. Llewellyn, The Common Law Tradition: Deciding Appeals (Boston: Little, Brown and Company, 1960), p. 99.

² MacPherson v. Buick Motor Co. [1916], reported in J. Newton Fiero, Reports of Cases Decided in the Court of Appeals of the State of New York 27 (Albany: J. B. Lyon Company, 1916), pp. 382-401.

³ For a lively account of how the U.S. Supreme Court reached its decisions on the major cases heard during the 1969-73 terms, see Bob Woodward and Scott Armstrong, The Brethren: Inside the Supreme Court (New York: Simon and Schuster, 1979).

⁴ Benjamin N. Cardozo, The Nature of the Judicial Process (New Haven: Yale University Press, 1921), p. 9.

⁵ Karl N. Llewellyn, "The Status of the Rule of Judicial Precedent," University of Cincinnati Law Review 14 (1940):216.

⁶ MacPherson v. Buick Motor Co., pp. 384-85.

⁷ U.S. Department of Commerce, Interagency Task Force on Product Liability--Final Report (Washington, D.C.: U.S. Government Printing Office, 1978), pp. VI-30 - VI-31.

⁸ A. W. B. Simpson, A History of the Common Law of Contract: the Rise of the Action of Assumpsit (Oxford: Clarendon Press, 1975), pp. 218-20.

⁹ For example, suppose that Mrs. A buys a pressure canning kettle from Company B. Contrary to the manufacturer's instructions, Mrs. A

tries to put up a mixture of cooked pears and champagne. The pressure which results from cooking this already bubbly combination exceeds the structural limits of the vessel, and Mrs. A is injured in the ensuing explosion. Although there was a sales contract between Company B and Mrs. A, this would not be an L^+ case, I think, because Company B discharged its obligations by providing the warning which Mrs. A chose to ignore.

My own intuitions are not clear about the status of the following sort of case involving causal overdetermination. Again Mrs. A buys a pressure canning kettle from Company B, but this time the kettle explodes while Mrs. A is using it strictly in accordance with the manufacturer's directions. The explosion is the result of a defect in the kettle resulting from the manufacturer's negligence in constructing it. Fragments of the exploding kettle enter Mrs. A's brain and kill her. Now suppose that at the very instant the fragments of the exploding kettle penetrate her brain, she is already being killed by poisons from a bad batch of beans which she put up last year and also by a falling iron pot which has been shaken loose from its ceiling hook by a slight earth tremor. Each of the three lethal agents (kettle shards, bean poisons, falling pot) was sufficient to cause death at precisely the instant at which Mrs. A died. In the absence of the other two causes of death, I would have no trouble classifying the case as an L^+ case. But here, it is not so clear. Mrs. A's survivors are probably entitled to a new kettle, but are they also entitled to collect for the loss of Mrs. A's life? Company B's negligence provided a sufficient cause of

Mrs. A's demise, but in this case, she was doomed to die when she did regardless of anything Company B did or failed to do in the construction and sale of the kettle.

Even if we have trouble deciding the status of the second case, the first case demonstrates that not every case in which the injured person is party to a contract with the supplier of the article which caused the injury is an L^+ case.

¹⁰ Thomas M. Cooley, A Treatise on the Law of Torts or the Wrongs Which Arise Independently of Contract, 3d ed., 2 vols. (Chicago: Callaghan & Company, 1906), 2:1486.

¹¹ I shall follow current legal practice in using "the court" to designate any (and sometimes all) Anglo-American courts. If it is important to talk about a particular court, I shall use a more specific name or definite description, e.g., "the U.S. Supreme Court," "the New York Court of Appeals."

¹² Langridge v. Levy [1837], reported in R. Meeson and W. N. Welsby, Reports of Cases Argued and Determined in the Courts of Exchequer & Exchequer Chamber 2 (London: S. Sweet, 1838), pp. 519-31.

¹³ Ibid., p. 519.

¹⁴ Ibid., p. 523.

¹⁵ Dixon v. Bell [1816], reported in George Maule and William Selwyn, Reports of Cases Argued and Determined in The Court of King's Bench 5 (London: J. Butterworth and Son, 1823), p. 198.

¹⁶ Langridge v. Levy, p. 525.

¹⁷ Ibid., p. 522.

- 18 Ibid., p. 523.
- 19 Ibid., p. 525.
- 20 Ibid., p. 526.
- 21 Ibid.
- 22 Ibid., p. 527.
- 23 Ibid., p. 528.
- 24 Ibid., pp. 527-28.
- 25 Ibid., p. 528.
- 26 Ibid., p. 530.
- 27 Ibid., p. 531.
- 28 Winterbottom v. Wright [1816], reported in R. Meeson and W. N.

Welsby, Reports of Cases Argued and Determined in the Courts of Exchequer & Exchequer Chamber 10 (London: S. Sweet, 1843), pp. 109-116.

- 29 Ibid., p. 110.
- 30 Ibid., p. 111.
- 31 Ibid., p. 113.
- 32 Ibid., p. 112.
- 33 Ibid., p. 114.
- 34 Ibid., p. 115.
- 35 Longmeid v. Holliday [1851], reported in W. N. Welsby, et.

al., The English Reports: Exchequer Division, ed. Max. Robertson, 155 (Edinburgh: W. Green & Son, Limited, 1916), pp. 752-55.

- 36 Ibid., p. 753.
- 37 Ibid., pp. 753-54.
- 38 Ibid., p. 754.

- 39 Ibid.
- 40 Ibid.
- 41 Ibid., p. 755.
- 42 Ibid.
- 43 Ibid.
- 44 Ibid., pp. 754-55.
- 45 Llewellyn, Common Law, pp. 77-91.
- 46 Arthur L. Goodhart, "Determining the Ratio Decidendi of a Case," Yale Law Journal 40 (1930):183.
- 47 Ibid., p. 169.
- 48 Ibid., p. 178.
- 49 Ibid., p. 170.
- 50 Ibid., p. 169.
- 51 Oliver W. Holmes, "The Path of the Law," Harvard Law Review 10 (1897):461.
- 52 Langridge v. Levy, p. 530.
- 53 Ibid.; see also, Pasley v. Freeman [1789], reported in Charles Durford and Edward Hyde East, Term Reports in the Court of King's Bench 3 (London: J. Butterworth and Son, 1817), pp. 51-65.
- 54 Pippin v. Sheppard [1822], reported in George Price, Reports of Cases Argued and Determined in the Court of Exchequer and in the Exchequer Chamber 11 (London: Henry Butterworth, 1928), pp. 400-410.
- 55 Gladwell v. Steggall [1839], reported in John Scott, Cases in the Court of Common Pleas and Exchequer Chamber 8 (London: S. Sweet, 1841), pp. 60-70.

56 Longmeid v. Holliday, p. 755.

57 Ibid.

58 Winterbottom v. Wright, pp. 114-15.

59 Thomas v. Winchester [1852], reported in Henry R. Selden,

Reports of Cases Argued and Determined in the Court of Appeals of the State of New-York 2 (Albany: W. C. Little & Co., 1854), pp. 397-411.

60 Ibid., p. 405.

61 Ibid., p. 408.

62 Ibid., p. 410.

63 Ibid.

64 Loop v. Litchfield [1870], reported in Samuel Hand, Reports of

Cases Argued and Determined in the Court of Appeals of the State of New York 3 (New York: Banks & Brothers, 1874), pp. 351-61.

65 Losee v. Clute [1873], reported in H. E. Sickels, Reports of Cases Decided in the Commission of Appeals of the State of New York 6 (New York: Banks & Brothers, 1874), pp. 494-97.

