

Sex, Culture, and Conflict: The Emergence of History

The basic aim of this chapter is to “grow” a very simple caricature of history—a “proto-history,” if you will. History unfolds on the “twin-peaked” sugarscape familiar from Chapter II, with the sugarscape following the unit growback rule, G_1 . The agents once again move according to rule M but now have other behavioral modes as well, including sex, cultural exchange, and combat. The social story is as follows:

In the beginning, there is a small population of agents, randomly distributed both in space and with respect to their genetic characteristics. Over time spatial agglomeration into two groups occurs as each agent—guided by the primal sugar drive—migrates to one of the two sugar peaks. There, in the midst of plenty, the pioneer agents interact sexually, producing children, who in turn beget children, and so on. All the while processes of cultural evolution are operating within each group producing culturally distinct “tribes” of agents on the two mountains. Ultimately, as population pressures mount from overexploitation of the sugar resources, each tribe spreads down into the central sugar lowlands between the two mountains. When the two tribes ultimately collide, processes of assimilation occur and feed back on the reproductive and cultural activities of the tribes, yielding complex social evolutions.

Our goal, as always, is to grow this history “from the bottom up.” Can the entire social history—along with all sorts of variants—be made to emerge from the interaction of agents operating under simple local rules?¹

In what follows we consider matters of sex first, followed by cultural transmission and the formation of groups, then combat between indi-

1. Since it is not a quantitatively exact story, the “proto-history” admits many realizations. The issue is whether we can grow one of these with the ingredients developed in this chapter.

viduals belonging to different groups, concluding with the promised “proto-history.”

Sexual Reproduction

Can we develop an agent-based demography in which the main dynamics observed in populations *emerge* from the local interactions of the individuals? Minimally, we would like to be able to “grow” the full range of observed aggregate population trajectories, including relatively stable population levels and large oscillations. Variables like “fertility rate,” which are treated as exogenous, often fixed, coefficients in many standard “top-down” demographic models, are in fact highly heterogeneous and should emerge as a result of agent-agent and agent-environment couplings. When fertility rates and population densities begin to interact on the sugarscape, we in fact find that extinction events—central topics of paleontology and evolutionary biology generally—can arise endogenously, without the aid of meteor impacts or other outside agencies. Finally, we expect natural selection to be observable over long times, which indeed it is in Sugarscape.

Imagine that some agent has just arrived at a new sugarscape location as a result of following some movement rule, for example, **M**. After moving, agents are permitted to engage in sexual reproduction with their neighbors. But they must be fertile.

Fertility

First, to have offspring, agents must be of childbearing age. Second, children born with literally no initial endowment of sugar would instantly die. We therefore require that parents give their children some initial endowment. Each newborn’s endowment is the sum of the (usually unequal) contributions of mother and father. Dad contributes an amount equal to one half of whatever his initial endowment had been, and likewise for mom.² To be parents, agents must have amassed at least the amount of sugar with which they were endowed at birth. (In Chapter IV we will permit agents to borrow from other agents to meet this need.) Agents meeting these age and wealth requirements are

2. Agents of the very first generation are assigned random initial endowments.

defined as “fertile.” Each fertile agent executes sex rule **S**, which may be stated algorithmically as follows:

Agent sex rule **S**:

- Select a neighboring agent at random;
- If the neighbor is fertile *and* of the opposite sex *and* at least one of the agents has an empty neighboring site (for the baby), then a child is born;
- Repeat for all neighbors.

It might be that the agent has four neighbors and each is a viable partner. In that case the agent mates with each of them before it is the next agent’s turn to move.³ (However, this possibility is realized infrequently in that agents, if fertile, rarely possess enough wealth to have multiple children in a single period.)

From birth, the baby agent follows **M** just as mature agents do—looking around, accumulating sugar, and so on. The sex of each child is random—males and females are equally likely. The child’s genetic makeup (metabolism, vision, maximum age, and so forth) is determined from parental genetics through Mendelian rules. As the simplest illustration, consider only metabolism and vision and imagine one parent to be of type (m, v) while the other parent is genetically (M, V). Then there are four equally likely genotypes for their child: (m, v), (m, V), (M, v), and (M, V), the combinations given in table III-1.

The agents produced by sexual reproduction are genetically *heterogeneous*. They are *heterogeneous*, too, from the perspective of their environmental attributes; for example, their positions and sugar accumulations. They are *homogeneous* with regard to their behavioral rules, since they all execute {**M**, **S**}. But their *behavior*—as distinct from their *behavioral rules*—is *heterogeneous* since each agent has somewhat different opportunities (depending on which part of the sugarscape it occupies) as well as different abilities (since genetic attributes are parameters of behavioral rules).

