NOTE: This material may be protected by copyright law (Title 17 U.S. code).
This copy is being furnished for private research use only.

NYU DIBNER LIBRARY
3 1155 06050484 7

...ent..." —Douglas Hofstadter

JOHN H. HOLLAND

HIDDEN ORDER

How Adaptation Builds Complexity

TJ
217
.H64
1995
c.1

Many of the designations used by manufacturers and sellers to distinguish their products are claimed as trademarks. Where those designations appear in this book and Addison-Wesley was aware of a trademark claim, the designations have been printed in initial capital letters.

Library of Congress Cataloging-in-Publication Data

Holland, John H. (John Henry), 1929-
Hidden order : how adaption builds complexity / John H. Holland.
p. cm.—(Helix books)
Includes bibliographical references and index.
ISBN 0-201-40793-0 (hardcover)
ISBN 0-201-44230-2 (paperback)
1. Adaptive control systems—Mathematical models. 2. Adaptation (Biology)—Mathematical models. I. Title. II. Series: Helix books.
TJ217.H64 1995
003.7—dc20 95-20042
CIP

Copyright © 1995 by John H. Holland

All rights reserved. No part of this publication may be reproduced, stored in a retrieval system, or transmitted, in any form or by any means, electronic, mechanical, photocopying, recording, or otherwise, without the prior written permission of the publisher. Printed in the United States of America. Published simultaneously in Canada.

Cover design by Suzanne Heiser
Text design by Diane Levy
Set in 10½-point Bembo by Pagesetters, Incorporated
1 2 3 4 5 6 7 8 9 10-DOH-0099989796
First printing, July 1995
First paperback printing, August 1996

·1·

Basic Elements

On an ordinary day in New York City, Eleanor Petersson goes to her favorite specialty store to pick up a jar of pickled herring. She fully expects the herring to be there. Indeed, New Yorkers of all kinds consume vast stocks of food of all kinds, with hardly a worry about continued supply. This is not just some New Yorker persuasion; the inhabitants of Paris and Delhi and Shanghai and Tokyo expect the same. It's a sort of magic that everywhere is taken for granted. Yet these cities have no central planning commissions that solve the problems of purchasing and distributing supplies. Nor do they maintain large reserves to buffer fluctuations; their food would last less than a week or two if the daily arrivals were cut off. How do these cities avoid devastating swings between shortage and glut, year after year, decade after decade?

The mystery deepens when we observe the kaleidoscopic nature of large cities. Buyers, sellers, administrations, streets, bridges, and buildings are always changing, so that a city's coherence is somehow imposed on a perpetual flux of people and structures. Like the standing wave in front of a rock in a fast-moving stream, a city is a pattern in time. No single constituent remains in place, but the city persists. To enlarge on the previous question: What enables cities to retain their coherence despite continual disruptions and a lack of central planning?

There are some standard answers to this question, but they really do

not resolve the mystery. It is suggestive to say that Adam Smith's "invisible hand," or commerce, or custom, maintains the city's coherence, but we still are left asking How?

Other patterns in time exhibit similar riddles. For instance, if we shift to the microscopic level, we find another community every bit as complicated as New York City. The human immune system is a community made up of large numbers of highly mobile units called *antibodies* that continually repel or destroy an ever-changing cast of invaders called *antigens*. The invaders—primarily biochemicals, bacteria, and viruses—come in endless varieties, as different from one another as snowflakes. Because of this variety, and because new invaders are always appearing, the immune system cannot simply list all possible invaders. It must change or adapt (Latin, "to fit") its antibodies to new invaders as they appear, never settling to a fixed configuration. Despite its protean nature, the immune system maintains an impressive coherence. Indeed, your immune system is coherent enough to provide a satisfactory scientific definition of your *identity*. It is so good at distinguishing you from the rest of the world that it will reject cells from any other human. As a result, a skin graft even from a sibling requires extraordinary measures.

How does the immune system develop its exquisite sense of identity, and what makes that identity vulnerable? How does an immune disease such as AIDS manage to destroy the identity? We can say that the identifications, and the misidentifications, are a product of "adaptation," but the "how" of this adaptive process is far from obvious.

It is more than an academic quest to try to understand the persistence and operation of these two complex communities. Pressing problems, such as prevention of inner-city decay and control of diseases such as AIDS, turn on this understanding. Once we look in this direction, we see that there are other complex systems that pose similar questions, and they too present troubling, long-range problems.

Consider the mammalian central nervous system (CNS). Like the immune system, the CNS consists of a large number of component cells, called neurons, that occur in a wide range of forms. Even a simple CNS consists of hundreds of millions of neurons, of hundreds of types,

and each neuron directly contacts hundreds, even thousands, of other neurons to form a complex network. Pulses of energy flash over this network, producing Sherrington's (1951) "enchanted loom." This network is similar to the immune system, with an aggregate emergent identity that *learns* speedily and with great facility. Though the activity of an individual neuron can be complex, it is clear that the behavior of the CNS aggregate identity is much more complex than the sum of these individual activities. The behavior of the central nervous system depends on the *interactions* much more than the actions. The sheer number of interactions—hundreds of millions of neurons, each undergoing thousands of simultaneous interactions in thousandths of a second—takes us well beyond any of our experience with machines. The most sophisticated computer, in comparison, seems little more than an automated abacus. The myriad interactions, modified by learned changes, yield the unique ability of canids, felines, primates, and other mammals to anticipate the consequences of their actions by modeling their worlds.

After more than a century of intensive effort, we still cannot model many basic capabilities of the CNS. We cannot model its ability to parse complex unfamiliar scenes into familiar elements, let alone its ability to construct experience-based internal models. The relation between the distributed, diverse CNS and the phenomenon we call consciousness is largely unknown, a mystery that leaves us with few guidelines for the treatment of mental diseases.

Ecosystems share many of the features and puzzles presented by an immune system or a CNS. They exhibit the same overwhelming diversity. We have yet to assay the range of organisms present in a cubic meter of temperate-zone soil, let alone the incredible arrays of species in a tropical forest. Ecosystems are continually in flux and exhibit a wondrous panoply of interactions such as mutualism, parasitism, biological arms races, and mimicry (more about these later). Matter, energy, and information are shunted around in complex cycles. Once again, the whole is more than the sum of its parts. Even when we have a catalog of the activities of most of the participating species, we are far from understanding the effect of changes in the ecosystem. For

example, the stupendous richness of the tropical forest biome contrasts with the poverty of its soil. The forest can only maintain its diversity through a complex set of interactions that recycle sparse nutrients through the system, over and over again.

Until we have an understanding of these complicated, changing interactions, our attempts to balance extraction of ecosystem resources against sustainability will remain at best naive, at worst disastrous. We, as humans, have become so numerous that we perforce extensively modify ecological interactions, with only vague ideas of longer-range effects. Yet our well-being, even our survival, depends on our being able to use these systems without destroying them. Attempts to turn tropical forest into farmland, or to fish the Grand Banks "efficiently," are only symptoms of a problem that year by year becomes more serious.

Many other complex systems show coherence in the face of change. But we can already begin to extract some of the commonalities, and we will later examine additional systems in this light. We can see, for instance, that the coherence and persistence of each system depend on extensive interactions, the aggregation of diverse elements, and adaptation or learning. We have also noted that several perplexing problems of contemporary society—inner-city decay, AIDS, mental disease and deterioration, biological sustainability—are likely to persist until we develop an understanding of the dynamics of these systems. We will see that economies, the Internet, and developing embryos offer similar challenges—trade balances, computer viruses, and birth defects, for example—and we will encounter still others.

Even though these complex systems differ in detail, the question of coherence under change is the central enigma for each. This common factor is so important that at the Santa Fe Institute we collect these systems under a common heading, referring to them as *complex adaptive systems* (*cas*). This is more than terminology. It signals our intuition that general principles rule *cas* behavior, principles that point to ways of solving the attendant problems.

Our quest is to extract these general principles. The quest is new, so this book can only begin to map the territory. And much of that map

will consist of terra incognita and legends such as "monsters exist here." Nevertheless, we have come far enough to do more than make casual comparisons. In this first chapter, we can observe some of the prominent landmarks and we can estimate what kinds of apparatus will be needed to come to a broad understanding of complex adaptive systems.

Objectives

The purpose of this book is to explore the ways in which our intuitions about *cas* can be transformed into a deeper understanding. Theory is crucial. Serendipity may occasionally yield insight, but is unlikely to be a frequent visitor. Without theory, we make endless forays into uncharted badlands. With theory, we can separate fundamental characteristics from fascinating idiosyncrasies and incidental features. Theory supplies landmarks and guideposts, and we begin to know what to observe and where to act.

One specific piece of understanding that theory could supply is a more principled way of locating "lever points" in *cas*. Many *cas* have the property that a small input can produce major predictable, directed changes—an amplifier effect. A familiar example is a vaccine. An infection into our bloodstream of a small amount of an incapacitated antigen, say the measles virus, can stimulate the immune system to produce enough antibodies to make us completely immune to the disease. The vaccine "levers" the immune system into learning about the disease, saving the costly, uncomfortable procedure of learning about the disease "on line." We know of other lever points in other *cas*, but to date we have no comprehensive method of searching them out. Theory is our best hope of finding such a method.

The task of formulating theory for *cas* is more than usually difficult because the behavior of a whole *cas* is more than a simple sum of the behaviors of its parts; *cas* abound in nonlinearities (more about this shortly). Nonlinearities mean that our most useful tools for generalizing observations into theory—trend analysis, determination of equilibria, sample means, and so on—are badly blunted. The best way to

compensate for this loss is to make cross-disciplinary comparisons of *cas*, in hopes of extracting common characteristics. With patience and insight we can shape those characteristics into building blocks for a general theory. Cross-comparisons provide another advantage: characteristics that are subtle and hard to extract from one system can be salient and easy to examine in another. This chapter is about seven characteristics that cross-disciplinary comparisons suggest are central to a broad understanding of *cas*. Subsequent chapters will weave these characteristics into the elements of a theory.

Agents, Meta-Agents, and Adaptation

Before going on to a description of the characteristics themselves, I should say something more about the general setting. *Cas* are, without exception, made up of large numbers of active elements that, from the examples we've seen, are diverse in both form and capability (see Figure 1.1). Think of the great array of firms in New York City or the exquisitely tuned antibodies in the immune system. To refer to active elements

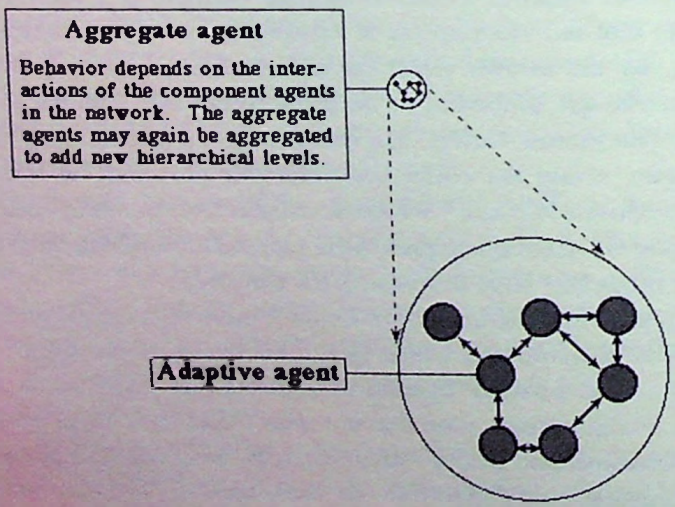


Figure 1.1 A Complex Adaptive System.

without invoking specific contexts, I have borrowed the term *agent* from economics. The term is descriptive but avoids preconceptions.

If we are to understand the interactions of large numbers of agents, we must first be able to describe the capabilities of individual agents. It is useful to think of an agent's behavior as determined by a collection of rules. Stimulus-response rules are typical and simple: IF stimulus *s* occurs THEN give response *r*. IF the market goes down THEN sell. IF the car has a flat tire THEN get the jack. And so on. To define the set of stimulus-response rules possible for a given agent, we must describe the stimuli that agent can receive and the responses it can give (see Figure 1.2).