66 Devlin v. Smith [1882], reported in H. E. Sickels, Reports of Cases Decided in the Court of Appeals of the State of New York 44 (Albany: Weed, Parsons & Co., 1883), pp. 470-79.

67 Heaven v. Pender [1883], reported in John E. Hall, et. al., The Law Reports: Cases Determined in the Queen's Bench Division and on Appeal therefrom in the Court of Appeal, and Decisions on Crown Cases Reserved, ed. J. R. Bulwer, 11 (London: William Clower and Sons, Limited, 1882-83), pp. 503-17.

⁶⁸ Huset v. J. I. Case Threshing Machine Co. [1903], reported in The Federal Reporter 120 (St. Paul: West Publishing Co., 1903), pp. 865-73.

⁶⁹ Torgeson v. Schultz [1908], reported in Alvah S. Newcomb, Reports of Cases Decided in the Court of Appeals of the State of New York 195 (Albany: J. B. Lyon Company, 1909), pp. 156-61.

⁷⁰ Statler v. Ray [1909], reported in J. Newton Fiero, Reports of Cases Decided in the Court of Appeals of the State of New York 195 (Albany: J. B. Lyon Company, 1909), pp. 478-85.

⁷¹ Huset v. J. I. Case Threshing Machine Co., pp. 868-69.

⁷² See, for example, the language in Huset v. J. I. Case Threshing Machine Co. in which the emphasis is placed on an article's being "dangerous to the lives and limbs of those who should undertake to use it for the purpose for which it was constructed" (p. 872). In the same case, it is stated that an exception to the general rule of nonliability occurs when the supplier's act of negligence "occurs in the preparation or sale of articles, like foods and poisons, whose primary use is to preserve, destroy, or affect life and health" (p. 865). We are to determine whether or not something is inherently dangerous by considering what is likely to happen during its normal use. This is corroborated by the list of paradigm examples of inherently dangerous articles given by Judge Ward (Second Court of Appeals, New York) in Cadillac Motor Car Co. v. Johnson [1915], reported in The Federal Reporter 221 (St. Paul: West Publishing Co., 1915): "Articles inherently dangerous, e.g., poison, dynamite, gunpowder, torpedos, bottles of water under gas pressure" (p. 802).

- 73 MacPherson v. Buick Motor Co., p. 397.
- 74 Ibid.
- 75 Ibid., p. 400.
- 76 Ibid.
- 77 Cadillac Motor Car Co. v. Johnson, pp. 801-4.
- 78 MacPherson v. Buick Motor Co., p. 401.
- 79 Llewellyn, Common Law, p. 434.
- 80 MacPherson v. Buick Motor Co., p. 389.
- 81 Ibid., p. 387.
- 82 Ibid., p. 389.
- 83 Ibid.
- 84 Ibid., pp. 389-90.
- 85 Ibid.
- 86 Ibid.
- 87 Ibid., p. 386.
- 88 Ibid., pp. 390-91.
- 89 Ibid., p. 391.
- 90 Ibid., p. 385.
- 91 Ibid., p. 391.
- 92 Ibid., p. 393.
- 93 Cadillac Motor Car Co. v. Johnson, p. 805.
- 94 Ibid., p. 803.
- 95 Edward H. Levi, in An Introduction to Legal Reasoning (Chicago: University of Chicago Press, 1949), pp. 8-27, discusses in considerably less detail the precedents leading to MacPherson v. Buick

Motor Co. in order to trace the breakdown of the "inherently dangerous rule." He discusses the notion of a life cycle for concepts which comprises three stages: creation, fixation, breakdown. He notes that a similar discussion of MacPherson v. Buick Motor Co. appears in Max Radin, "Case Law and Stare Decisis: Concerning Präjudizienrecht in Amerika," Columbia Law Review 33 (1933):199-212.

⁹⁶ Roscoe Pound, "Mechanical Jurisprudence," Columbia Law Review 8 (1908):606. See also arguments against legal positivism in Ronald Dworkin, Taking Rights Seriously (Cambridge: Harvard University Press, 1978) and Richard A. Wasserstrom, The Judicial Decision: Toward a Theory of Legal Justification (Stanford: Stanford University Press, 1961).

CHAPTER V

RATIONAL CONCEPTUAL ACTIVITY

In this final chapter, I propose to reap the fruits of our labors in the earlier chapters by constructing a causal theory of rational conceptual activity. The theory will be causal in the sense that the various stages of activity relating to a given concept will be linked into a causal chain extending back to a particular type of event. Despite the changes which a concept may undergo as we gain experience with it, it possesses a certain stability or continuity which results from the fact that the transformations are orderly. At any given time, a concept is causally linked to its earlier formulations. The theory will be a theory of rational conceptual activity in the following sense. We have some choice in the selection of our ends. We also seem to have some choice in the strategies we use for organizing our experiences of the world, our concepts. It seems to be a fact about the world that, viewed as means, certain ways of classifying give us a better chance of achieving our ends than others. The mountain climber who doesn't have the concepts of up and down or loose and secure will probably fail. What the rational conceptualizer wants to know is, given a certain set of ends and priorities, how can he develop or lay hold of concepts which will best serve his purposes. In contrast, an irrational conceptualizer would try to adopt concepts

randomly (or, if he were really sick, he might perversely seek those concepts which would frustrate his purposes) so that success in achieving his ends would be a matter of luck.

I take it that natural selection has not been kind to the irrational conceptualizer nor to the rational, but unsuccessful, conceptualizer. Thus, if we can give an account of the way in which concepts actually originate and develop in the public ken, we will have something of both factual and normative import. At least, we shall have provided a starting point for discussions of how to optimize our conceptual activity. The assumption that actual conceptual activity is at least an approximation of ideally rational conceptual activity justifies the "case study" approach which I have taken in this dissertation. Notice, however, that concepts which are the output at a given time of ideally rational concept development need not be "perfect" concepts, guaranteed to maximize our success in attaining our ends. The world changes, and our conceptual activity--however rational--is temporally limited; the venture carries the usual inductive risks.

There is a sense in which the rational conceptualizer's question is strictly the question of an individual: What method of reasoning should I employ in order to develop a set of concepts which will be optimally useful to me in pursuing my ends? But there is another sense in which the rational conceptualizer's question is a question asked on behalf of a whole community of language users. As generations of philosophers have observed, our concepts and our languages are parallel enterprises. When a person becomes a member of a linguistic community, he acquires a set of concepts which are shared by the members of that

community and which are the product of conceptual activity at the societal level. In so far as we share ends with other members of our linguistic communities, we share an interest in answering the societal version of the rational conceptualizer's question: How can we develop concepts which will help us maximize the attainment of our goals?¹

But the answer to the societal version of the rational conceptualizer's question will concern the activity of individuals. Hilary Putnam's image of the linguistic division of labor is apt. In so far as our society may be said to have the concept of, e.g., a feedback control system, we have it because of the activity of certain individual pioneers and experts who understand the fine points of control theory and to whose judgments we defer in hard cases. For the rest of us, we wield the concept with varying degrees of inexpertness, relying on the belief that expert testimony is available if we should really need it. The control engineers and appellate court judges of my case studies are or were experts for the concepts I have been concerned with. It is their activity we shall try to track in answering the societal version of the rational conceptualizer's question. Clearly, our theory must permit us to talk about the joint efforts of and interactions among members of a community.