Now all of this may seem a cumbersome apparatus. But it is the simplest one we could devise. In fact, it turns out that this set of *local* repro-

3. After observing that the sex rule **S** occasionally yielded multiple births per time period—an outcome we had not considered before running the model—we were tempted to augment the rule with additional conditions such as “females can have only one baby per time period.” However, given our wish to keep the rules as simple as possible, we decided not to add more conditions. Similarly, we have not prohibited sexual relations between close relatives.

Table III-1. Crossover of Genetic Attributes in Sexual Reproduction

Vision	Metabolism	
	m	M
v	(m, v)	(M, v)
V	(m, V)	(M, V)

ductive rules and regulations—the sex code—gives rise to a rich variety of *global*, or macroscopic, population dynamics.

To study demography on the sugarscape it is not particularly revealing to look down on the agents as they gather sugar and reproduce sexually. Rather, a time series of the total number of agents succinctly summarizes the overall dynamics. We present several time series of this type below for various parameterizations of the rule **S**.

Calling each time period of the model a “year,” we first study a population of agents having the following characteristics:⁴

- for both men and women, childbearing begins between the ages of 12 and 15;
- for women, childbearing terminates between the ages of 40 and 50;
- for men, childbearing terminates between the ages of 50 and 60;
- for both men and women, the age of natural death is between 60 and 100;
- members of the initial population have initial endowments in the range 50 to 100 sugar units.

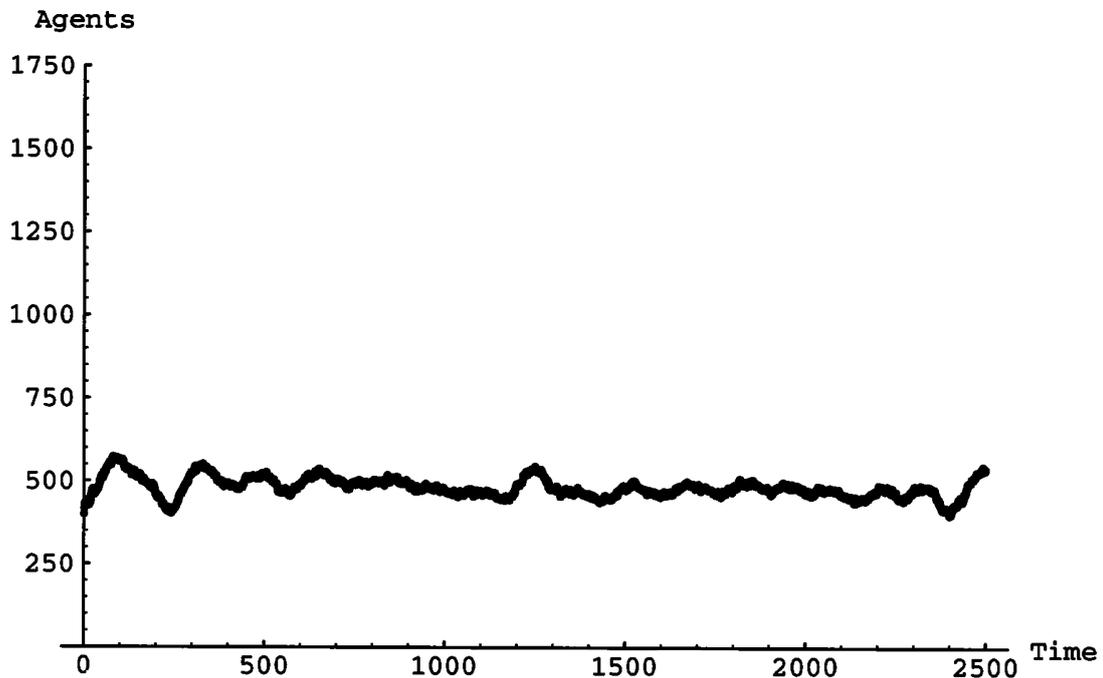
Combining all of this with the movement rule **M** and the sugarscape growback rule **G**₁, population dynamics result. A typical aggregate population time series is shown in figure III-1.

Note that the total number of agents is more or less constant. There seem to be some small quasi-oscillations but these have magnitude less than 10 percent of the overall population level.

While the population is essentially constant in this case it is important to remember that many distinct generations of agents make up figure III-1. That is, the constancy of the total population is actually the result of an approximately *stationary age distribution* of agents. In the artificial society of Sugarscape we can study this distribution directly. Animation

4. For any particular agent the actual value of a parameter is a random variable in the stated range.

Figure III-1. Time Series of Aggregate Population under Rules $(\{G_1\}, \{M, S\})$



III-1 is the age distribution corresponding to figure III-1, changing over time.