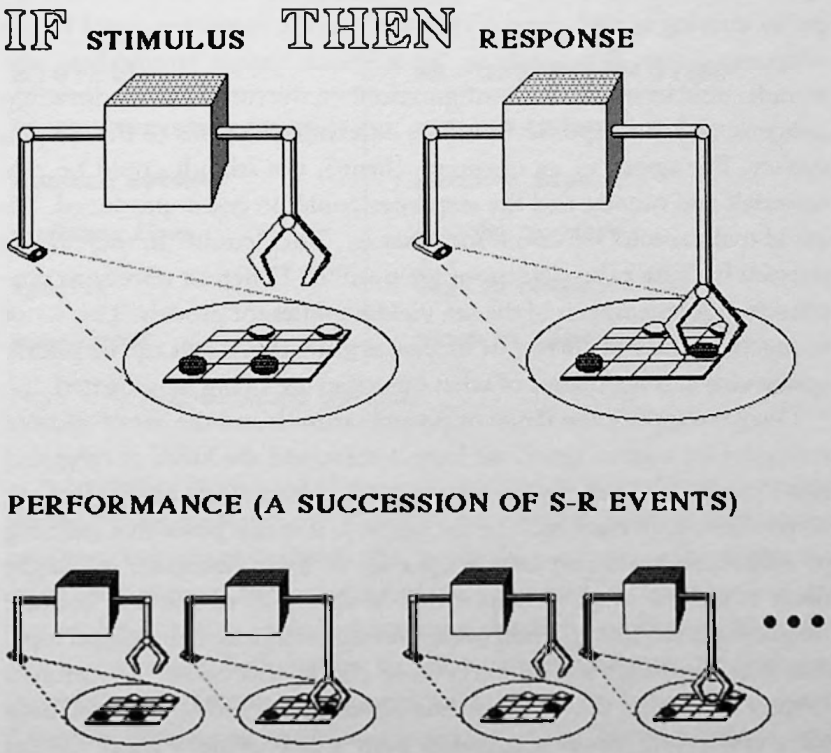
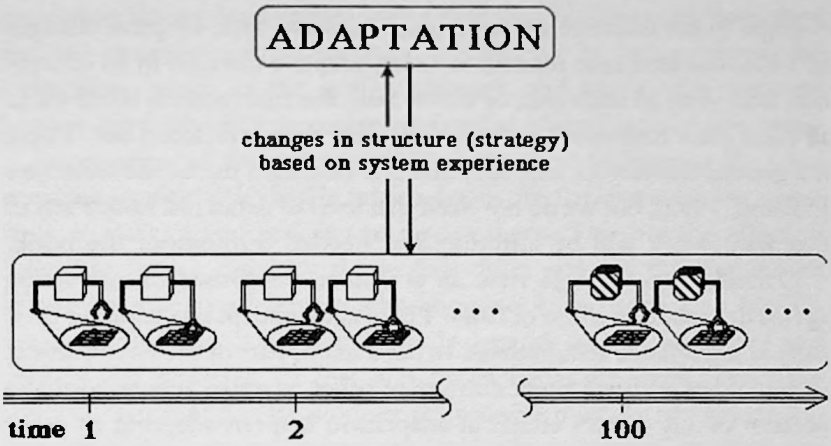


Figure 1.2 A Rule-Based Agent.

Though stimulus-response rules are limited in scope, there are simple ways of expanding that scope. Indeed, with minor changes, the scope can be enlarged sufficiently that clusters of rules can generate any behavior that can be computationally described. In the definition of these rules, our intent is *not* to claim that we can locate the rules explicitly in the real agents. Rules are simply a convenient way to describe agent strategies. In the next chapter I will say more about this rule-based approach to agent behavior; for now, let us treat it as a descriptive device.

A major part of the modeling effort for any *cas*, then, goes into selecting and representing stimuli and responses, because the behaviors and strategies of the component agents are determined thereby. For agents in the central nervous system (neurons), the stimuli could be pulses arriving at each neuron's surface, and the responses could be the outgoing pulses. For agents in the immune system (antibodies), the stimuli could be molecular configurations on the surface of the invading antigens, and the responses could be differing adhesions to the antigen surface. For agents in an economy (firms), the stimuli could be raw materials and money, and the responses could be goods produced. We could make similar selections for other *cas*. The "could" in each case is relevant because other selections are possible. Different selections emphasize different aspects of the *cas*, yielding different models. This is not so much a matter of correct or incorrect (though models can be poorly conceived) as it is a matter of what questions are being investigated.

Once we specify the range of possible stimuli and the set of allowed responses for a given agent, we have determined the kinds of rules that agent can have. Then, by looking at these rules acting in sequence, we arrive at the behaviors open to the agent. It is at this point that learning or adaptation enters. In setting up a list of basic elements, we might think it natural to put "adaptation" at the head of the list, because adaptation is the sine qua non of *cas*. But adaptation is such a broad topic that it encompasses almost everything else in this book. The present chapter centers on the more specific characteristics of *cas*, so I will only say a few words about adaptation here and provide a more careful discussion in the next chapter.



SYSTEM	MODIFICATION TIME
central nervous system	seconds to hours
immune system	hours to days
business firm	months to years
species	days to centuries
ecosystem	years to millennia

Figure 1.3 Adaptation and Learning.

Adaptation, in biological usage, is the process whereby an organism fits itself to its environment. Roughly, experience guides changes in the organism's structure so that as time passes the organism makes better use of its environment for its own ends (see Figure 1.3). Here we expand the term's range to include learning and related processes. With this extension, adaptation applies to all *cas* agents, despite the different timescales of different *cas* processes. And indeed, the timescales do vary. Adaptive changes in individual neurons in the nervous system take place over an interval that ranges from seconds to hours; adaptive

changes in the immune system require hours to days; adaptive changes in a business firm take months to years; adaptive changes in an ecosystem take years to millennia or more. Still, the mechanisms involved in all these cases have much in common, once time is factored out. There is a general framework that supports this extended use of the term (see Holland, 1992), but we do not need that level of detail just now. Parts of that framework will be introduced as needed throughout the book.

Overall, then, we will view *cas* as systems composed of interacting agents described in terms of rules. These agents adapt by changing their rules as experience accumulates. In *cas*, a major part of the environment of any given adaptive agent consists of other adaptive agents, so that a portion of any agent's efforts at adaptation is spent adapting to other adaptive agents. This one feature is a major source of the complex temporal patterns that *cas* generate. To understand *cas* we must understand these ever-changing patterns. The rest of this book is devoted to developing such an understanding by filling in this rough sketch, adding detail, content, and relevant pieces of theory. Now to our seven basics.

Seven Basics

The seven basics consist of four properties and three mechanisms that are common to all *cas*. They are not the only basics that could be selected from a list of common characteristics; the selection process is, in part, a matter of taste. Still, all the other candidates of which I am aware can be "derived" from appropriate combinations of these seven.

In presenting the basics, I have ordered them in a way that emphasizes their interrelations rather than grouping them into properties and mechanisms.

Aggregation (*Property*)

Aggregation enters into the study of *cas* in two senses. The first refers to a standard way of simplifying complex systems. We aggregate similar things into categories—trees, cars, banks—and then treat them as equivalent. Humans analyze familiar scenes in this way with the greatest of ease. Not too surprisingly, the categories we choose are reusable; we

can almost always decompose novel scenes into familiar categories. And we can generate scenes we have never seen by recombining the categories—much as the griffin, chimera, and harpy of a medieval bestiary are formed by recombining familiar animal parts in new ways.

Aggregation in this sense is one of our chief techniques for constructing models. We decide which details are irrelevant for the questions of interest and proceed to ignore them. This has the effect of collecting into a category things that differ only in the abandoned details; the category becomes a building block for the model. Modeling, it should be clear, is an art form. It depends on the experience and taste of the modeler. In this it is much like cartooning, especially political cartooning. The modeler (cartoonist) must decide which features to make salient (exaggerate), and which features to eliminate (avoid), in order to answer the questions (make the political point).

The second sense of aggregation is closely related to the first, but it is more a matter of *what cas do*, rather than *how* we model them. It concerns the emergence of complex large-scale behaviors from the aggregate interactions of less complex agents. An ant nest serves as a familiar example. The individual ant has a highly stereotyped behavior, and it almost always dies when circumstances do not fit the stereotype. On the other hand, the ant aggregate—the ant nest—is highly adaptive, surviving over long periods in the face of a wide range of hazards. It is much like an intelligent organism constructed of relatively unintelligent parts. Douglas Hofstadter's wonderful chapter "Ant Fugue" in his 1979 book makes this point better than anything else I have read. In it the ant nest provides a comprehensible version of more spectacular emergent phenomena, such as the intelligence of large numbers of interconnected neurons, or the identity provided by a diverse array of antibodies, or the spectacular coordination of an organism made of myriad cell types, or even the coherence and persistence of a large city.

Aggregates so formed can in turn act as agents at a higher level—*meta-agents*. The interactions of these meta-agents are often best described in terms of their aggregate (first sense) properties (see Figure 1.4). Thus we speak of the gross domestic product of an economy, or

the identity provided by the immune system, or the behavior of a nervous system. Meta-agents can, of course, aggregate (second sense) in turn to yield meta-meta-agents. When this process is repeated several times, we get the hierarchical organization so typical of *cas*.

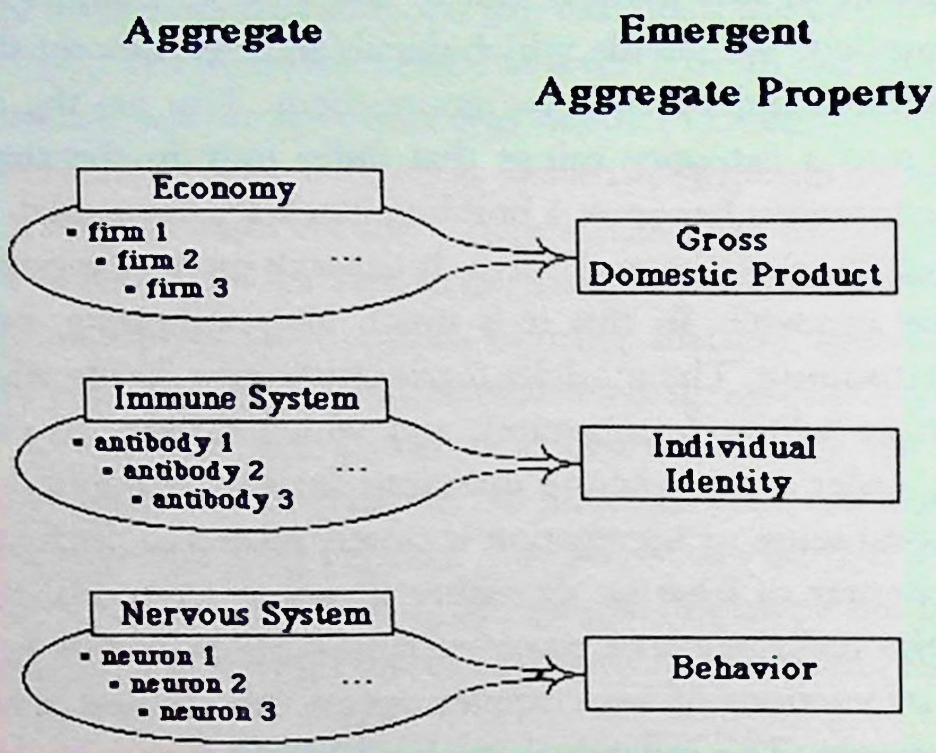


Figure 1.4 Aggregation and Aggregate Properties.

Aggregation in the second sense is indeed a basic characteristic of all *cas*, and the emergent phenomena that result are the most enigmatic aspect of *cas*. The study of *cas* turns on our ability to discern the mechanisms that enable simple agents to form highly adaptive aggregates. What kind of "boundaries" demarcate these adaptive aggregates? How are the agent interactions within these boundaries directed and coordinated? How do the contained interactions generate behaviors that transcend the behaviors of the component agents? We must be able to answer such questions if we are to resolve the mysteries, and the difficulties, that attend *cas*.

TAGGING (*MECHANISM*)

There is a mechanism that consistently facilitates the formation of aggregates—a mechanism that in this book will go by the name *tagging*.

The most familiar example is a banner or flag that is used to rally members of an army or people of similar political persuasion. A more operational version of a tag, in these days of Internet, is the header on a message that knits together members of a bulletin board or conference group. Still more operational are the "active sites" that enable antibodies to attach themselves to antigens. The sophistication of this particular version of tagging is well described in Edelman's (1988) discussion of cell adhesion molecules. We can continue with the visual patterns and pheromones that facilitate selective mating in animals, and the trademarks, logos, and icons that facilitate commercial interaction (see Figure 1.5). It soon becomes apparent that tagging is a pervasive mechanism for aggregation and boundary formation in *cas*.

When we closely examine different instances of tagging, we see there is a common feature: *cas* use tags to manipulate symmetries. Because symmetries are common, we often use them in perceiving or modeling our day-to-day world, sometimes quite unconsciously. They enable us to ignore certain details, while directing our attention to others. Weyl (1952) gives a rich exposition of this point. The classic example of a full-blown symmetry is a perfect sphere, say the white cue ball in billiards. A cue ball exhibits complete rotational symmetry, so that rotation in any direction produces no noticeable change. If we put a stripe around the cue ball's "equator," turning it into one of the other billiard balls, we break the symmetry, allowing us to distinguish the previously indistinguishable. For example, if we start the striped ball spinning, we can easily observe whether or not the ball's axis of rotation defines the equator marked out by the stripe. Most rotations produce noticeable changes, except for those around the axis that defines the cue ball's equator. That is, some symmetries are broken and others remain. In general, tags enable us to observe and act on properties previously hidden by symmetries.