There is yet another way in which the answer to the societal version of the rational conceptualizer's question concerns the activity of individuals. Although the fundamental concepts of a linguistic community come pre-packaged in the language, it seems that each individual must, in acquiring the language, recapitulate the process of developing the sameness relation for a given concept from a set of exemplars. We

do not "teach" concepts by opening up the student's brain and hard wiring them in, nor do we give "explications" in terms of necessary and sufficient conditions. Instead, as Thomas Kuhn has observed, we teach natural concepts by exposing the student to certain group-licensed exemplars. The exemplars may be material objects (samples of gold) or abstract ones (examples of musical counterpoint or of applications of a particular formula or process to a number of different situations). After being given the hypothesis that the samples exhibited bear some sameness relation to one another, the student really must work out for himself what that relation might be. Problems of indeterminacy prevent him from ever knowing with complete certainty that he has discovered the same sameness relation as the rest of the community. We can stipulate necessary and sufficient conditions for the application of certain (artificial) concepts, but such conditions themselves are given in terms of other concepts. These analysans concepts could also be artificial, but the stipulations must themselves be in terms of other concepts. If any of these be natural concepts, then even the artificial concepts are infected with indeterminacy.

When we attempt to teach a person a natural concept by means of exemplars, we hope that the student will be able to work his way from the exemplars and the sameness hypothesis to an understanding of the sameness relation which matches the community's present understanding (at some level of expertise) and which has the potential for future development. This is a tall order, and our chances of success rest on the hope that, in so far as the student's goals are similar to our own, he will reason from the exemplars to the sameness relation in the same

manner that we did, viz., in the way that I am calling "rational."

As we noted in chapter one, Kuhn has observed that scientists exhibit group-sanctioned behavior when applying complex scientific concepts, despite the absence of generally known rules of application. He has tried to explain this by noting the fact that new scientists are initiated into the linguistic community by exposure to group-licensed exemplars. But Kuhn's explanation is (admittedly) incomplete in the absence of an account of how the education-by-exemplar technique results in similar behavior. We can easily imagine possible worlds in which students are shown the same exemplars but come away with different understandings of the concept which was being exemplified. Anyone who has watched a young child trying to catch on to the concept of a poem, the concept of a joke, or the concept of a magic trick will know what I mean. If we suppose that scientists and other persons whose apparent agreement on concepts is to be explained are rational in the sense that they try to construct concepts which will help them attain their ends, and we assume that persons who are learning concepts have certain ends in common with other members of the linguistic community which possesses the concept, then I think the following causal theory of rational conceptual activity supplies part of what is needed to complete Kuhn's explanation.

Rational activity concerning a certain concept begins with an intentional event: we notice or give our attention to some property. Perhaps the property is one which makes things which have it dangerous to us or pleasurable; perhaps we see the property as one which would be useful to us in attaining something we desire or as one which could

frustrate our attempts to achieve some end.

We needn't set the stage for this initial event with William James's "blooming, buzzing confusion;" the person who first paid attention to the property of self-regulation probably had all the normal conceptual furniture of his culture.

For convenience, let us use "i" to designate a property which seems interesting to us. Such properties in general, let us call "I-properties," and let us call the items which have them "I-things." There seem to be three ways in which an I-property might come to our attention. First, we might notice that some actual item is an I-thing. The natural kinds--gold, tigers, and acids--provide examples of this sort of case. Second, we might concern ourselves with possibles which, if they existed, would be I-things. We might concern ourselves with possible I-things if such things helped us to explain something or to achieve some other end. Electrons and black holes are possible I-things, and so are self-washing clothes, perennial vegetable crops, and money trees. Of course, we also can fantasize possible I-things which fulfill no particular needs, e.g., talking squirrels, planets with green skies, and centaurs. Finally, in addition to concerning ourselves with actual I-things or possible nonactual I-things, we might want to consider some impossible I-things. By combining concepts, we can think about square circles and things that are red all over at the same time they are green all over.

Some of the things I have listed as possible I-things might, for one reason or another, actually be physically impossible. But I intend the division between possible and impossible I-things to apply

along lines of logical possibility and impossibility. Although logically impossible I-things are mere curios which we construct as exercises in the limits of logic, our notions of which things are physically possible or impossible change with the advance of knowledge.

Which properties a person or society will find interesting will depend on his or its interests. Let us call the set of ends relative to which some i is of interest "D". The mappings between D-ends and I-properties are not necessarily one-to-one. D may consist of a single member, giving value to many I-properties, or D may have many members, but pick out only one i. Thus, the initial event in the life of a concept occurs when a person who has certain ends D gives his attention to some property i, which he believes to have value for him in his caring for or valuing of D.

From the beginning of this dissertation, I have characterized concepts as capacities to react selectively to things in our environment, and it is admittedly a long way from being the subject of the sort of intentional event I have been describing to having a capacity to classify things in a certain way. The next step in the journey from giving one's attention to i to being able to react toward things as N's seems to be to form a model of the actual or possible I-thing or things which have been presented to us in imagination or perception. I shall henceforth omit discussion of concepts of logically impossible things or events, since these seem to be artificial capacities at best. Having the concept of a square circle seems to involve having the capacity to pick out squares and circles and to see why nothing can be both at once. Accounting for the concepts based on actual and possible

things will, as a by-product, take care of the concepts having to do with impossible things.

When we form a model of an actual or possible I-thing, we would like it to include things which will reliably indicate the presence of i. Thus we find the early control engineers searching for the earmarks of instability in governors and jurists trying to arrive at a "telling" description of the L⁻ cases. Our models of the I-things are hypotheses which we test and refine with experience, but they are also our attempts to place the I-things in the world of our understanding, and so they are essentially conservative. Our initial attempts to place something unfamiliar amount to regarding it as a variation of something familiar. In "placing" something in the world of our understanding, we attempt to explain it to ourselves, and so our model of a thing includes those properties which, given our other beliefs, seem relevant to our perceiving or imagining it as an I-thing.

In forming such models, we draw upon our theory or theories of the world in general. These include ideas of the world as a place in which certain kinds of regularities may be depended upon. So, for example, early control engineers tried to analyze control mechanisms by regarding them as little planetary systems or as pendulums. Jurists regarded L⁻ cases as simple cases of relations between noncontracting parties or between parties with an implied contract. These approaches are all attempts to regard the phenomena in question as falling under known regularities.

By appealing to such regularities or to analogies with other familiar items, we fashion a model of the I-thing or things we are

considering which includes certain features which we suppose to be relevant to our perception or imagination of it as an I-thing. This model is relative to our interests and purposes D in at least two ways. First, the properties which we will find interesting, the I-properties, are those which seem to have some value for us in the pursuit of our ends. Second, our model of the things which have those properties, the I-things, will be acceptable or unacceptable to us, depending on our purposes. For most of us, our model of a hawk, for example, does not include details about nesting habits, migration routes, flight characteristics, color and size of eggs, type and number of feathers, and so forth. Instead, most of us represent hawks to ourselves as large birds with fierce eyes and nasty looking beaks and talons. But, if we were being frustrated in our efforts to raise chickens by the attacks of a particular kind of hawk, our model of hawks might need to get a lot more detailed as we sought ways to protect our chickens from them. On the other hand, if we were artists, trying to sketch realistic pictures of hawks, again our model would grow a lot more complex as we noticed more and more details, and yet, the details of the artist's model would be different from the details of the poultry farmer's model.

We might represent a particular model of this type in the following way: $E_{\{e\}, i, D}$. "D" stands for the ends or purposes which motivate the model; "i" designates a property of some actual or possible thing which has value relative to D; "e" designates other properties which are included in the model; and "E" stands for the relationship in which e, i, and D stand to one another in the model. It is difficult to say much more about E, but the familiar relationship of explanation

seems to come close to the sort of relation which E is. Viewed in this way, the presence of i in the actual or possible exemplar is the explanandum, and the presence of the e's constitutes the explanans.