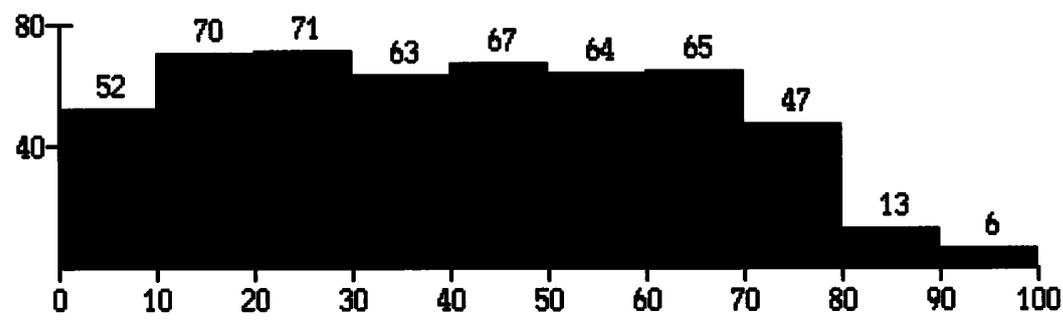
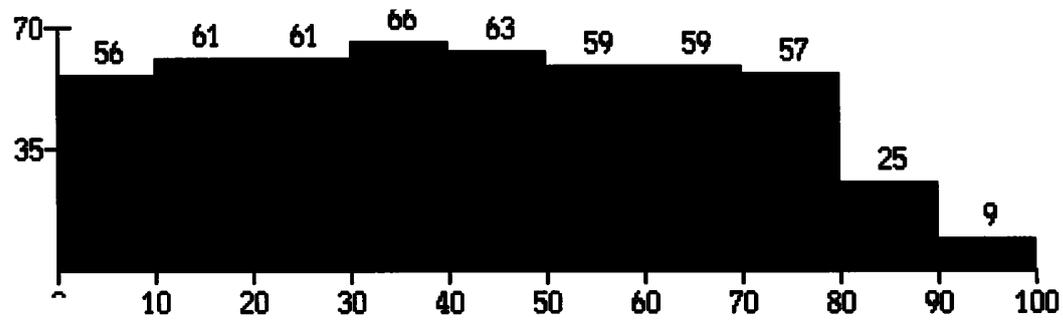
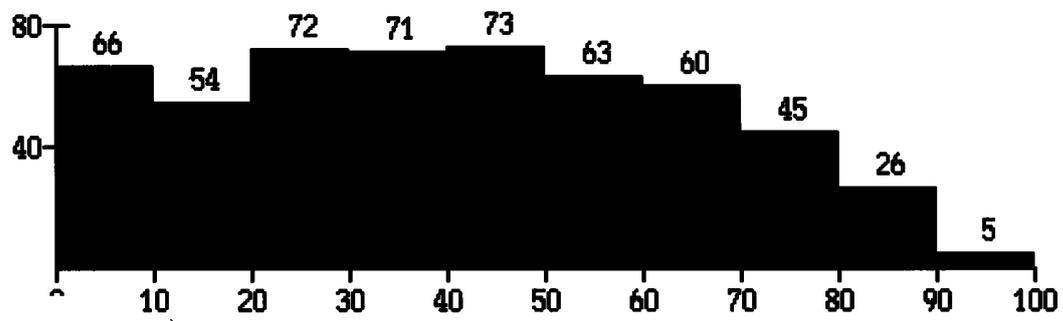
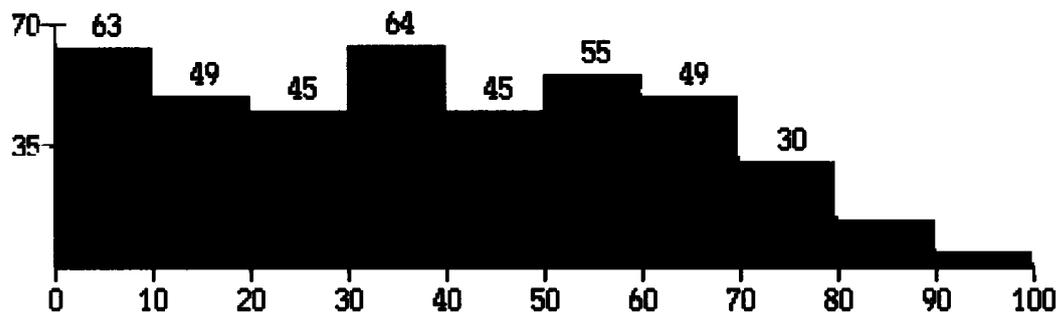
Along the horizontal axis of the figure are age cohort bins, while the vertical axis gives the number of agents falling into each bin. Age distributions highly skewed to the left represent young societies in which there is a high rate of childbirth, while those highly skewed to the right represent aging societies. This age distribution, as mentioned above, assumes an approximately stationary configuration.

Although the total number of agents is approximately constant in this case, it is not true that the characteristics of the agents are unchanging. We explore this presently.