To carry the example a bit further, consider a *set* of cue balls in rapid motion on a billiard table, say after a strong "break." We cannot distinguish the individual cue balls unless we keep a careful record of their trajectories. But again, we can break the symmetry via a tag. If we

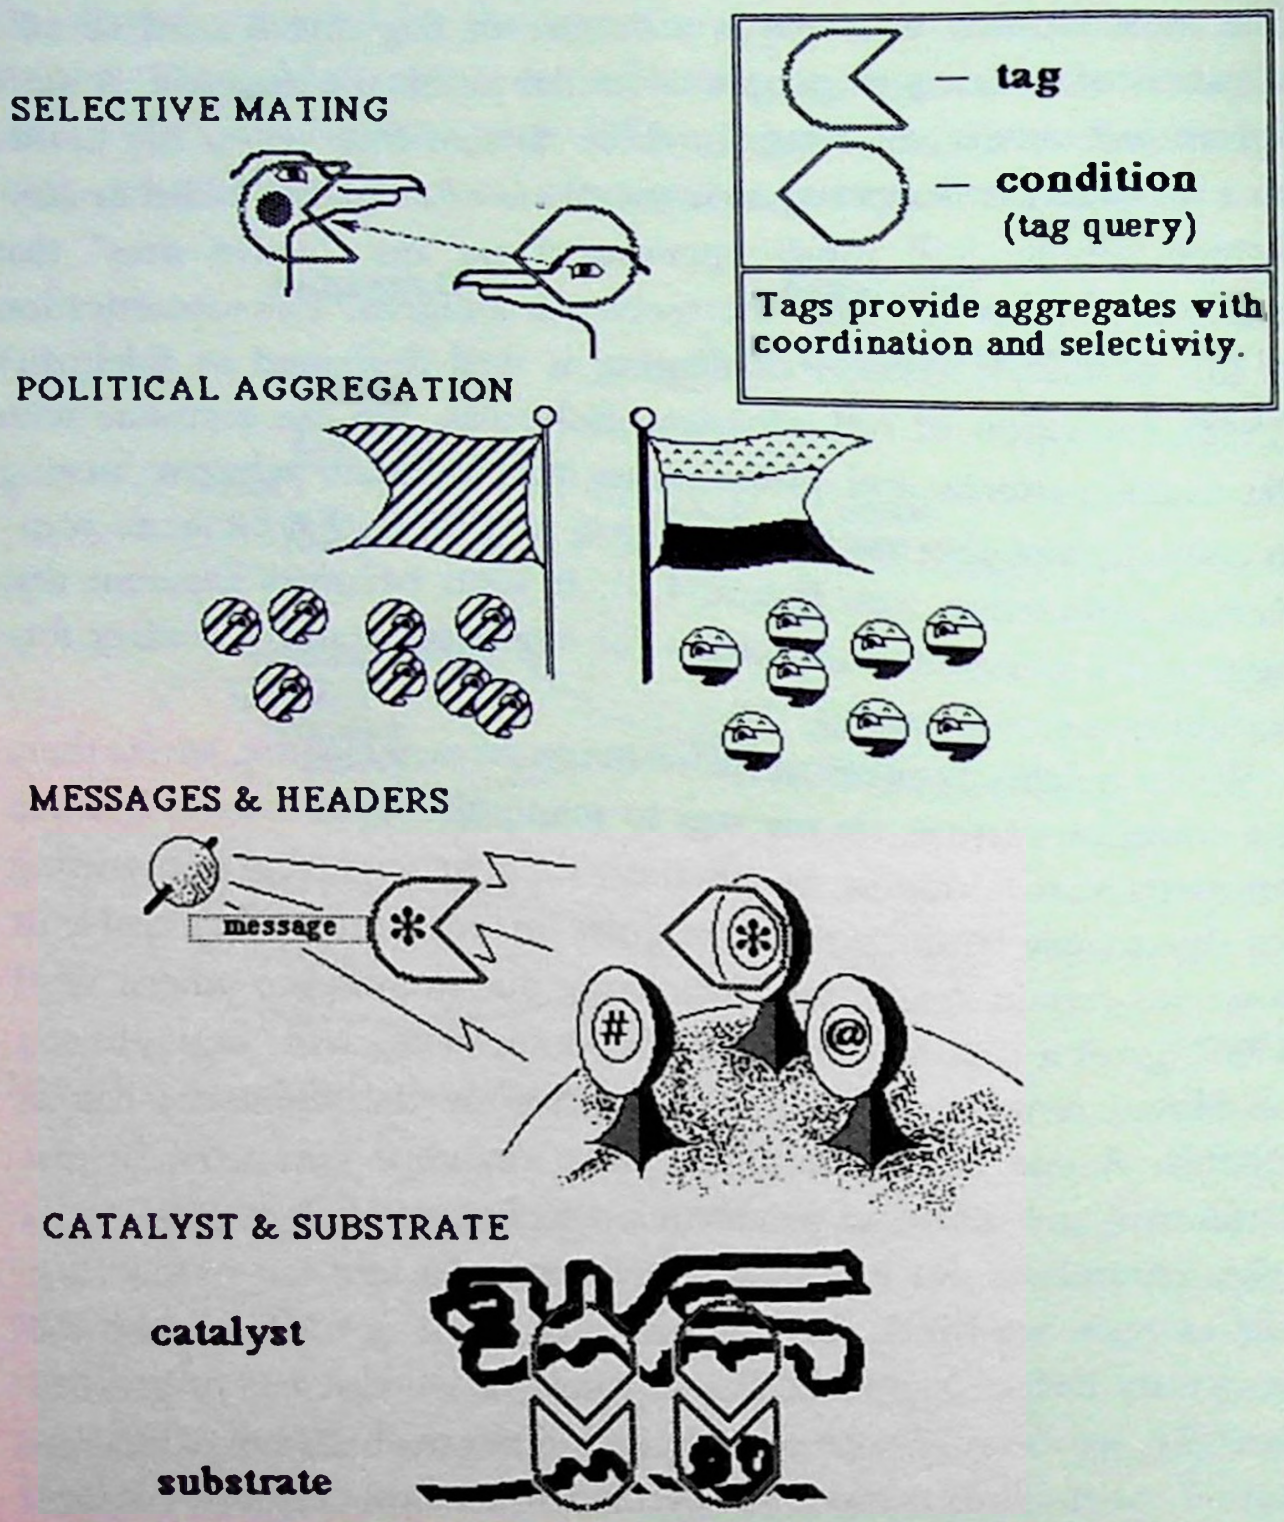


Figure 1.5 Tags and Aggregates.

put a striped cue ball in with the other cue balls, we can easily track it despite its motion.

Tags are a pervasive feature of *cas* because they facilitate selective interaction. They allow agents to select among agents or objects that would otherwise be indistinguishable. Well-established tag-based inter-

actions provide a sound basis for filtering, specialization, and cooperation. This, in turn, leads to the emergence of meta-agents and organizations that persist even though their components are continually changing. Ultimately, tags are the mechanism behind hierarchical organization—the agent / meta-agent / meta-meta-agent / . . . organization so common in *cas*. We'll see many examples of the origin and intervention of tags as we go along.

Nonlinearity (*Property*)

It is little known outside the world of mathematics that most of our mathematical tools, from simple arithmetic through differential calculus to algebraic topology, rely on the assumption of linearity. Roughly, linearity means that we can get a value for the whole by adding up the values of its parts. More carefully, a *function* is *linear* if the value of the function, for any set of values assigned to its arguments, is simply a weighted sum of those values. The function $3x + 5y + z$, for example, is linear.

We say some numerical property of a system is linear, relative to numerical values assigned to its parts, if the property is a linear function of those values. Consider, for example, the fuel consumption c of a plane as a function of its velocity v and its altitude x. Given suitable units for fuel consumption, altitude, and velocity, we might be able to establish that

$$c = (0.5)v + (-0.1)v.$$

Fuel consumption then would be linear in terms of velocity and altitude.

Polls, or project trends, or industrial statistics, all of which employ summation, are only useful if they describe linear properties of the underlying systems. It is so much easier to use mathematics when systems have linear properties that we often expend considerable effort to justify an assumption of linearity. Whole branches of mathematics are devoted to finding linear functions that are reasonable approximations when linearity cannot be directly established. Unfortunately, none of this works well for *cas*. To attempt to study *cas* with these

techniques is much like trying to play chess by collecting statistics on the way pieces move in the game.

Let me illustrate the difficulty by starting with one of the simplest *non*linear interactions, that between a predator population and its prey. The model we look at, despite its simple assumptions, does a satisfactory job of describing real data, such as the centuries-long record of lynx-hare interactions derived from the Hudson Bay Company's yearly records of pelt acquisitions. In putting this model together, we sketch a typical procedure for building mathematical models. When we have finished, we'll have a clear example of the complications caused by nonlinearities.

We begin with the commonsense observation that increases in either the predator population or the prey population increase the likelihood of encounters between predator and prey. Symbolically, if U represents the number of predators in a given area, say a square mile, and V represents the number of prey in the same area, then the number of interactions per unit time, say a day, is given by cUV, where c is a constant that reflects the efficiency of the predator (for example, the average rate at which it searches the territory). If $c = 0.5$, $U = 2$, and $V = 10$, then there will be

$$cUV = 0.5(2)(10) = 10 \text{ encounters}$$

per day per square mile. If the number of predators increases by 2, so that $U = 4$, and the number of prey increases by 10, so that $V = 20$, then the number of encounters will be quadrupled to

$$cUV = 0.5(4)(20) = 40 \text{ encounters}$$

per day per square mile. This expression involves a nonlinearity, one of the simplest, because it entails the *product* of two distinct variables instead of their *sum*. That is, the overall predator-prey interaction cannot be obtained merely by adding predator activity to prey activity.

Our next step is to take explicit account of the fact that the populations change over time. Notationally, we let $U(t)$ stand for the population of predators at time t; similarly $V(t)$ stands for the population of

prey at time t. We augment the predator-prey interaction by allowing for births and for deaths from causes other than predation. Taking the simplest approach, we set a common birthrate b for all predators, so that the number of predator births at time t is $bU(t)$. Deaths can be handled similarly by using a common death rate d, so that the number of predator deaths at time t is $dU(t)$.

If we ignore predator-prey interactions for a moment, we arrive at a simple model of the way the population of predators changes over time. The size of the population after one unit of time has elapsed is the population at time t, minus the deaths, plus the births, or

$$U(t + 1) = U(t) - dU(t) + bU(t).$$

This equation, with allowances for aging, is the foundation for population projections and such mundane things as life insurance premiums. We use exactly the same argument to get a similar equation for the prey,

$$V(t + 1) = V(t) - d'V(t) + b'V(t),$$

where b' and d' are the respective birth and death rates for the prey (again without interactions).

To reintroduce the effect of predator-prey interactions, we incorporate the intuitive idea that the predator enhances its well-being each time it catches prey. Ultimately this process exerts a positive effect on the predator's production of offspring. To capture this idea mathematically, introduce another constant r that represents the efficiency of transforming captured prey (food) into offspring. More interactions mean more births, so using the interaction rate $cU(t)V(t)$, we get

$$r[cU(t)V(t)]$$

as the enhancement in births because of predator-prey interaction. The population change for predators then becomes

$$U(t + 1) = U(t) - dU(t) + bU(t) + r[cU(t)V(t)].$$

For prey, capture by a predator increases the number of deaths. Using r' to indicate the vulnerability to capture and death during interactions, we obtain a population equation for the prey,

$$V(t + 1) = V(t) - d' V(t) - r'[cU(t) V(t)] + b' V(t).$$

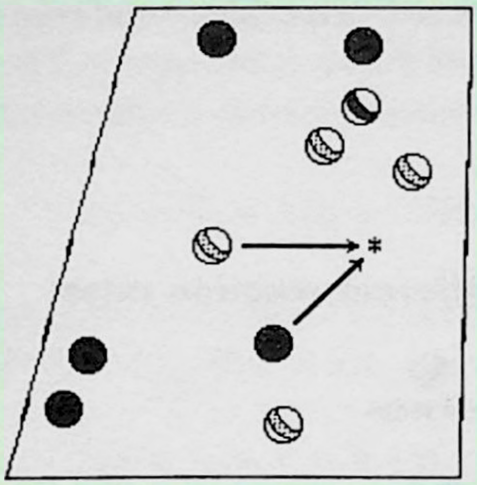
This pair of equations for $U(t + 1)$ and $V(t + 1)$ is a version of the famous Lotka-Volterra model (see Lotka, 1956). Standard ways of simplifying and solving the Lotka-Volterra equations show that, under most conditions, the predator population will go through a series of oscillations between feast and famine, as will the prey population. This prediction is borne out by the Hudson Bay Company's records. In the long run, extensions of such models should help us understand why predator-prey interactions exhibit strong oscillations, whereas the interactions that form a city are typically more stable. For now we are only interested in the effect of nonlinearities on such modeling efforts.