Of course, the expression given above is a simplification. The more interesting concepts have exemplars which have a large number of I-properties. For some purposes, an automobile, for example, has one set of I-properties; for other purposes, it has others. Different sets of I-properties require different sets of explainers. Thus, the models which we have for our familiar concepts are much more complex than is implied by the simple three-part relation we have been considering. If we gather up all the ends we could have relative to which a feedback control system might have interesting properties, and we add in all of those interesting properties, and, in addition, we add all of the features which we might take to account for those properties; the result illustrates something of the complexity of our mature conceptual models.

But we have been talking about concepts at the very start of their existence. Somebody directs his attention to a property which is or might be had by something, and he formulates a model of the things which have that property in his experience or imagination. The exact shape which the model takes will depend on his interests, the level of sophistication of his theories about the way the world is, and his ingenuity. The model then serves as a sort of template which he applies to future experiences. He can certainly use it to reidentify the original exemplars as things which have the property which interests him, and so the possession of such a model seems to be all that need be added to our original capacities in order for us to be able to act selectively

toward things in our environment, i.e., for us to have a concept. But concepts whose application is limited to the exemplars from which they were originally drawn are not as useful to us as more general concepts which can be applied to particular things not yet in our experience or imagination.

We extend our proto-concepts by taking risks. If we find some thing which has each of the e's of our model, then, even though it is not one of the exemplars from which the model was drawn, we have a strong justification for expecting the new specimen to have i as well. There is some risk involved here, since we have no guarantee that our model really captures all of the features of the exemplars which are relevant to their having i, and there might be previously unknown things which have all of the e's but lack i. But this risk is balanced by the value of having a new supply of I-things if we are right.

But it is more likely that we will find nothing which our "template" fits exactly. Instead, we will find things which come close to fitting, and we must judge whether or not to regard these as being the same as the original exemplars. Indeed, there may be a continuum of items which fit the template with varying degrees of looseness. This sort of situation is a lot more risky than the first, since we must try to guess which of the ill-fitting items will have i or will have an i-like property in a degree sufficient for our purposes D. We learn from both our mistakes and our successes, and we improve in our ability to get the results we want. As we work, we develop an increasingly complex model.

A part of such a mature model apparently deals with near misses.

When all the indications which we know about lead us to expect a thing to have i sufficiently for our purposes D, and then it turns out not to, we have a superior model if we can explain why the item under consideration failed to have i. If we value gold for its resistance to rust, e.g., and some mineral which seems, according to our model, to be gold rusts, we have a better model if we can explain where we went wrong than if we can't.

But sometimes, what goes wrong is that the model is defective. Perhaps the original set of exemplars was not homogeneous: we were dealing with more than one kind of thing. Any attempt to model such a set of exemplars as one kind of thing is doomed to failure. A further possibility is that some of the background beliefs on which the model is based are in error. Our hypotheses about what e's account for the presence of i in the exemplars may then be false. Finally, the original set of exemplars may be defective in that it contains noncentral or borderline cases: we may have learned what bread is solely from samples of Wonderbread. Later, we may discover that Wonderbread is a marginal example of what bread is.

The immediate "point" of the conceptual activity I have been describing is apparently to achieve our ends by improving our access to certain I-properties. Any conceptual activity which further improves access to the I-properties is what I have called "optimizing" activity. The whole inductive process of learning what kinds of judgments will yield acceptable results aims at optimizing our ability to take advantage of the available I-properties. Once we have noticed something in our experience or imagination with some property of value to us, i,

success and failure has been defined for us in regard to one aspect of our conceptual activity. Success is maximizing our access to i, and failure is losing our grasp on i.

Virtually all of our conceptual activities are optimizing activities in this sense. I have already given illustrations of a number of these activities: construction of a model (e.g., the activity which resulted in the Parke model of \underline{L}^- cases or Maxwell's model of the centrifugal governor); applying the model to non-exemplars (e.g., the use of the Parke model in Cadillac Motor Car Co. v. Johnson and applications of the Maxwell model to novel types of governors); fine-tuning the model to accomodate differing purposes and circumstances (e.g., attempts to refine the notion of "inherently dangerous" in the Parke model and Hazen's extension of the Minorsky model of rotational servo-mechanisms to include dynamically similar systems including translation or electric current); accounting for failures (e.g., the judicial technique of insulating one's model from a bad result by pointing out disanalogies to show that the model did not apply or Maxwell's plea that he could get more acceptable results from his model if he knew how to find the roots of the higher order equations); and correcting a model for a nonhomogeneous or noncentral exemplar set (e.g., Parke's judgment that Dixon v. Bell was a central \underline{L}^- case and Maxwell's insistence on dividing the set of speed regulators into governors and moderators). There is another type of optimizing activity which is available to us in some cases. We may be able to maximize our access to i by actually changing what it is to be a central case of an I-thing.

Science is not solely descriptive. We are not only interested

in discovering the best way of understanding what the world is like; there is also a normative aspect to science. In so far as we are able to change the world, we want to know which of various possible future worlds we should try to produce, given our ends. If we believe that something which conforms more closely to a given model would be more useful to us than what is presently available, in many cases, we can produce items which do conform to our model. We can cause possible I-things to be actual. For example, modern hybrid corn plants bear only the slightest resemblance to their wild ancestors. The same is true of livestock, houses, and various types of machinery. In time, our efforts actually change what it is to be a central case of a certain kind of thing, not only in the sense of being closer to our model of that type of thing but also in the sense of being a typical example.

Stated in general terms, the process I am describing is the following. Given a particular set of I-things, P, we form a model, $E\{e\}, i, D = M$. M explains the I-things by describing them in a way which may only be roughly correct. We construct or breed a new set of I-things, P', which are closer to the description given by M, and the transaction gives us better access to the I-properties we are interested in. P' then becomes the new exemplar set for us and replaces P. The elements of P are regarded as less central cases of what is in P' and is modeled by M.

An especially fascinating example of this process is to be found in the history of architecture. In The Classical Language of Architecture, John Summerson discusses the origin of the "Five Orders of Architecture," the five types of columns derived from classical Roman

architecture. The general practice of using columns surmounted by elaborate entablature began with wooden buildings. As stone construction techniques became available and the old timber constructions began to decay, the more important buildings were replaced with stone replicas. These, in turn, became paradigms for other buildings. Different regions developed different styles. In the first century A.D., Vitruvius, a Roman architect, wrote a ten-volume essay on architectural practice, De Architectura. In this work, Vitruvius described the Doric, Ionic, and Corinthian orders and included a few notes on the Tuscan order. He apparently did not intend to give an exhaustive taxonomy of all the different types of column styles then in use, but merely to mention a few of them and their origins. Fourteen centuries later, the Florentine architect Leon Battista Alberti, relying on Vitruvius and his own observations of the Roman remains, described the same four orders and added a fifth from his observations, the Composite, a combination of the Corinthian and the Ionic. In 1537, Sebastiano Serlio wrote a book which began with an engraving in which all five orders were shown standing side by side, arranged according to their proportions, from the squat Tuscan to the slim and lofty Composite. The text portrays these five orders as a complete, canonical system, reifying and freezing for all time the forms of harmony and perfection in architecture. Serlio's innovative manner of presenting his material was wildly successful; even today, architects pay at least token obeisance to the five orders.²