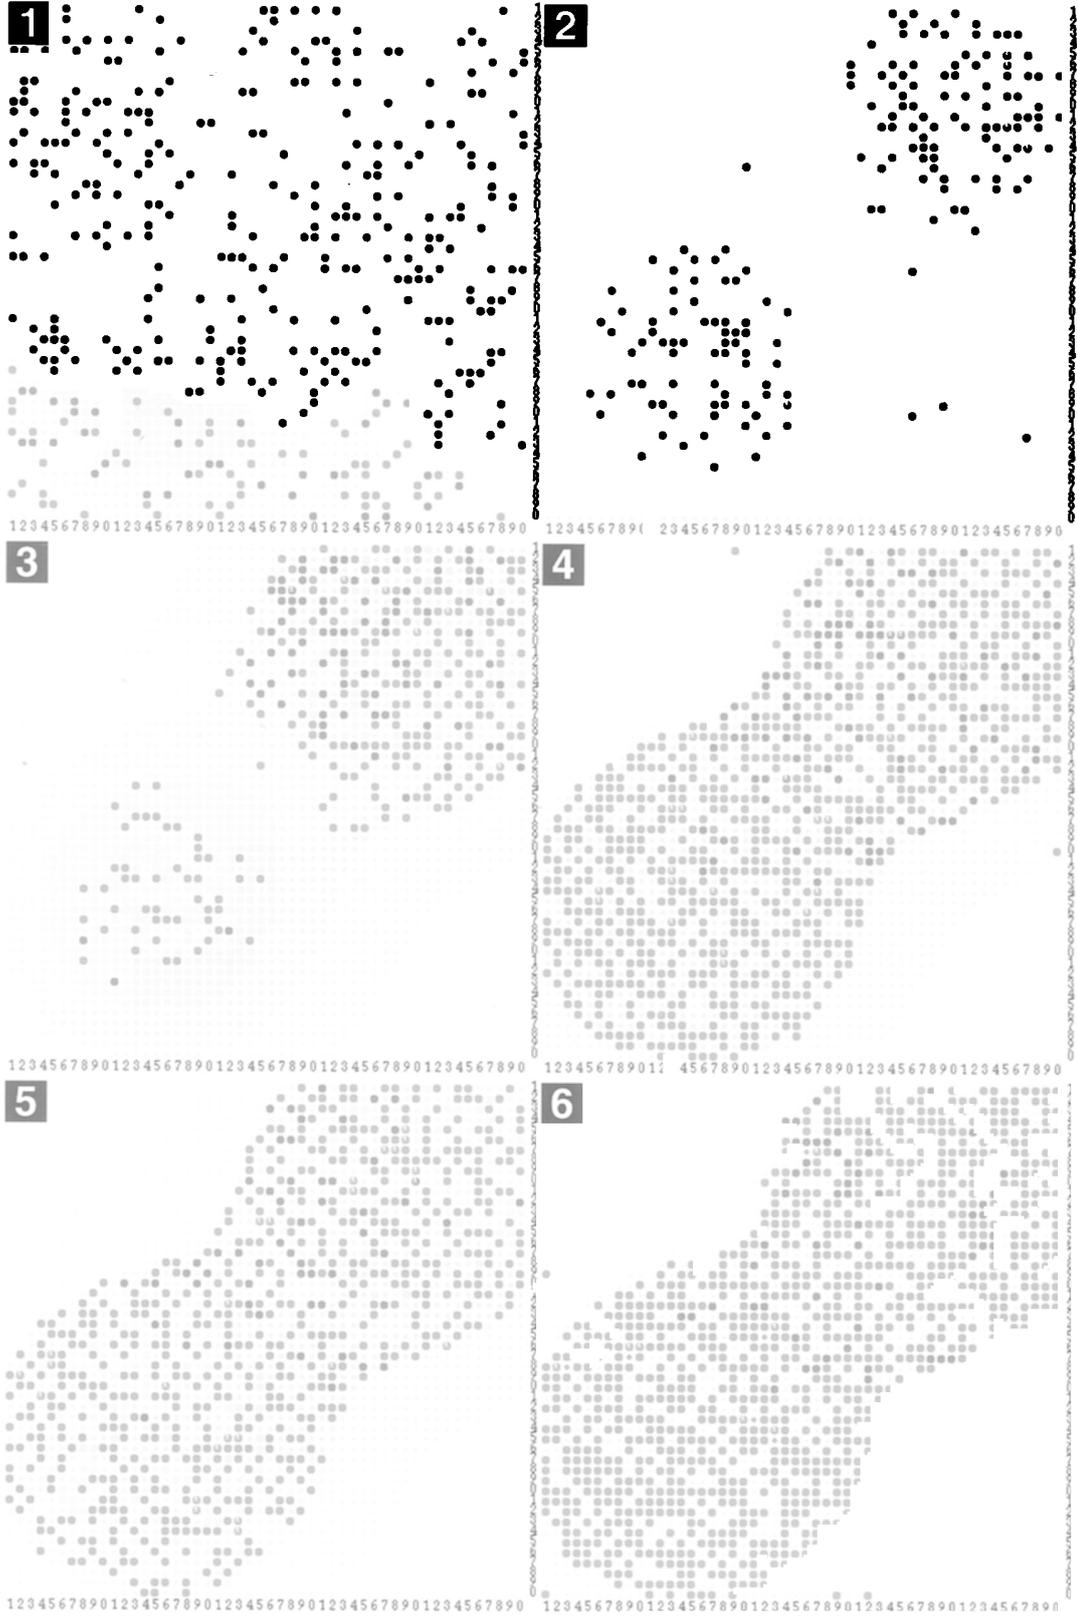
The Theory of Evolution Brought to Life

Agents with relatively low metabolism and high vision enjoy a selective advantage on the sugarscape. Indeed, one can actually watch evolution in action here by coloring agents according to their genetic attributes. First, we color the agents according to their vision. At the start of the run there is a uniform distribution of vision from 1 to 6 among the agents. We color an agent blue if it has vision 1, 2, or 3 and red if it has vision 4, 5, or 6.

Animation III-1. Age Histogram Evolution under Rules $(\{G_1\}, \{M, S\})$



**Animation III-2. Societal Evolution under Rules $(\{G_1\}, \{M, S\})$,
Coloring by Agent Vision**



Watch animation III-2 and note how the colors of the agents change.

At first there are nearly equal numbers of red and blue agents, but over time there is a trend toward more red (high vision) agents. Evolution happens here because of the fertility advantage high vision bestows on agents who have it—high vision yields more sugar income, which begets more children.⁵

Next we do the same thing for metabolism; agents are colored blue if their metabolism is below the initial population mean, red if it is above. Metabolism for sugar is initially uniformly distributed among the agents in the range 1 to 4. But over time evolution eats away at the high metabolism elements of the population, producing, in the end, agents having a uniformly low metabolism. The reader can watch this in animation III-3.