Let's return to the interactive part of the model. The $cU(t)V(t)$ formulation is actually a starting point for many other models, including interactions between atoms or molecules or even billiard balls. To study the effect of nonlinear interactions in the simplest possible setting, we shift back to the billiard balls (see Figure 1.6).

Let's restrict the model to just three "species" of billiard balls: white balls with a red stripe, white balls with an orange stripe, and solid-blue balls. Assume that there are several of each on the table and that they are in random motion—a kind of "big bang" or, better, "big break." Also assume, somewhat fancifully, that the "stripes" sometimes stick to the "solids" when they collide, as if they had dots of Velcro on their surfaces. The earlier formula cUV can now be used to model the rate at which the "stripe/solid compounds" form.

To see this, let's begin with the red-stripe/blue-solid combination. Our U gives the proportion of red-stripes on the table, V gives the proportion of blue-solids, and the constant c now gives a reaction rate that depends on the stickiness of the red-stripe/blue-solid combination. Using $Z(t)$ to represent the proportion of red-stripes stuck to blue-solids at time t, we get a simpler version of the Lotka-Volterra equation,

The simplest models of interactions use *random collisions*
(e.g. atomic, chemical, and predator-prey models)



Total number of balls: 10

Proportion of ◍: 4/10 = 0.4

Proportion of ●: 5/10 = 0.5

Some collisions produce a compound (a product).

The proportion of collisions resulting in a compound is set by a *reaction rate* using the (nonlinear) equation:

[propor.◍] x [propor.●] x *rate* = [propor.◍●]
[0.4] x [0.5] x 0.5 = [0.1]

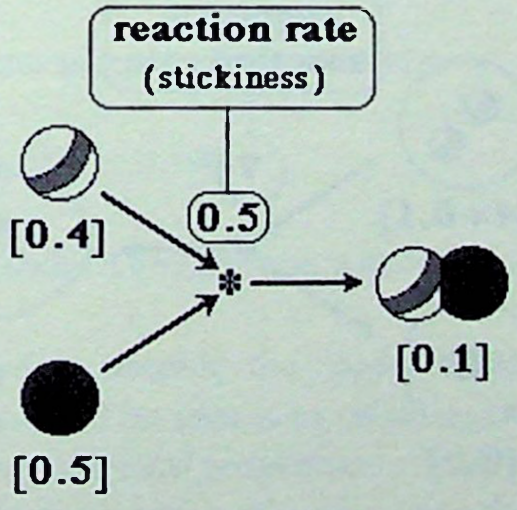


Figure 1.6 A Billiard Ball Model of Interaction.

$$Z(t + 1) = Z(t) + cU(t)V(t).$$

For example, if $Z(t) = 0$, $U(t) = 0.4$, $V(t) = 0.5$, and $c = 0.5$, then after one unit of time the proportion of the red-stripe/blue-solid compound is

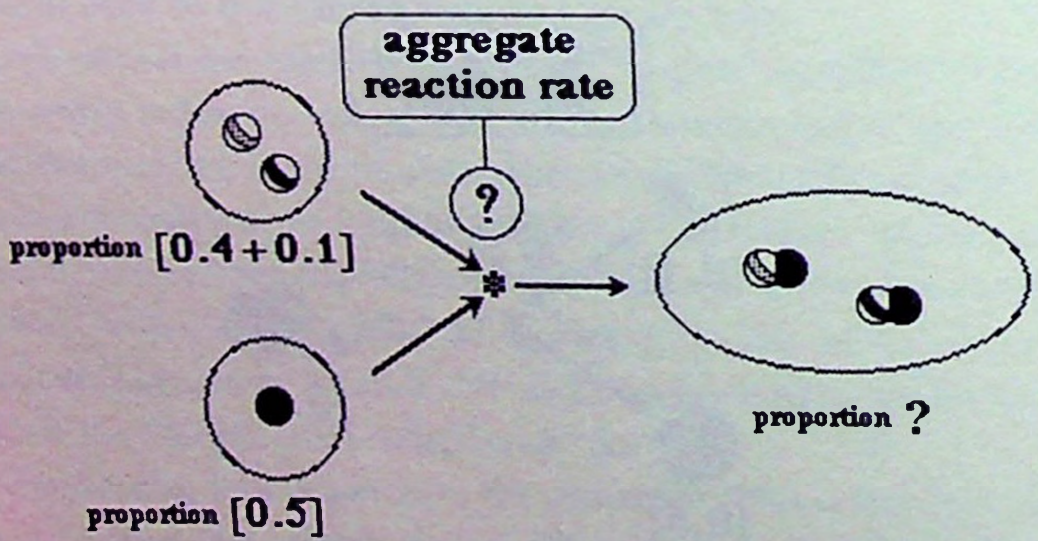
$$Z(t + 1) = 0 + 0.5(0.4)(0.5) = 0.1.$$

Different kinds of balls may have different reaction rates:

REACTION RATES	orange stripe	red stripe
blue solid	0.5	0.1

Suppose we want to know the proportion of collisions resulting in a stripe-solid compound {and }.

Can we build a simple model by assigning an *aggregate* (average) *reaction rate* to the overall process?



This reaction aggregates the stripes so it uses only the total proportion of stripes {proportion of + proportion of }.

FOR THE MODEL TO WORK: Two mixes of stripes with the same total must produce the same result.

Figure 1.7 Aggregate Reactions.

We can proceed similarly with the orange-stripes, allowing for the fact that stickiness of the orange-stripes may be different from that of the red-stripes (see Figure 1.7). Let $W(t)$ be the proportion of orange-stripes at time t, let $Y(t)$ be the proportion of the orange-stripe/blue-solid compound at time t, and let c' be the reaction rate determined by the stickiness of orange-stripes. Then, as for the red-stripes, the formula

$$Y(t + 1) = Y(t) + c' W(t) V(t)$$

gives the outcome of the interaction of orange-stripes with blue-solids. If $Y(t) = 0$, $W(t) = 0.1$, and $c' = 0.1$, then

$$Y(t + 1) = 0 + 0.1(0.1)(0.5) = 0.005.$$

We can get the total stripe-solid compound (red-stripe/blue-solid *plus* orange-stripe/blue-solid), $X(t)$, by adding the results of the separate reactions,

$$\begin{aligned} X(t + 1) &= Y(t + 1) + Z(t + 1) \\ &= Y(t) + Z(t) + cU(t) V(t) + c' W(t) V(t). \end{aligned}$$

Using the numerical values given earlier, we obtain

$$X(t + 1) = 0.1 + 0.005 = 0.105.$$

This part of the model is indeed linear—the whole does equal the sum of the parts!

Now suppose we want to simplify the model by aggregating the stripes into a single category. The idea is to calculate the total stripe-solid compound using only the total proportion of stripes on the table. Even when there are only two species of stripes, as in the present case, this aggregation cuts the complication (the number of equations) in half. When there are large numbers of species (as with an ecosystem or a city), aggregation makes the difference between feasibility and infeasibility when it comes to analysis. The simplification occurs because the aggregate equation uses a single variable $S(t)$ for the *total* population of

stripes, along with a single reaction coefficient c'', giving the single equation

$$X(t + 1) = X(t) + c''S(t)V(t).$$

There is a problem about the validity of this equation, however. For it to be useful, we must find a coefficient c'' that works for all mixes of stripes.

Under a standard linear approach, we would obtain c'' by averaging the coefficients of the individual stripe-solid interactions. However, this is the point at which the nonlinearities interfere. Consider two different mixes of stripes. In mix 1, the proportion of red-stripes is $U = 0.4$ and the proportion of orange-stripes is $W = 0.1$; in mix 2, the proportion is reversed, so that $U = 0.1$ and $W = 0.4$. In both cases the total number of stripes is $S = U + W = 0.5$. It follows that in both cases the equations for X must give the same answers for the proportion of stripe-solid compound, since all the numbers on the right side are the same. But what actually happens? Do the interactions of the two different mixes really yield the same total proportion of stripe-solid compound?

To check, let's carry out the detailed computation for the two mixes. For mix 1, we have already calculated the result when $X(t) = 0$,

$$X(t + 1) = Y(t + 1) + Z(t + 1) = 0.105.$$

For mix 2, we have

$$\begin{aligned} X(t + 1) = Y(t + 1) &= Z(t + 1) \\ &= Y(t) + Z(t) + cU(t)V(t) + c'W(t)V(t) \\ &= 0 + 0 + 0.5(0.1)(0.5) + 0.1(0.4)(0.5) \\ &= 0.025 + 0.020 = 0.045. \end{aligned}$$

And there's the rub. The two mixes produce different compound totals, 0.105 versus 0.045, even though the total number of stripes is the same. No summing or averaging of the reaction coefficients of the aggregate's

parts will work, because there is *no* coefficient that will work for both mixes. The nonlinear interactions prevent us from assigning an aggregate reaction rate to the aggregate reaction.

We are now at the end of this particular tale. We've seen that even in the simplest situations nonlinearities can interfere with a linear approach to aggregates. That point holds in general: nonlinear interactions almost always make the behavior of the aggregate more complicated than would be predicted by summing or averaging.

Flows (*Property*)

The idea of flows extends well beyond the movement of fluids. In everyday usage, we speak of the flow of goods into a city or the flow of capital between countries. In more sophisticated contexts, we think of flows over a network of nodes and connectors. The nodes may be factories, and the connectors transport routes for the flow of goods between the factories. Similar {node, connector, resource} triads exist for other *cas*: {nerve cells, nerve cell interconnections, pulses} for the central nervous system; {species, foodweb interactions, biochemicals} for ecosystems; {computer stations, cables, messages} for the electronic Internet; and so on (see Figure 1.8). In general terms, the nodes are processors—*agents*—and the connectors designate the possible interactions. In *cas* the flows through these networks vary over time; moreover, nodes and connections can appear and disappear as the agents adapt or fail to adapt. Thus neither the flows nor the networks are fixed in time. They are patterns that reflect changing adaptations as time elapses and experience accumulates.

Tags almost always define the network by delimiting the critical interactions, the major connections. Tags acquire this role because the adaptive processes that modify *cas* select *for* tags that mediate useful interactions and *against* tags that cause malfunctions. That is, agents with useful tags spread, while agents with malfunctioning tags cease to exist. Later on we will look at this process in some detail.

There are two properties of flows, well known from economics, that are important to all *cas*. The first of these is the *multiplier effect* (see, for

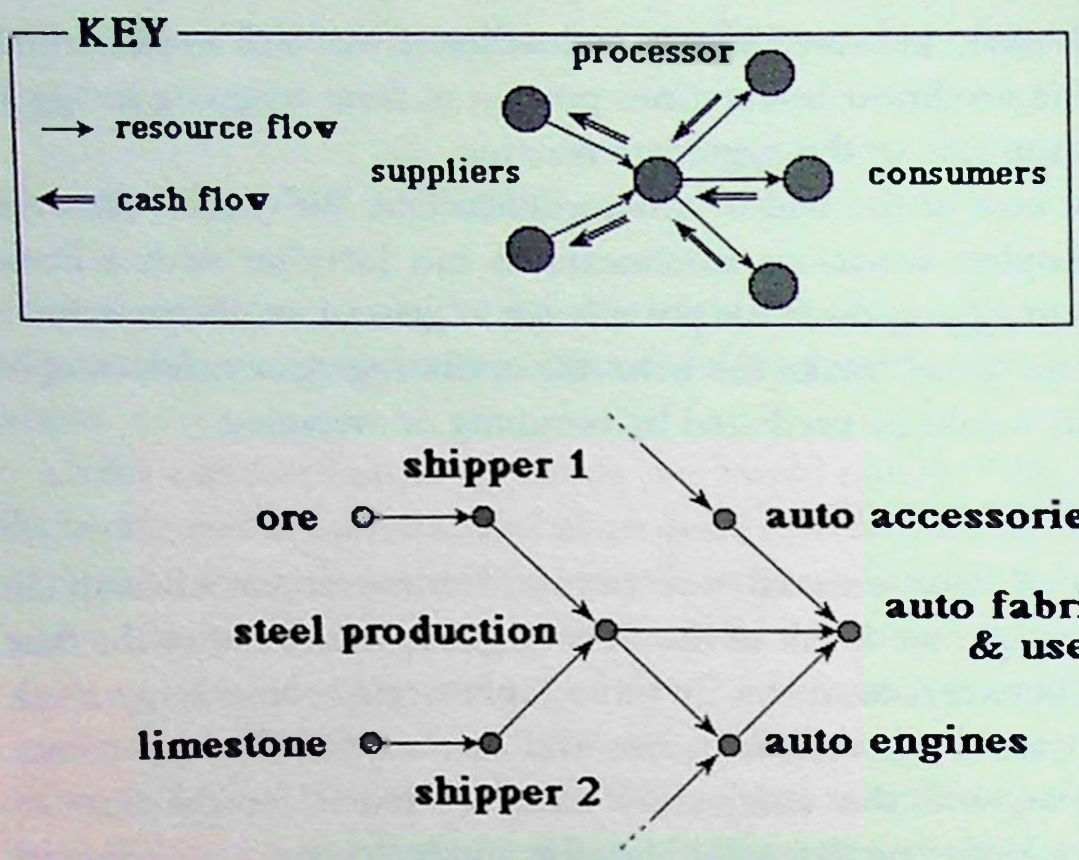


Figure 1.8 Flows.

example, Samuelson, 1948), which occurs if one injects additional resource at some node. Typically this resource is passed from node to node, possibly being transformed along the way, and produces a chain of changes (see Figure 1.9).