Now the interesting thing about all this is that the five orders which had so much to do with the way our large public buildings have been built for centuries are only partly the result of observations of

what the Romans built. They are actually a model which Vitruvius, Alberti, and Serlio constructed by fitting together selected features of the Roman works. As Summerson says:

Serlio puts [the orders] before us with a tremendous air of authority giving dimensions for each part as if to settle the profiles and proportions once and for all. But in fact, Serlio's orders, while obviously reflecting Vitruvius to some extent, are also based on his own observation of ancient monuments and thus, by a process of personal selection, to quite a considerable degree his own invention. It could hardly be other wise. Vitruvius' descriptions have gaps in them and these can only be filled from knowledge of surviving Roman monuments themselves. The orders as exemplified in these monuments vary considerably from one to the other so it is open to anybody to abstract what he considers the best features of each in order to set out what he considers his ideal Corinthian, Ionic or whatever it is.³

And so Serlio managed to convince the world that these models he had in his mind were perfect and that we should design and build according to them. Today, the buildings which were constructed according to Serlio's models are the central cases from which the more adventurous architecture somewhat self-consciously deviates. Although Summerson is able to list a few cases in which architects elected to ignore Serlio's models and to build literal copies of specific Roman buildings,⁴ for the most part, Serlio's models were accepted as true or proper descriptions of the proper use of columns in architecture, and the central cases originally provided by the Roman remains were eventually supplanted by monuments constructed according to Serlio's models. As more and more of these buildings were constructed, the "typical" building with columns came to be one which had columns resembling Serlio's designs. Although these columns had a causal ancestry which could be traced back to the Roman stonework and even back to the original wooden buildings, Serlio's

pictures of the orders and the buildings inspired by his drawings became the exemplars which were shown to aspiring architects.

In addition to the optimizing activities, some of our conceptual activity is aimed at standardization. As we noticed earlier, a person who is learning a language or one who is learning concepts which are new to him does so by studying exemplars licensed by the linguistic community. In reasoning from the exemplars to the capacities which are exemplified, he must repeat for himself the discovery of the concept. If he wants to acquire the same concept as the one in use in the particular community of speakers, he must attempt to reason from the exemplars in a manner which will parallel the development of the concept in that community. I claim that the salient approach for the learner to take is to assume that the community at large has tried to achieve a capacity for selective behavior which maximizes its access to the I-properties involved. If so, then the appropriate stance for the person who is interested in standardizing his understanding of a community's concepts is optimization. Of course, the opportunity to collect experience relevant to optimizing a given concept is much greater for a community than for an individual, but the optimizing stance will at least put the learner on the right track. His capacities for selective behavior will approach those of the community as he gains experience, and there are other cues available to the learner of a concept which are unavailable to the original discoverer of one. Whereas much of a community's experience with a particular concept may be repetitive or wasteful, that of the learner who is merely "catching up" with the community can be directed. By giving the learner definitions of a community's concepts,

even if those "definitions" cannot be sets of necessary and sufficient conditions, we direct his attention and focus his efforts. In this way, we can use a learner's mastery of simpler concepts to help shorten his discovery time for the more complex ones. For example, texts on feedback control systems not only give paradigm examples of such systems, they also attempt to give definitions of their subject. Although the definitions cannot capture the complexity and vital variability of the concepts, both definitions and exemplars are important to the standardization process. The definitions help the student think of the right beginning model or proto-model; the exemplars set him on the road to developing that proto-model into the more sophisticated model possessed by the community's current experts on the subject.

Since rational conceptual activity is aimed at optimization, we are interested in making our models more and more useful in maximizing our access to the I-properties we value. Sometimes this may be achieved by adding details to the models to achieve greater realism, but we must balance the gains made by making our models more realistic against our own finite abilities to use such models. In chapter three, we examined a number of arguments designed to justify various degrees of unrealism in the models which were being considered. Advances in mensuration, mathematics, and information processing permit the experts who have access to these aids to work with increasingly complex models. The controlling factor is whether or not a move to a more or a less idealized model will give us better access to properties which have value for us relative to our ends.

Close scrutiny of the legal cases detailed in chapter four reveals only four types of arguments being used by either lawyers or judges. Each of these types of arguments may be seen as optimizing arguments.

The most frequently found of these four types of arguments is the argument based on a claimed analogy or disanalogy with a previously decided case. A lawyer typically appears in court armed with a supply of previous cases which bear similarities to the pending case. If his client is, e.g., the plaintiff, the lawyer tries to show that the pending case has a very strong similarity to the previous cases which were decided for the plaintiff, but only a weak and superficial similarity to those which were decided for the defendant. The counsel for the defense tries to use the same strategy to show that the pending case is most similar to earlier cases which were decided in favor of the defendant. Arguments of this sort appear in each of the cases we considered in chapter four.

In law, as in science or engineering, judging a case is a matter of exercising a capacity to behave selectively toward certain kinds of things. It is a matter of applying a model by judging that the case is sufficiently like (or unlike) the model, given the ends and purposes which set the context of judgment. In chapter four, we speculated that the ends of our legal system are predictability and justice. In science and engineering, our goals seem to be explanation and control. In deciding that a particular case is (or is not) an \underline{L} case, a judge further elaborates a model of the \underline{L} cases in a way which, he hopes, will

tend to maximize our satisfaction of those ends. Since, we may presume, an inconsistent model which both includes and excludes the same case will not advance either of our legal goals, if a judge becomes convinced that a pending case is similar to previous cases, consistency dictates a similar judgment.

So the conclusion of this first type of argument is that a particular case A is of a particular type N. This conclusion is justified by means of the familiar "negligible discrepancy" argument which we first identified in chapter three. By showing the strength of the analogy between the pending case and a model of certain previous cases, jurists attempt to show that any discrepancy between them is negligible. In the legal cases, as in the engineering cases, this type of argument does double duty, being used first to justify the use of an idealized model to represent members of a given exemplar set and, second, to justify the application of the model to nonexemplar cases. In both uses, the argument serves as an optimizing argument, since the goal of increased access to I-properties is what legitimizes the argument and causes the premises to be reasons for accepting the conclusion.

A second kind of optimizing argument to be found in the court cases we examined concerns the effect which a particular decision would have on public policy. Again and again, we found judges making a certain decision on the grounds that the opposite decision would have disastrous effects on society. To cite just one example, Lord Abinger argued in Winterbottom v. Wright against giving noncontracting

parties a right of action against negligent suppliers of merchandise by giving just this sort of argument. He said:

There is no privity of contract between these parties; and if the plaintiff can sue, every passenger, or even any person passing along the road, who was injured by the upsetting of the coach, might bring a similar action. Unless we confine the operation of such contracts as this to the parties who entered into them, the most absurd and outrageous consequences, to which I can see no limit would ensue.⁵

This is clearly an optimizing argument to the effect that there is no model of L^- cases which could lead to satisfaction of our ends of predictability and justice. But even this argument assumes some such model, for the purpose of identifying the kinds of case being discussed.

A second example of this type of argument is Judge Coxe's argument in Cadillac Motor Car Co. v. Johnson:

It is, I think, doubtful whether, in the circumstances disclosed, an action can be brought to a successful termination against the Pennsylvania company where the wheel was manufactured. If this be so, it follows that an injury may be occasioned by the grossest negligence and no one be legally responsible. Such a situation would, it seems to me, be a reproach to our jurisprudence.⁶

Here the argument is that the model of L^- cases which was then in use actually frustrated the goal of justice. Both arguments are optimizing arguments, since both obtain their persuasive force by pointing out the expected consequences of an action, although one is employed in constructing or selecting a model and the other concerns applying a model.