These genetic characteristics—vision and metabolism—affect the ability of agents to survive on the sugarscape in an unambiguous way; *ceteris paribus*, high vision out-gathers low vision and high metabolism makes survival more difficult. Thus the direction in which the mean values of these characteristics move evolutionarily is intuitively clear.⁶ A time series plot of the mean vision and metabolism from animations III-2 and III-3 is given in figure III-2.

This plot makes clear the speed and power of evolutionary processes on the sugarscape. Interestingly, the average metabolism falls faster than the average vision rises. Moreover, the increase in vision—evolutionary “progress”—is not strictly monotone; there are short periods during which the mean actually decreases!

Our population of agents (each of whom possesses its bundle of genetic attributes) can be fruitfully thought of as evolving according to a kind of genetic algorithm. This is so because our sex rule, **S**, involves crossover of the parents’ genetic attributes. However, we have nowhere explicitly defined a fitness function. Rather, we have merely stated rules of reproduction and, from this, endogenous fitness emerges (locally) on the sugarscape. It emerges from agent-environment and agent-agent interactions.⁷

In fact, it is not clear that “fitness” abstracted from all environmental

5. Notice that evolution occurs even though we have *not* defined an exogenous fitness function. We will return to this point.

6. In Chapter IV we will introduce another genetic parameter that has a more ambiguous effect on the ability of agents to survive—a foresight parameter—and we will study its complicated evolution there.

7. On endogenous or intrinsic fitness, see Packard [1989] and Langton [1989: 38].

Animation III-3. Societal Evolution under Rules $(\{G_1\}, \{M, S\})$, Coloring by Agent Metabolism