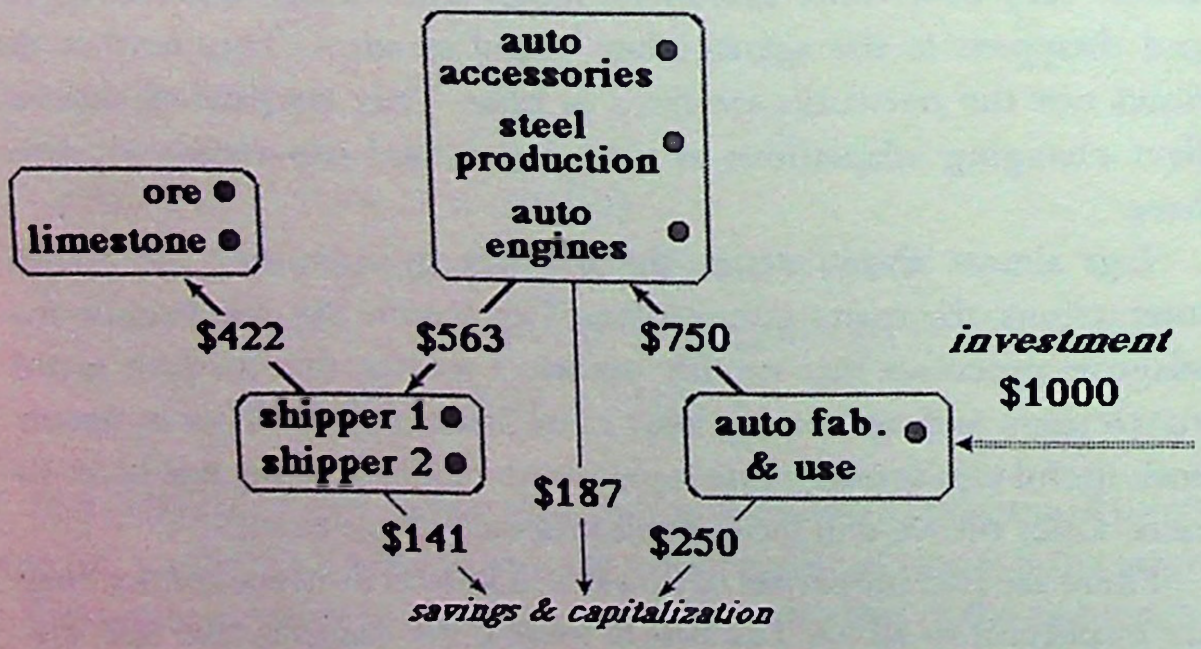


Figure 1.9 Multiplier Effect.

The simplest examples come from economics. When you contract to build a house, you pay the contractor, who pays the tradesmen, who in turn buy food and other commodities, and so on, stage by stage through the economic network. In order to make a simple computation, let's assume that at each stage one-fifth of the new income is saved, and the other four-fifths is paid to the next stage. Then for each dollar you pay, 80 cents will be passed on by the contractor to the tradesmen, who in turn pass on 64 cents, and so on. More generally, a fraction r is passed on at each stage. So at stage 2, a fraction r of the original amount is available. At stage 3, a fraction r of that r is available, so r^2 is available at stage 3. And so it goes for each successive stage. We can calculate the total effect by using the fact that $1 + r + r^2 + r^3 + \ldots = 1/(1 - r)$. In this example, $r = 0.8$, so the total effect is approximately $1/(1 - 0.8) = 5$. That is, the initial effect, your contract, is multiplied by five when we trace its total effect as it passes through the network.

This multiplier effect is a major feature of networks and flows. It arises regardless of the particular nature of the resource, be it goods, money, or messages. It is relevant whenever we want to estimate the effect of some new resource, or the effect of a diversion of some resource over a new path. It is particularly evident when evolutionary changes occur, and it typically jeopardizes long-range predictions based on simple trends.

The second property is the *recycling effect*—the effect of cycles in the networks (see Figure 1.10). This too is most easily explained using an example from economics. Consider a network involving three nodes, say an ore supplier, a steel producer, and a node that stands for auto fabrication and use. For simplicity, we adjust the resource measures (weights) so that one unit of ore produces one unit of steel which in turn produces one unit of automobile. Further, we'll have the steel producer send exactly half its output to the auto fabrication/use node. That is, if the ore supplier ships 1000 units of ore, that will translate through the network to become $0.5(1000) = 500$ units of auto. If we assume that the autos produced are used until they turn into unrecoverable rust, then the return for each 1000 units of ore mined is 500 units of automobile. How do things change if we manage to recycle three-quarters of the steel in autos? Some of the material now goes through a

cycle from the fabrication/use node to the shipper through the steel producer and back to the fabrication/use node. Under this arrangement, with the same 1000 units of ore from the miner, steel production settles down at 1600 units output, which in turn yields 800 units of auto at the fabrication/use node. Recycling, with the same raw input, produces more resource at each node.

That recycling can increase output is not particularly surprising, but the overall effect in a network with many cycles can be striking. A tropical rain forest illustrates the point. The soil there is extremely poor because tropical downpours have a leaching effect that quickly moves resources from the soil into the river system. For this reason ordinary agriculture, which does not recycle resources, fares poorly when the tropical forest is cleared. Yet the forest itself is rich in both species and

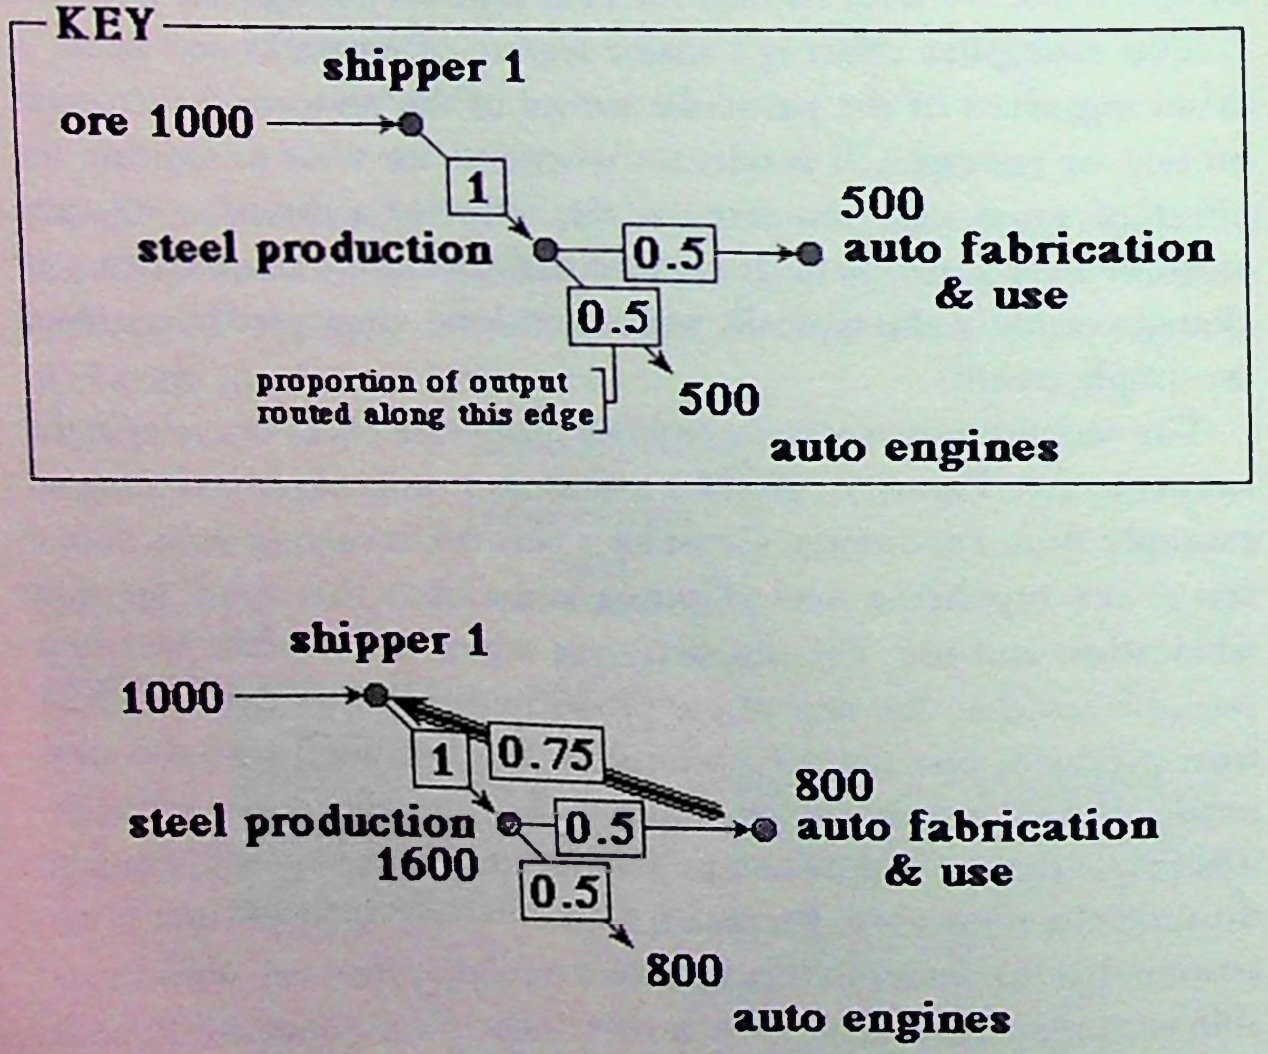


Figure 1.10 Recycling.

numbers of individuals. This state of affairs depends almost entirely on the forest's ability to capture and recycle critical resources. For the forest departs from the simple version of a food web, in which resources are only passed upward to the top predator. Instead, cycle after cycle traps the resources so that they are used and reused before they finally make their way into the river system. The resulting system is so rich that a *single* rain forest tree may harbor over 10,000 distinct species (!) of insects.

Diversity (*Property*)

In that same tropical rain forest, in addition to the diversity of insects, it is possible to walk half a kilometer without twice encountering the same species of tree. But the rain forest is not an isolated example. The mammalian brain consists of a panoply of neuron morphologies organized into an elaborate hierarchy of nuclei and regions; New York City consists of thousands of distinct kinds of wholesalers and retailers; and so it goes for each *cas* in turn.

This diversity is neither accidental nor random. The persistence of any individual agent, whether organism, neuron, or firm, depends on the context provided by the other agents. Roughly, each kind of agent fills a niche that is defined by the interactions centering on that agent. If we remove one kind of agent from the system, creating a "hole," the system typically responds with a cascade of adaptations resulting in a new agent that "fills the hole." The new agent typically occupies the same niche as the deleted agent and provides most of the missing interactions. This process is akin to the phenomenon called *convergence* in biology. The ichthyosaur of the ancient Triassic seas filled much the same niche as the porpoise in modern seas. Though the ichthyosaur is no kin to the porpoise, it is surprisingly similar in form and habit. It even preyed on cephalopods (squid and octopuses). And here you have another convergence. The eye of a squid exhibits all the parts and complexity of a mammalian eye, yet the two are derived from entirely different tissues. The two eyes fill the same niche in different physiologies, a niche determined by the interactions eyes must provide.

Convergence of a kind also occurs when an established species enters

virgin territory. The islands of Hawaii, newly arisen a few million years ago, constituted virgin territory for a pregnant fruit fly (genus *Drosophila*) that drifted or was blown there from elsewhere. Over 600 indigenous species of fruit fly have arisen from that founder. Still more remarkable, these new species fill all sorts of niches that are occupied by very different fly species elsewhere in the world. The ecosystem interactions are largely re-created, although the agents are quite different.