A third type of optimizing argument found in the court cases we studied is the attempt to support a particular model by appealing

to legal rules or principles. I gave examples of such principles in the discussion of Langridge v. Levy. They seem to be attempts to codify our findings about what properties are relevant to what other properties. If we look over a set of exemplars, there will always be a great number of true descriptions which the various members of the set have in common. But only certain of those descriptions are relevant to the properties which interest us, the I-properties. The principles seem to serve in the model construction process as filters to suggest to us which elements to include in our models. Unfortunately, as we saw in Langridge v. Levy, different principles suggest competing models. Nevertheless, such principles seem to assist in the process of discovering the most fruitful models by providing a heuristic for selecting models which are the most likely to develop optimally.

The legal rules and principles form a part of our stock of beliefs which provide a background for our model making activity. In science, engineering, and other human endeavors, laws of nature and rules of thumb are the analogues of these legal rules and principles.

The fourth and final type of argument visible in the court cases tries to show that the exemplar set needs to be reassessed in the light of judgments which have been made using the model based on the original exemplar set. This is the kind of argument made by Judge Cardozo to motivate the change from the old Parke model to the Cardozo model. As we apply a model to various nonexemplars, we build up a stock of cases which we are committed to saying are in some sense the "same" as the original exemplars. Each new case represents a bit of conceptual risk

taking; we make the hypothesis that it will be more productive to treat the new case as though our "template" fits even though it doesn't exactly, than to treat the case differently. If a number of these close fits give us access to the same I-properties that we noticed in the original exemplar set, then it is efficient for us to broaden the exemplar set to include these new cases, since we are no longer really taking a big risk when we make judgments about them. This will have the effect of increasing the tolerances on our "template," and this will be an optimizing move for us, since the "template" will fit a lot more cases with no greater risk than we had before the change.

This argument is a member of a class of arguments which motivate us to make corrections (or at least tend to justify the making of such corrections) in an exemplar set or in a set of background beliefs, based on the success or failure of a given model. These arguments are, thus, a sort of feedback mechanism for improving our concepts, based on their performance.

In summary, we have found that arguments employed in our conceptual activity perform basically one of two functions. One set of arguments justify the decisions we make in constructing, selecting, and applying models. The specific arguments of this type which we have found are the: (1) negligible discrepancy argument; (2) possibility of improvement argument; (3) limiting case argument; (4) appeal to expected consequences; and (5) appeal to laws, principles, or rules of thumb. Although the limiting case argument may have applications only in model construction and selection, the others--it appears--may be

used alone or in combination in model construction, selection, and application. Finally, it should be noted, that I have no argument for the exhaustiveness of this list of arguments. Other types of arguments may be used for these same purposes; these are the ones revealed in the cases I have studied.

The other function performed by arguments in our conceptual activity is the feedback function mentioned earlier. In this role, they serve to justify alterations in background beliefs or exemplar sets, given our success or failure in using a concept in a given form. The example we discussed earlier, Judge Cardozo's argument, we might call a "conservation of energy" argument. Rather than use a model based on a very narrow exemplar set and continue to deal with what the model tells us are risk-laden poor fits, we might as well acknowledge that our success rate in applying the concept to the ill-fitting nonexemplars has been excellent (assuming that it has). We can save our energy for dealing with the really risky cases by simply broadening the exemplar set (and the model) to include some of the previously decided cases.

Other arguments involved in the feedback function justify our attempts to correct a concept which is not producing satisfactory results. This might be accomplished by altering our background beliefs, by dividing the exemplar set, or by replacing some of the exemplars with more central cases. One might object that it is not proper to speak of "arguments" which justify these alterations, since it is clearly our success or failure which justifies them, but surely some

such argument as the following is involved:

1. When person S replaced item A with B in his exemplar set for concept C, improved access to C's I-properties for S resulted.
2. Any change which improves access to I-properties is justified.
3. Therefore, S was justified in replacing A with B in his exemplar set for concept C.

Certainly it would be interesting to know how to discover before the fact exactly which alterations would be optimal. Perhaps a "logic of discovery" will someday be discovered for inventing scientific hypotheses which have a high probability of confirmation. That sort of a logic would supply what is needed to find optimal alterations without resorting to simple trial and error testing.

With this we come to the end of our investigation of paradigm-based conceptual activity. The end product of this investigation is a theory of rational conceptual activity which may be seen as a continuation of Thomas Kuhn's remarks about the role of paradigms in science. I have diagrammed the theory in figure 2 on the following page. It is a theory which makes much of the importance of models, explanation, human ends, and the manner in which the drive to achieve those ends manifests itself in our conceptual activity. It is offered here as a theory which accounts for the data presented in chapters three and four. I believe that it will serve equally well to describe conceptual activity in any field.⁷