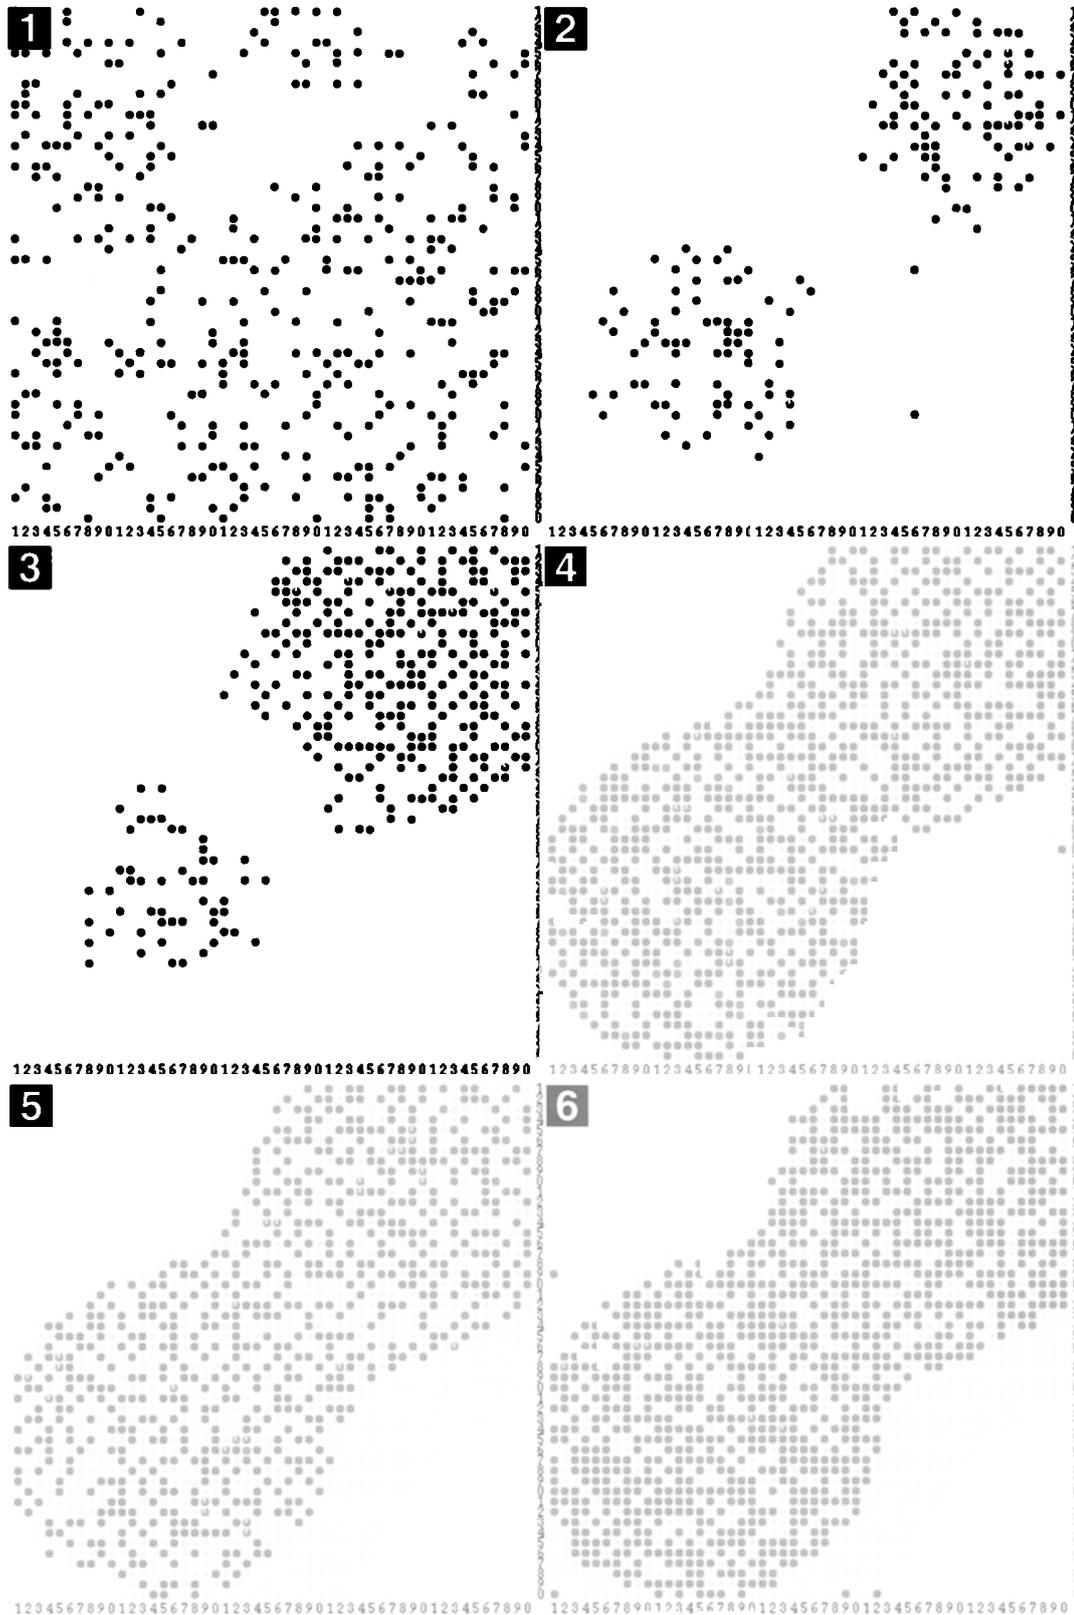
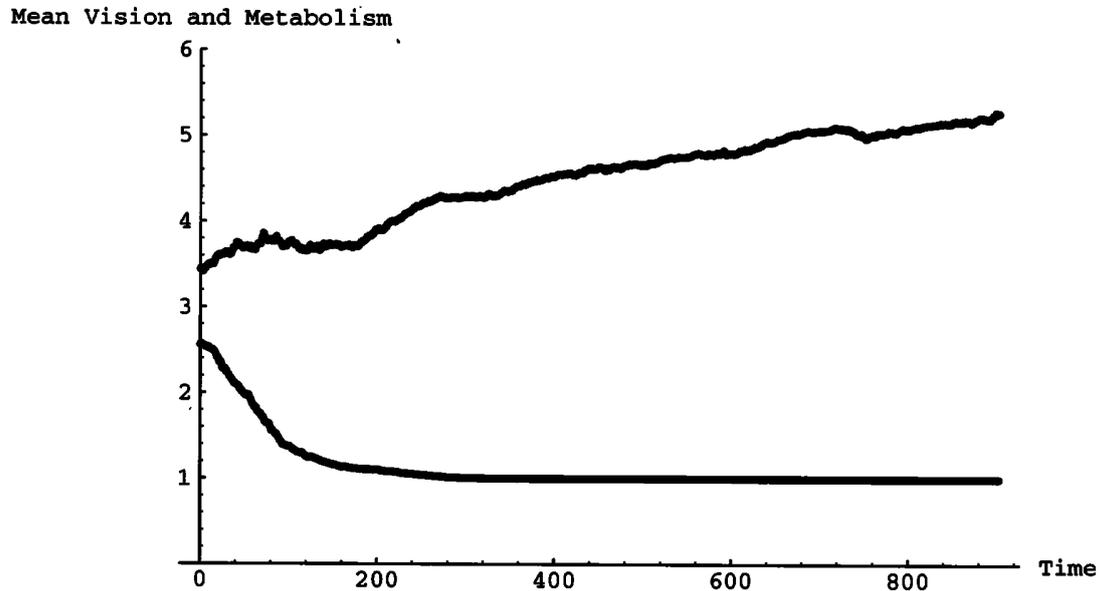