Diversity also arises when the spread of an agent opens a new niche—opportunities for new interactions—that can be exploited by modifications of other agents. Mimicry, a pervasive biological phenomenon, is a good example. In North America the most familiar example of mimicry involves the monarch butterfly (see Figure 1.11). The monarch is marked by a striking orange and black pattern, but it flies quite openly in the fields, unlike most butterflies that flit quickly from cover to cover to avoid predators. The monarch can move so freely because its caterpillar accumulates a bitter alkaloid from the milkweed plant; birds quickly learn that the monarch butterfly induces vomiting. There is a second butterfly, the viceroy, that has a wing pattern almost identical to that of the monarch but it lacks the monarch's bitterness. It mimics the monarch, and thereby gains an important freedom. How can blind chromosomes generate a complicated pattern that mimics the pattern of an entirely different species? It's an important question that we'll look into later, when we have a better foundation. For now we simply note the new niche, and the diversity, provided by the presence of the monarch.

Mimicry exists at every turn in the rain forest. Insects mimic twigs,



Figure 1.11 Mimicry.

snakes, and even bird splat. Orchids mimic a wide range of pollinators so well that, as in the case of the bee orchid, they induce copulatory movements as a means of covering the insect with pollen. The orchid family itself consists of close to 20,000 species, exhibiting an extraordinary variety of shapes and mechanisms (including pollen-throwing and clasping devices). Each new species opens still newer possibilities for interaction and specialization, with still further increases in diversity.

The diversity of *cas* is a dynamic pattern, often persistent and coherent like the standing wave we alluded to earlier. If you disturb the wave, say with a stick or paddle, the wave quickly repairs itself once the disturbance is removed. Similarly in *cas*, a pattern of interactions disturbed by the extinction of component agents often reasserts itself, though the new agents may differ in detail from the old. There is, however, a crucial difference between the standing wave pattern and *cas* patterns: *cas* patterns evolve. The diversity observed in *cas* is the product of progressive adaptations. Each new adaptation opens the possibility for further interactions and new niches.

What mechanisms enable *cas* to generate and maintain temporal patterns with such diverse components? Answers to this question are pivotal to any deep understanding of *cas*. To have a comprehensive theory, we must answer this question in way that applies to all *cas*. A principle from paleontology applies *mutatis mutandis* here: to understand species, understand their phylogeny.

We can make some progress in comprehending the origins of diversity if we revisit flows in the light of this paleontological principle. Note first that the patterns of interaction familiar from ecology—symbiosis, parasitism, mimicry, biological arms races (see Figure 1.12; Dawkins, 1976, is worth reading on this subject), and so on—are all well described in terms of agent-directed flows of resources. Because these interactions have counterparts in other *cas*, we can extend this observation to them as well. From the earlier discussion of recycling, we know that agents that participate in cyclic flows cause the system to retain resources. The resources so retained can be further exploited—they offer new niches to be exploited by new kinds of agents. Parts of a *cas* that exploit these possibilities, particularly parts that further enhance

recycling, will thrive. Parts that fail to do so will lose their resources to those that do. This is natural selection writ large. It is a process that leads to increasing diversity through increasing recycling.

We can further enlarge this view if we add some thoughts about

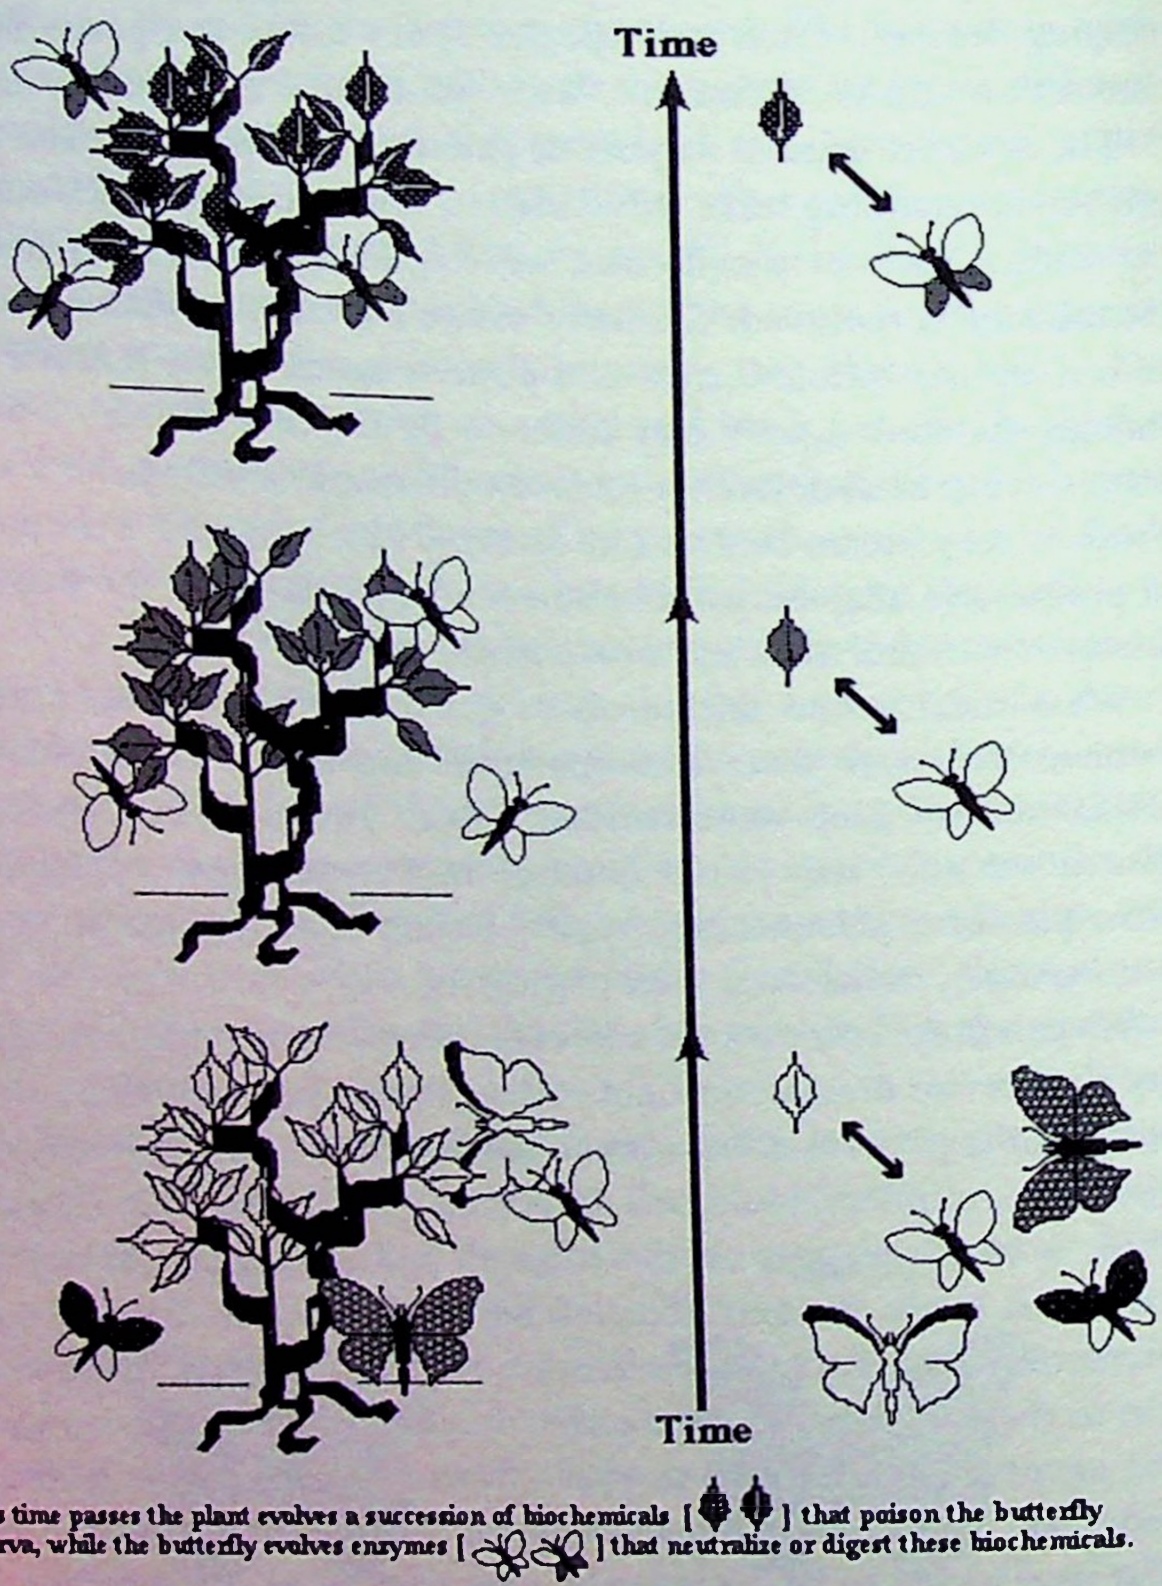


Figure 1.12 A Biological Arms Race.

nonlinearity. The recycling of resources by the aggregate behavior of a diverse array of agents is much more than the sum of the individual actions. For this reason it is difficult to evolve a single agent with the aggregate's capabilities. Such complex capabilities are more easily approached step by step, using a distributed system. This is a point to be emphasized later when we examine the emergence of *default hierarchies* in the next chapter. It should be evident then that we will not find *cas* settling to a few highly adapted types that exploit all opportunities. Perpetual novelty is the hallmark of *cas*.

Internal Models (*Mechanism*)

In introducing mimicry, I mentioned the role of learned avoidance in birds: insectivorous birds anticipate the bitter taste of butterflies with a particular orange and black wing pattern. Just how do they do this? This question, enlarged to encompass all *cas*, takes us to another hallmark of *cas*: they anticipate. To understand anticipation we have to understand a mechanism that is itself complex—an internal model. I use *internal model* to cover much the same ground that Gell-Mann (1994) covers with his *schema*. Unfortunately, the word "schema" has become a fixture in the study of genetic algorithms, designating a related but different topic. Since both topics appear in this book, I choose to avoid confusion by using the term "internal model" to refer to the mechanism for anticipation.

The use of models for anticipation and prediction is a topic that, in its broadest sense, encompasses much of science. It is a difficult topic, but not impenetrable. In the next chapter we will bring out sufficient apparatus to discuss the generation of models, but there are some simpler aspects that we can look at now.

The basic maneuver for constructing models was pointed up in our earlier examination of aggregation: eliminate details so that selected patterns are emphasized. Because the models of interest here are interior to the agent, the agent must select patterns in the torrent of input it receives and then must convert those patterns into changes in its internal structure. Finally, the changes in structure, the model, must enable the agent to anticipate the consequences that follow

when that pattern (or one like it) is again encountered. How can an agent distill experience into an internal model? How does an agent unfold the model's temporal consequences to anticipate future events?

To make a start on these questions, let's take a closer look at models as predictors. We usually ascribe prediction only to "higher" mammals, rather than taking it as a property of all organisms. Still, a bacterium moves in the direction of a chemical gradient, implicitly predicting that food lies in that direction (see Figure 1.13). The mimic survives because it implicitly forecasts that a certain pattern discourages predators. When