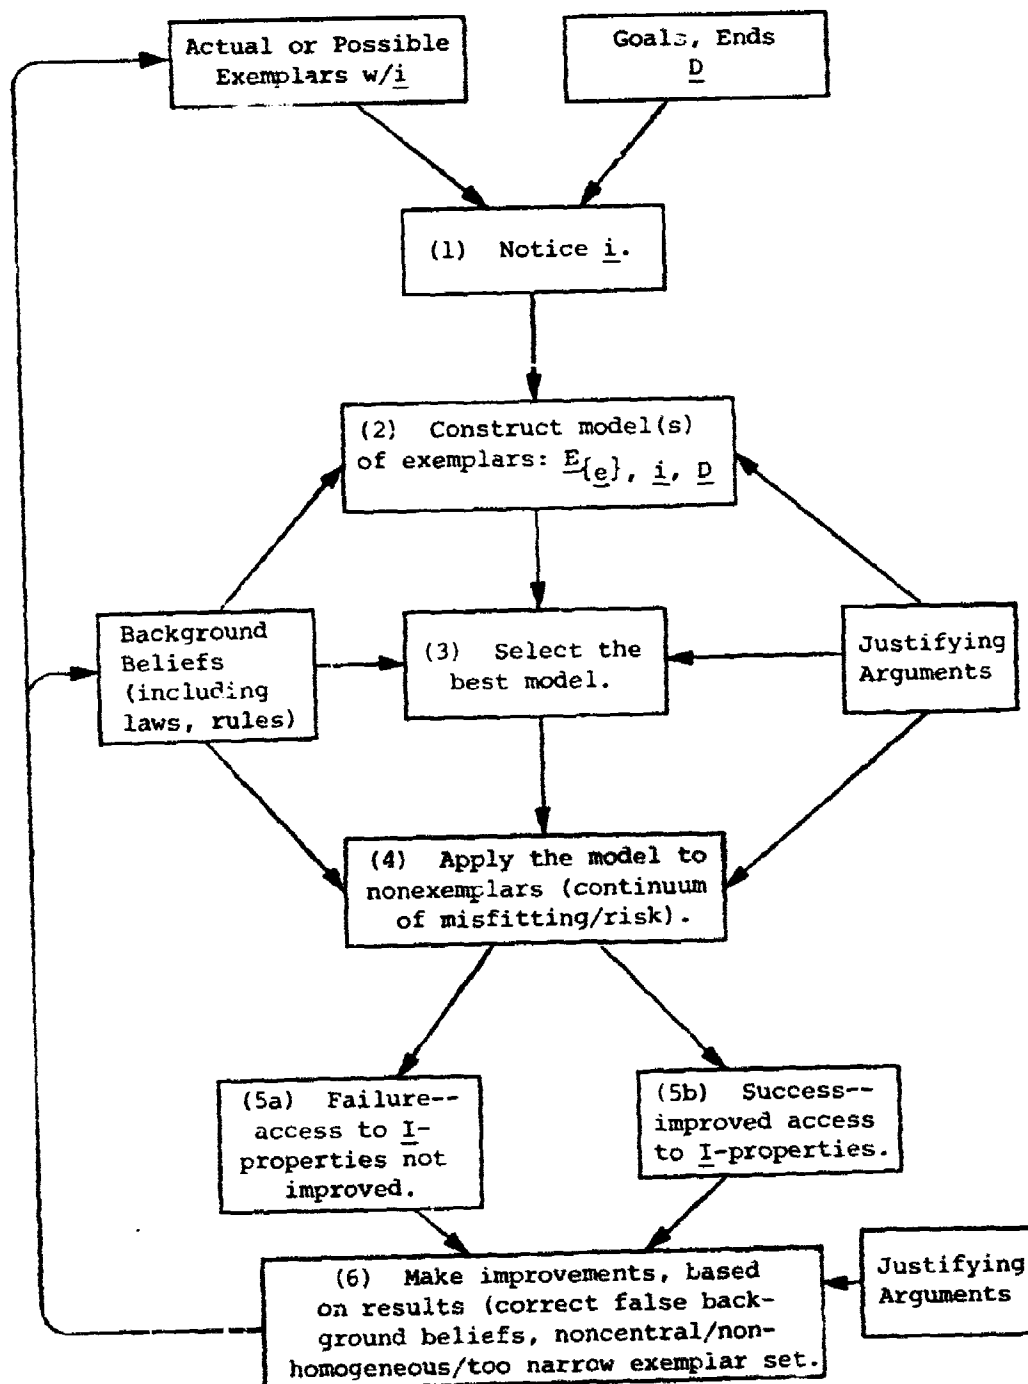


Fig. 2. Block Diagram of Conceptual Activity

NOTES TO CHAPTER FIVE

¹ For more on this question and on the way in which our concepts themselves influence our ends and expectations, see Stephen Toulmin, Human Understanding (Princeton: Princeton University Press, 1972), especially pp. 145-55.

² John Summerson, The Classical Language of Architecture (Cambridge: The M.I.T. Press, 1963), pp. 9-11.

³ Ibid., p. 11.

⁴ Ibid.

⁵ Winterbottom v. Wright [1816], reported in R. Meeson and W. N. Welsby, Reports of Cases Argued and Determined in the Courts of Exchequer & Exchequer Chamber 10 (London: S. Sweet, 1843), p. 114.

⁷ Recently, cognitive and developmental psychologists have reported empirical evidence which seems to confirm the thesis that we make judgments about whether or not a given object falls under a certain concept by determining the similarity between the object at hand and a stored model (they say "prototype") of a paradigmatic instance of the concept. See, for example, Eleanor Rosch, "On the Internal Structure of Perceptual and Semantic Categories," in Cognitive Development and the Acquisition of Language, ed. Timothy E. Moore (New York:

Academic Press, Inc., 1973), pp. 111-144; Eleanor Rosch, "Cognitive Representations of Semantic Categories," Journal of Experimental Psychology: General 104 (1975):192-233; Eleanor Rosch, "Classification of Real-world Objects: Origins and Representations in Cognition," in Thinking: Readings in Cognitive Science, eds. P. N. Johnson-Laird and P. C. Wason (Cambridge: Cambridge University Press, 1977), pp. 212-22; and Amos Tversky, "Features of Similarity," Psychological Review 84 (1977):327-52.

BIBLIOGRAPHY

Articles

- Adler, David. "Amorphous-Semiconductor Devices." Scientific American, May 1977, pp. 36-48.
- A. I. E. E. Committee Report. "Proposed Symbols and Terms for Feedback Control Systems." Electrical Engineering 70 (1951):905-909.
- Bellman, Richard. "Control Theory." Scientific American, September 1964, pp. 186-200.
- Black, Harold S. "Stabilized Feedback Amplifiers." Bell System Technical Journal 13 (1934):1-18.
- Bode, Hendrik W. "Relations between Attenuation and Phase in Feedback Amplifier Design." Bell System Technical Journal 19 (1940): 421-54.
- Callender, A.; Hartree, D.; and Porter, A. "Time Lag in a Control System." Philosophical Transactions of the Royal Society of London A235 (1936):415-44.
- Donnellan, Keith S. "Reference and Definite Descriptions." The Philosophical Review 75 (1966):281-304.
- Encyclopaedia Britannica. 15th ed. (1979), S.v. "Automation." By Morris Tanenbaum.
- Encyclopaedia Britannica. 15th ed. (1979), S.v. "Control Systems." By Theodore J. Williams.
- Haavind, Robert. "Glassy Semiconductors--a Missed Bet?" High Technology, February 1980, pp. 34-41.
- Hazen, Harold L. "Theory of Servo-Mechanisms." Journal of the Franklin Institute 218 (1934):279-331.
- Ivanoff, A. "Theoretical Foundations of the Automatic Regulation of Temperature." Journal of the Institute of Fuel, London 7 (1934):117-38.

- Kripke, Saul A. "Naming and Necessity." In Semantics of Natural Language, pp. 253-355. Edited by Donald Davidson and Gilbert Harman. Dordrecht-Holland: D. Reidel Publishing Company, 1972.
- Llewellyn, Karl N. "The Status of the Rule of Judicial Precedent." University of Cincinnati Law Review 14 (1940):208-17.
- Maxwell, James C. "On Governors." Proceedings of the Royal Society of London 16 (1868):270-83.
- Mayr, Otto. "Maxwell and the Origins of Cybernetics." Isis 62 (1971): 424-44.
- Minorsky, N. "Directional Stability of Automatically Steered Bodies." Journal of the American Society of Naval Engineers 42 (1922): 280-309.
- Nyquist, Harry. "Regeneration Theory." Bell System Technical Journal 11 (1932):126-47.
- Pound, Roscoe. "Mechanical Jurisprudence." Columbia Law Review 8 (1908):605-23.
- Radin, Max. "Case Law and Stare Decisis: Concerning Präjudizienrecht in Amerika." Columbia Law Review 33 (1933):199-212.
- Rosch, Eleanor H. "Classification of Real-world Objects: Origins and Representations in Cognition." In Thinking: Readings in Cognitive Science, pp. 212-22. Edited by P. N. Johnson-Laird and P. C. Wason. Cambridge: Cambridge University Press, 1977.
- Rosch, Eleanor H. "Cognitive Representations of Semantic Categories." Journal of Experimental Psychology: General 104 (1975):192-235.
- Rosch, Eleanor H. "On the Internal Structure of Perceptual and Semantic Categories." In Cognitive Development and the Acquisition of Language, pp. 111-44. Edited by Timothy E. Moore. New York: Academic Press, Inc., 1973.
- Tversky, Amos. "Features of Similarity." Psychological Review 84 (1977):327-52.