Figure III-2. Evolution of Mean Agent Vision and Metabolism under Rules ($\{G_1\}$, $\{M, S\}$)



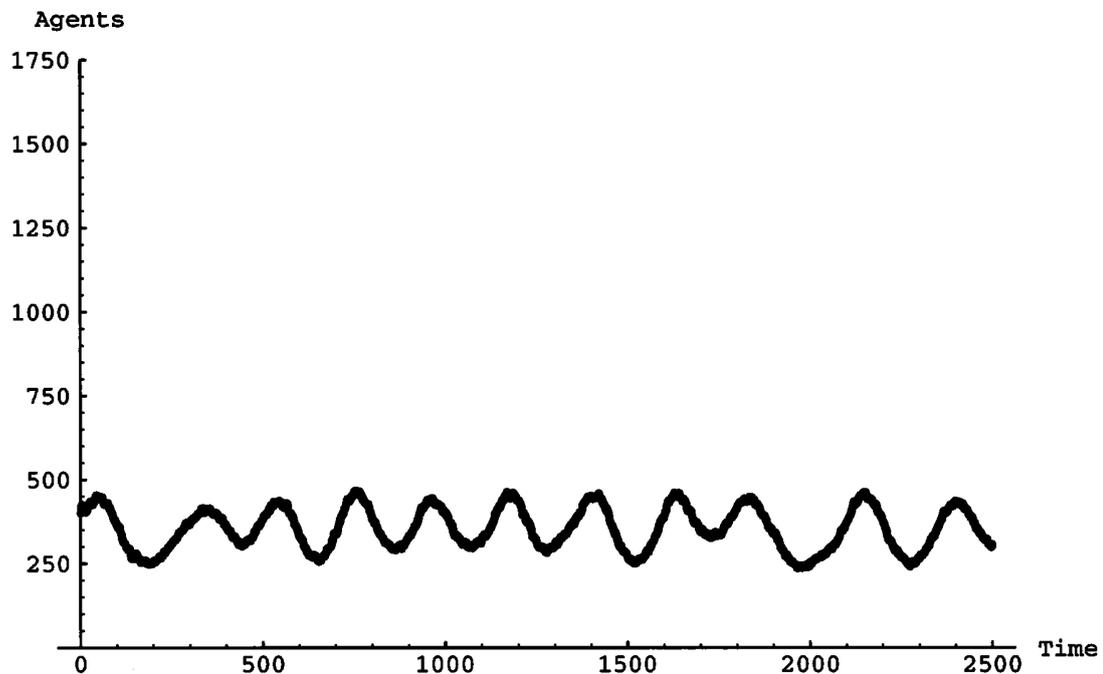
conditions can be usefully defined. To illustrate this point, suppose we—as gods of Sugarscape—simply set the agents’ vision to some high level, boosting their foraging efficiency and, through **S**, their reproductive rates. On one reading, the average fitness in society would appear to be higher than before. But suppose these “fitter” agents bring on their own extinction—through a combination of overgrazing and explosive reproduction! Then it obviously was not so “fit” after all for everybody to have exceedingly high vision.⁸ *Sustainable coevolution with one’s environment* is a necessary condition for “fitness,” if we wish to retain this term at all.⁹ Sugarscape invites us to conceive of fitness as another emergent property, not as something—such as vision—that can be determined by inspection of individuals in isolation.