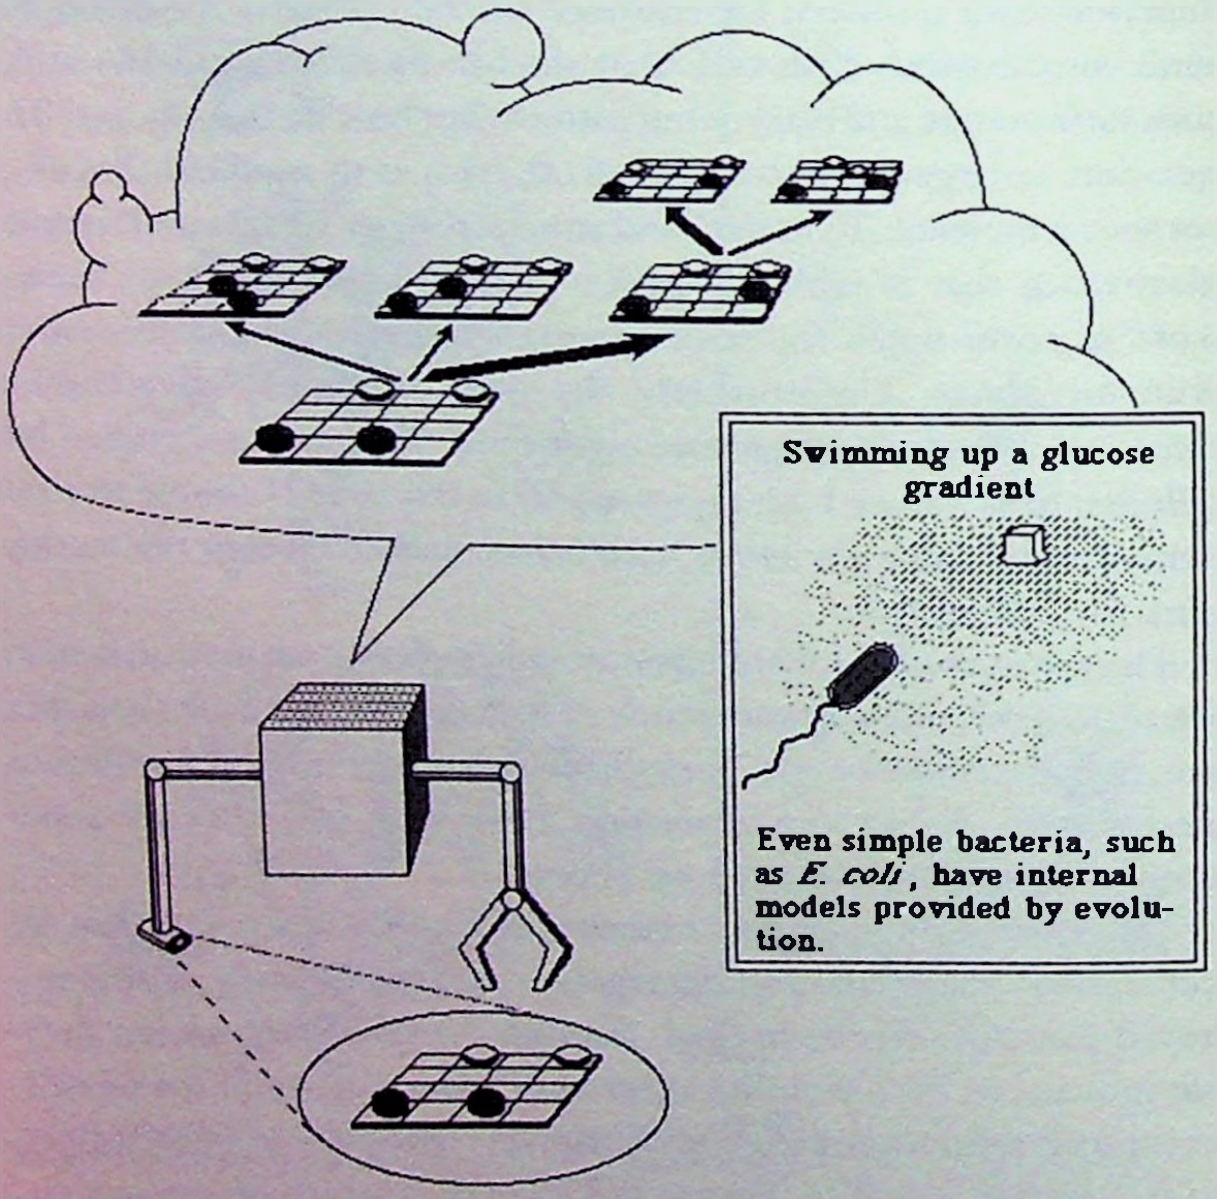


Figure 1.13 Internal Models.

we get to the so-called higher mammals, the models do depend more directly on the agent's sensory experience. A wolf bases its movements on anticipations generated by a mental map that incorporates landmarks and scents. Early humans built Stonehenge as an explicit, external model that helped predict the equinoxes. Now we use computer simulations to make predictions ranging from the flight characteristics of untried aircraft to the future gross domestic product. In all these cases prediction is involved, and in the last two cases external models augment internal models.

Taking these examples into account, we will find it useful to distinguish two kinds of internal models, *tacit* and *overt*. A tacit internal model simply prescribes a current action, under an *implicit* prediction of some desired future state, as in the case of the bacterium. An overt internal model is used as a basis for *explicit*, but internal, explorations of alternatives, a process often called *lookahead*. The quintessential example of lookahead is the mental exploration of possible move sequences in chess prior to moving a piece. Both tacit and overt models are found in *cas* of all kinds—the actions and identity supplied by an immune system fall at the tacit end of the scale, whereas the internal models of agents in an economy are both tacit and overt.

How do we distinguish an internal model from other pieces of internal structure that have nothing to do with modeling? We might start with the critical characteristic of a model: a model allows us to infer something about the thing modeled. Following this line, we could say that a given structure in an agent is an internal model if we can infer something of the agent's environment merely by inspecting that structure. Certainly we can infer a great deal about the environment of any organism by studying relevant pieces of morphology and biochemistry. Accordingly, we might say that those pieces constitute a tacit internal model. But, equally, we can infer a meteorite's history from its composition and surface condition. It is clearly fruitless, even metaphorically, to attribute an internal model to a meteorite, so we need something more in our definition.

There is an additional requirement that will eliminate meteorites and other inanimate structures. We can require that the structure from

which we infer the agent's environment also actively determine the agent's behavior. Then, if the resulting actions anticipate useful future consequences, the agent has an effective internal model; otherwise it has an ineffective one. With an appropriate way of connecting future credit to current actions, evolution can favor effective internal models and eliminate ineffective ones.

Despite the apparent and real differences between the bacterium's tacit model and mammalian overt models, there are important commonalities. In both cases the organism's chances of survival are enhanced by the predictions, implicit or explicit, that the model entails. Thus, variants of the model are subject to selection and progressive adaptation. The timescale for change of the implicit model of the bacterium or the mimic is orders of magnitude different from the timescale for change of a mammal's central nervous system, but the *process* of selective emphasis that generates these models is not so different as we shall see.

Building Blocks (*Mechanism*)

In realistic situations an internal model must be based on limited samples of a perpetually novel environment. Yet the model can only be useful if there is some kind of repetition of the situations modeled. How can we resolve this paradox?

We get the beginnings of an answer when we look to a commonplace human ability, the ability to decompose a complex scene into parts. When we do this, the component parts are far from arbitrary. They can be used and reused in a great variety of combinations (see Figure 1.14), like a child's set of building blocks. Indeed, it is evident that we parse a complex scene by searching for elements already tested for reusability by natural selection and learning.

Because reusability means repetition, we begin to see how we can have repetition while being confronted with perpetually novel scenes. We gain experience through repeated use of the building blocks, even though they may never twice appear in exactly the same combination. By way of example, consider the common building blocks for a human face: hair, forehead, eyebrows, eyes, and so on (see Figure

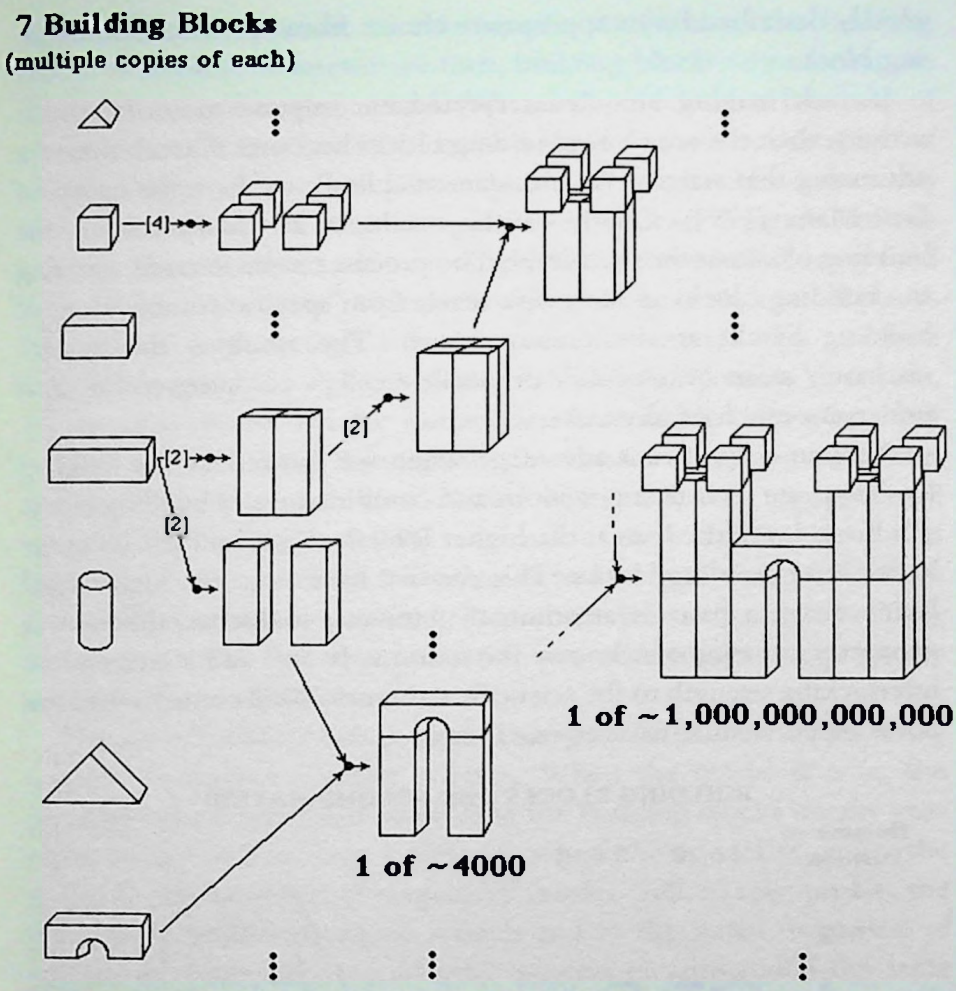


Figure 1.14 Building Blocks (Generators).

1.15). Let's decompose the face into ten components (one of which is "eyes"), and let's allow ten alternatives for each component (as in "blue eyes," "brown eyes," "hazel eyes," . . .). We can think of ten "bags" holding ten building blocks each, for a total of $10 \times 10 = 100$ building blocks. Then we can construct a face by choosing one building block from each bag. Because there are ten alternatives in each bag, we can construct any of $10^{10} = 10$ billion distinct faces with these 100 building blocks! Almost any new face we encounter can be

closely described by an appropriate choice from the set of 100 building blocks.

If model making, broadly interpreted, encompasses most of scientific activity, then the search for building blocks becomes *the* technique for advancing that activity. At a fundamental level, we have the quarks of Gell-Mann (1994). Quarks can be combined to yield nucleons, the building blocks at the next level. The process can be iterated, deriving the building blocks at successive levels from specific combinations of building blocks at the previous level. The result is the quark / nucleon / atom / molecule / organelle / cell / . . . progression that underpins much of physical science.

We gain a significant advantage when we can reduce the building blocks at one level to interactions and combinations of building blocks at a lower level: the laws at the higher level derive from the laws of the lower-level building blocks. This *does not* mean that the higher-level laws are easy to discover, any more than it is easy to discover theorems in geometry because one knows the axioms. It *does* add a tremendous interlocking strength to the scientific structure. We'll come back to this point when we discuss *emergence* in *cas.*

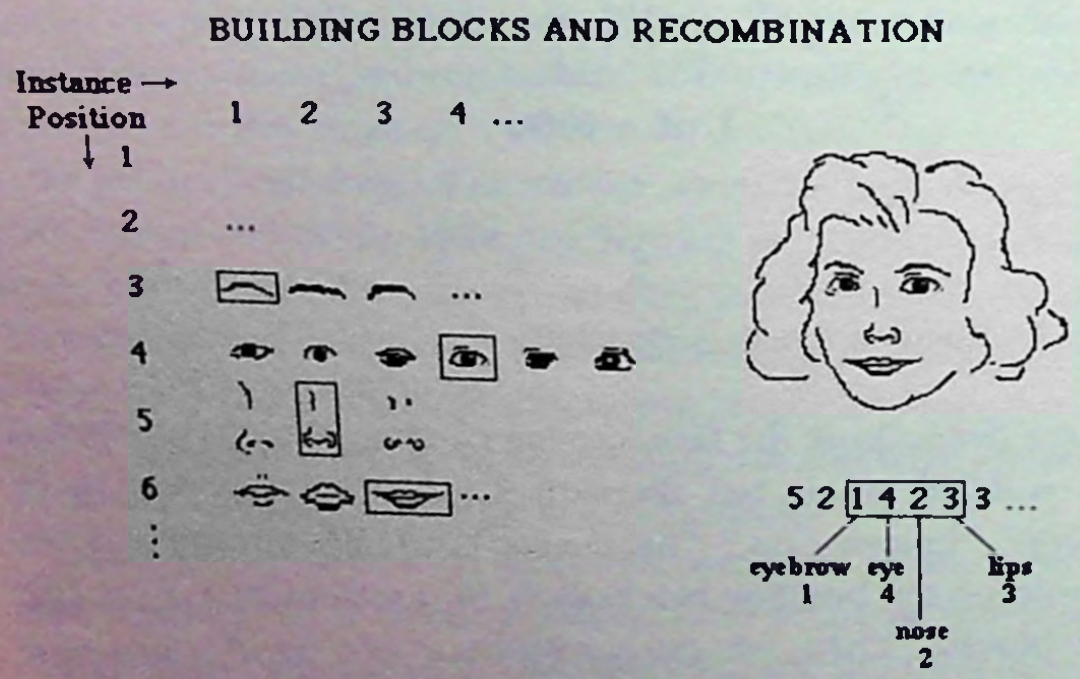


Figure 1.15 Building Blocks for Faces.