Books

- Alberti, Leon Battista. L'Architettura di Leonbatista Alberti. Translated from the Latin into Florentine Italian by Cosimo Bartoli. Firenze: L. Torrentino, 1550. Microfilm copy from the original in Biblioteca Vaticana, 1966.
- Armstrong, David M. Belief, Truth and Knowledge. Cambridge: Cambridge University Press, 1973.
- Bennett, S. A History of Control Engineering, 1800-1930. Stevenage, U.K.: Peter Peregrinus, Ltd., 1979.
- Brown, Gordon S., and Campbell, Donald P. Principles of Servomechanisms: Dynamics and Synthesis of Closed-Loop Control Systems. New York: John Wiley & Sons, Inc., 1948.
- Cardozo, Benjamin N. The Nature of the Judicial Process. New Haven: Yale University Press, 1921.
- Cooley, Thomas M. A Treatise on the Law of Torts or the Wrongs Which Arise Independently of Contract. 2 Vols., 3d. ed. Chicago: Callaghan & Company, 1906.
- D'Azzo, John J., and Houpis, Constantine H. Linear Control System Analysis and Design. New York: McGraw-Hill, Inc., 1975.
- Di Stefano, Joseph J.; Stubberud, Allen R.; and Williams, Ivan J. Feedback and Control Systems. New York: McGraw-Hill Book Company, 1967.
- Dworkin, Ronald. Taking Rights Seriously. Cambridge: Harvard University Press, 1978.
- Harré, [Horace] Rom The Principles of Scientific Thinking. Chicago: University of Chicago Press, 1970.
- Kuhn, Thomas S. The Essential Tension: Selected Studies in Scientific Tradition and Change. Chicago: The University of Chicago Press, 1977.
- Kuhn, Thomas S. The Structure of Scientific Revolutions. 2d. ed., enl. Chicago: The University of Chicago Press, 1970.
- Levi, Edward H. An Introduction to Legal Reasoning. Chicago: University of Chicago Press, 1949.
- Llewellyn, Karl N. The Common Law Tradition: Deciding Appeals. Boston: Little, Brown and Company, 1960.

- Mayr, Otto. Feedback Mechanisms in the Historical Collections of the National Museum of History and Technology. Washington, D.C.: Smithsonian Institution Press, 1971.
- Mayr, Otto. The Origins of Feedback Control. English ed. Cambridge: The M.I.T. Press, 1970.
- Mott, Neville F., and David E. A. Electronic Processes in Non-Crystalline Materials. Oxford: Oxford University Press, 1971.
- Putnam, Hilary W. Philosophical Papers. Vol. 2: Mind, Language and Reality. Cambridge: Cambridge University Press, 1975.
- Routh, Edward J. A Treatise on the Stability of a Given State of Motion, Particularly Steady Motion. London: Macmillan and Co., 1877.
- Serlio, Sebastiano. Libro Estraordinario di Sebastiano Serlio Bolognese. Venice: Francesco Senese & Zuane Kruger Alemano Compagni, 1566.
- Simpson, A. W. B. A History of the Common Law of Contract: the Rise of the Action of Assumpsit. Oxford: Clarendon Press, 1975.
- Summerson, John. The Classical Language of Architecture. Cambridge: The M.I.T. Press, 1963.
- Toulmin, Stephen. Human Understanding. Princeton: Princeton University Press, 1972.
- Vitruvius, Pollio. De Architectura. Translated from the Latin into Italian by Bono Mauro Bergomense and Benedicto Iovio. Commentary and illustrations by Cesare di Lorenzo Cesariano. Como, 1921. Facsimile reproduction: Bronx: B. Blom, 1968.
- Wasserstrom, Richard A. The Judicial Decision: Toward a Theory of Legal Justification. Stanford: Stanford University Press, 1961.
- Wiener, Norbert. Cybernetics, or Control and Communication in the Animal and the Machine. 2d. ed. Cambridge: The M.I.T. Press, 1961.
- Wittgenstein, Ludwig. Philosophical Investigations. Translated by G. E. M. Anscombe. New York: The Macmillan Company, 1958.
- Woodward, Robert U., and Armstrong, Scott. The Brethren: Inside the Supreme Court. New York: Simon and Schuster, 1979.

Court Cases

- Cadillac Motor Car Co. v. Johnson [1915]: The Federal Reporter. Vol. 221. St. Paul: West Publishing Co., 1915:801-6.
- Devlin v. Smith [1882]: Sickels, H. E. Reports of Cases Decided in the Court of Appeals of the State of New York. Vol. 44. Albany: Weed, Parsons & Co., 1883:470-79.
- Dixon v. Bell [1816]: Maule, George, and Selwyn, William. Reports of Cases Argued and Determined in The Court of King's Bench. Vol. 5. London: J. Butterworth and Son, 1823:198-200.
- Gladwell v. Steggall [1839]: Scott, John. Cases in the Court of Common Pleas and Exchequer Chamber. Vol. 8. London: S. Sweet, 1841: 60-70.
- Heaven v. Pender [1883]: Hall, John E., et. al. The Law Reports: Cases Determined in the Queen's Bench Division and on Appeal therefrom in the Court of Appeal, and Decisions on Crown Cases Reserved. Edited by J. R. Bulwer. Vol. 11. London: William Clowes and Sons, Limited, 1882-83:503-17.
- Huset v. J. I. Case Threshing Machine Co. [1903]: The Federal Reporter. Vol. 120. St. Paul: West Publishing Co., 1903:865-73.
- Langridge v. Levy [1837]: Meeson, R., and Welsby, W. N. Reports of Cases Argued and Determined in the Courts of Exchequer & Exchequer Chamber. Vol. 2. London: S. Sweet, 1838:519-31.
- Longmeid v. Holliday [1851]: Welsby, W. N., et. al. The English Reports: Exchequer Division. Edited by Max. Robertson. Vol. 155. Edinburgh: W. Green & Son, Limited, 1916:752-55.
- Loop v. Litchfield [1870]: Hand, Samuel. Reports of Cases Argued and Determined in the Court of Appeals of the State of New York. Vol. 3. New York: Banks & Brothers, 1874:351-61.
- Lossee v. Clute [1873]: Sickels, H. E. Reports of Cases Decided in the Commission of Appeals of the State of New York. Vol. 6. New York: Banks & Brothers, 1874:494-97.
- MacPherson v. Buick Motor Co. [1916]: J. Newton Fiero. Reports of Cases Decided in the Court of Appeals of the State of New York. Vol. 217. Albany: J. B. Lyon Company, 1916:382-401.
- Pasley v. Freeman [1789]: Durford, Charles, and East, Edward Hyde. Term Reports in the Court of King's Bench. Vol. 3. London: J. Butterworth and Son, 1817:51-65.

Pippin v. Sheppard [1822]: Price, George. Reports of Cases Argued and Determined in the Court of Exchequer and in the Exchequer Chamber. Vol. 11. London: Henry Butterworth, 1928:400-410.

Statler v. Ray [1909]: J. Newton Fiero. Reports of Cases Decided in the Court of Appeals of the State of New York. Vol. 195. Albany: J. B. Lyon Company, 1909:478-85.

Thomas v. Winchester [1852]: Selden, Harry R. Reports of Cases Argued and Determined in the Court of Appeals of the State of New-York. Vol. 2. Albany: W. C. Little & Co., 1854:397-411.

Torgeson v. Schultz [1908]: Newcomb, Alvah S. Reports of Cases Decided in the Court of Appeals of the State of New York. Vol. 192. Albany: J. B. Lyon Company, 1909:156-61.

Winterbottom v. Wright [1842]: Meeson, R., and Welsby, W. N. Reports of Cases Argued and Determined in the Courts of Exchequer & Exchequer Chamber. Vol. 10. London: S. Sweet, 1843:109-116.

Theses

Fleming, Robin S. "Patterns of Scientific Change: The Acid-Base Research Tradition." Ph.D. dissertation, University of Virginia, 1975.

United States Government Documents

U.S. Department of Commerce. Interagency Task Force on Product Liability--Final Report. Washington, D.C.: U.S. Government Printing Office, 1978.