It would be a mistake to confine our attention to the building blocks of physics. Wherever we turn, building blocks serve to impose regularity on a complex world. We need only look at the use of musical notation to transmit the endless variety of music, or the use of a limited range of morphologies to describe the tremendous spectrum of animal structures. The point applies with at least as much force to our everyday encounters. If I encounter "a flat tire while driving a red Saab on the expressway," I immediately come up with a set of plausible actions even though I have never encountered this situation before. I cannot have a prepared list of rules for all possible situations, for the same reason that the immune system cannot keep a list of all possible invaders. So I decompose the situation, evoking rules that deal with "expressways," "cars," "tires," and so on, from my repertoire of everyday building blocks. By now each of these building blocks has been practiced and refined in dozens or hundreds of situations. When a new situation is encountered, I combine relevant, tested building blocks to model the situation in a way that suggests appropriate actions and consequences.

This use of building blocks to generate internal models is a pervasive feature of complex adaptive systems. When the model is tacit, the process of discovering and combining the building blocks usually proceeds on an evolutionary timescale; when the model is overt, the timescale may be orders of magnitude shorter. Still, to reemphasize the point made both for internal models and in the initial discussion of adaptation, the underlying adaptive process remains much the same throughout the range of *cas*.

Where Next?

The next three chapters combine these seven basics (see Figure 1.16) in different ways to achieve two goals. The first goal, the object of the next chapter, is to provide a definition of "adaptive agent" that works for all the different kinds of agents found in *cas*. The second goal, to be pursued in Chapters 3 and 4, is to provide a computer-based model that has enough generality to allow us to carry out thought experiments

relevant to all *cas*. We'll see that the seven basics appear over and over again, suggesting mechanisms and directions (see Figure 1.17).

Beyond these two goals is a larger objective: to uncover general principles that will enable us to synthesize complex *cas* behaviors from simple laws. Complex adaptive systems are quite different from most systems that have been studied scientifically. They exhibit coherence under change, via conditional action and anticipation, and they do so

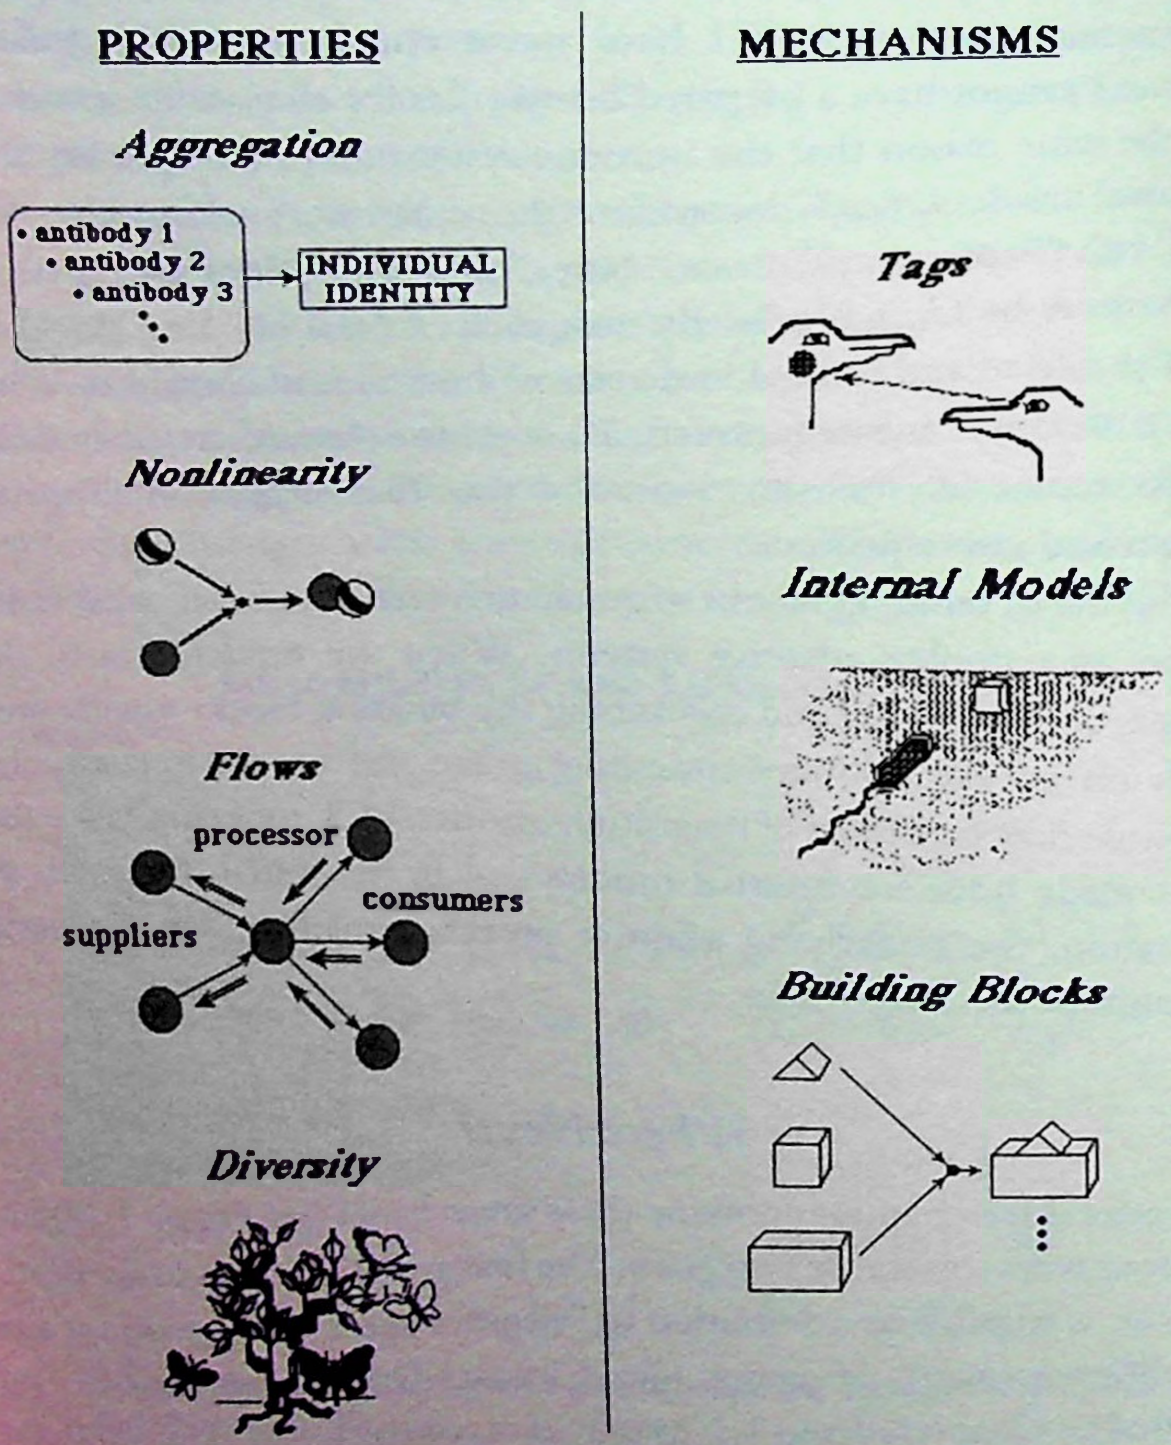


Figure 1.16 Seven Basics for Complex Adaptive Systems.

without central direction. At the same time, it would appear that *cas* have lever points, wherein small amounts of input produce large, directed changes. It should be easier to discover these lever points if we can uncover general principles that govern *cas* dynamics. Knowing more about lever points would, in turn, provide us with guidelines for effective approaches to *cas*-based problems such as immune diseases,

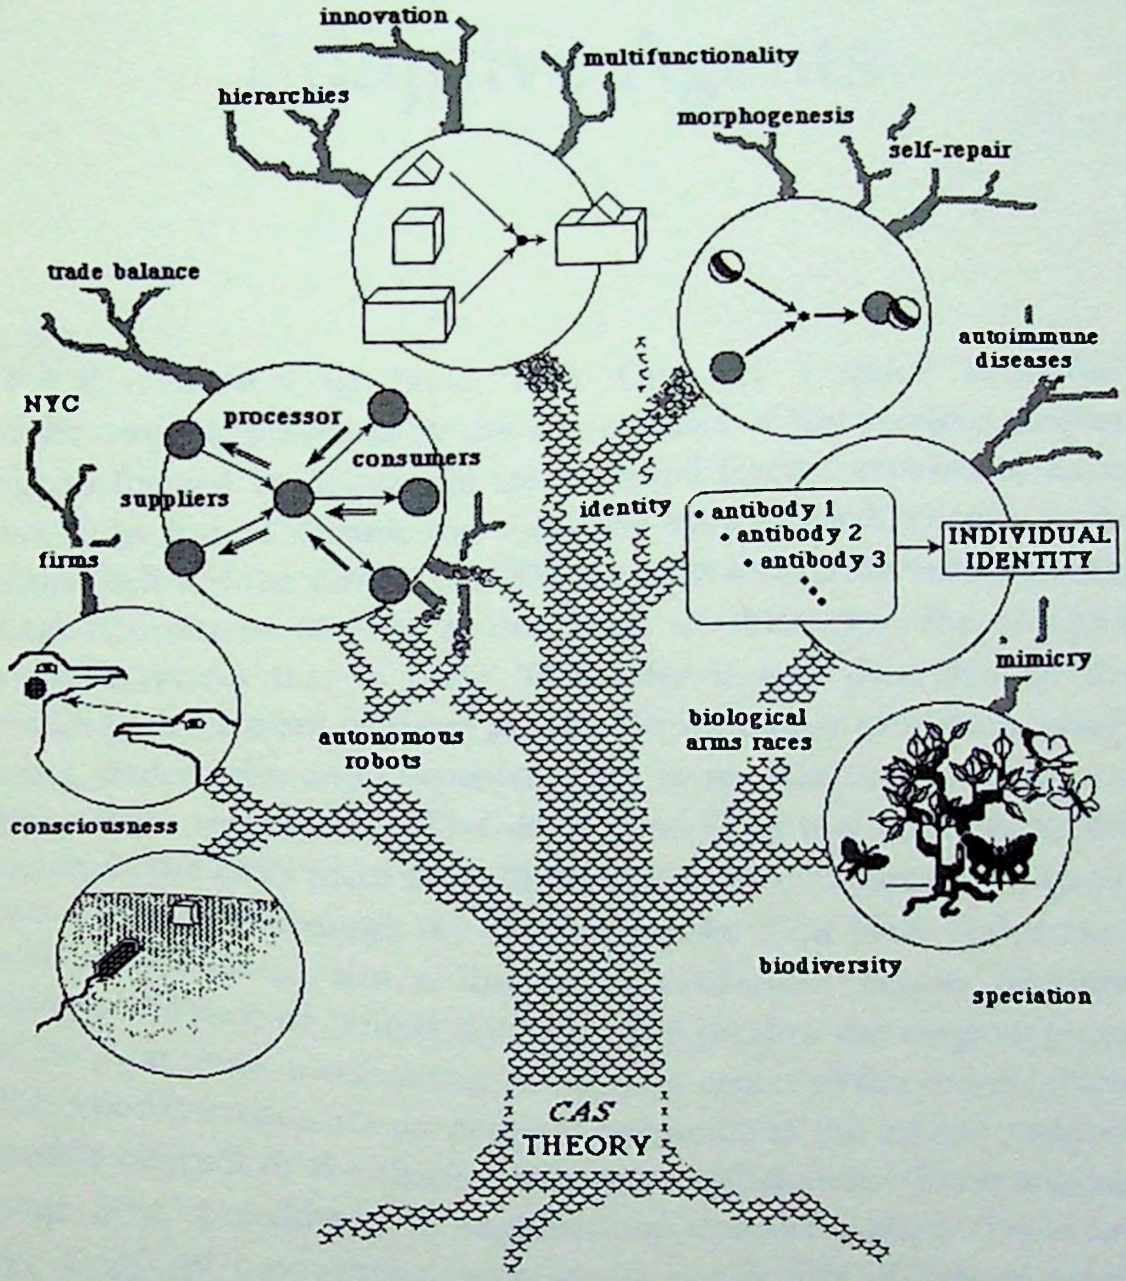


Figure 1.17 The Role of the Seven Basics in the Study of Complex Adaptive Systems.

inner-city decay, industrial innovation, and the like. For problems so complex, it is unlikely that we will make substantial progress without theoretical guidelines. We are only at the beginning of the search for general principles, but we do have some hints as to what those principles might be. I'll set down those hints, as I see them, in the concluding chapter.