Our concerns in this section have been with the genetic composition of the population. What of population levels? This is of course a crucial issue for policy and a major ingredient of the “proto-history.” What

8. Ackley and Littman [1992] perform a related experiment in their artificial ecology.

9. On conceptual difficulties surrounding the notion of fitness and related issues, see Sober [1994] and Cohen [1985].

Figure III-3. Small Amplitude Population Oscillations under Rules $(\{G_1\}, \{M, S\})$



determines population growth and how does it interact with other variables, such as fertility? We now return to population dynamics proper.

Regimes of Population Dynamics

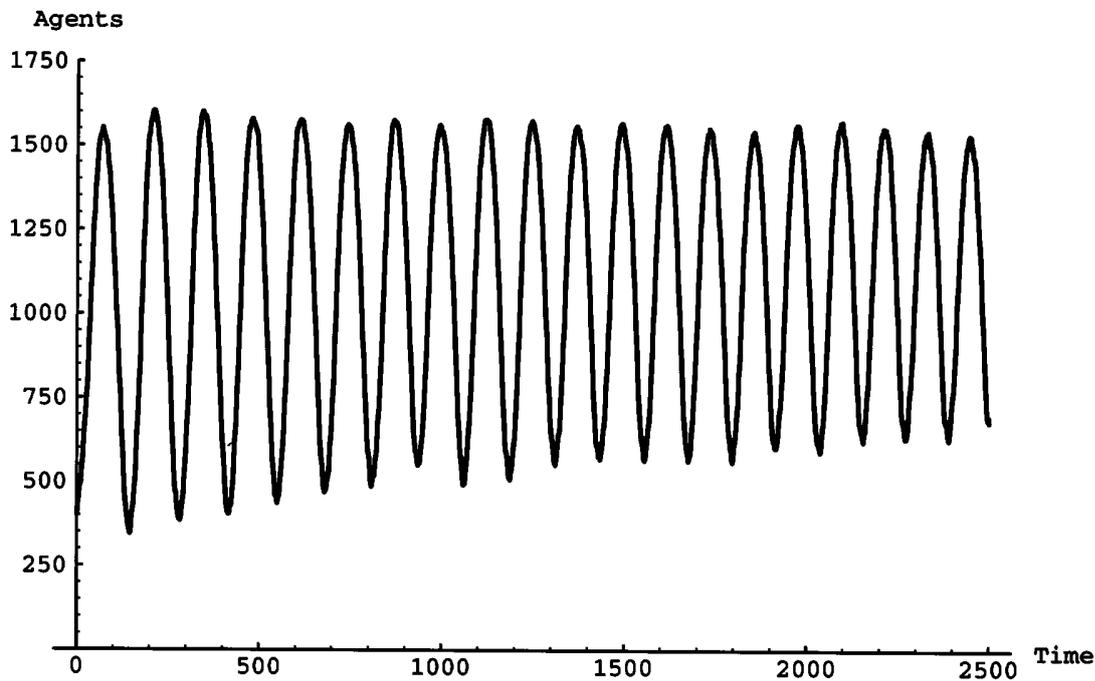
More interesting than the constant population level of figure III-1 and the corresponding stationary age distribution shown in animation III-1 are situations in which birth rates and population size may vary periodically. This leads to a distribution of ages oscillating between two limiting distributions. It turns out that such an outcome can be produced by employing sex rule **S** with a single change to the parameters: For both men and women, fertility terminates 10 years earlier than before; that is,

- for women, childbearing terminates between the ages of 30 and 40;
- for men, childbearing terminates between the ages of 40 and 50.

In this case the sex rule produces the time series given in figure III-3.

Note the fairly regular oscillations. These are bounded in amplitude, and the population never falls below about 250. The period of the oscillations appears to be around 200 years. Perhaps the most interesting feature of this cyclic aggregate behavior is that it is produced completely “from the bottom up”—through the individual actions and interactions

Figure III-4. Large Amplitude Population Oscillations under Rules $(\{G_1\}, \{M, S\})$



of myriad agents.¹⁰ Furthermore, given that the average agent's lifetime is approximately 80 years, the 200-year period of the oscillation means that no single agent ever participates in an entire population swing.

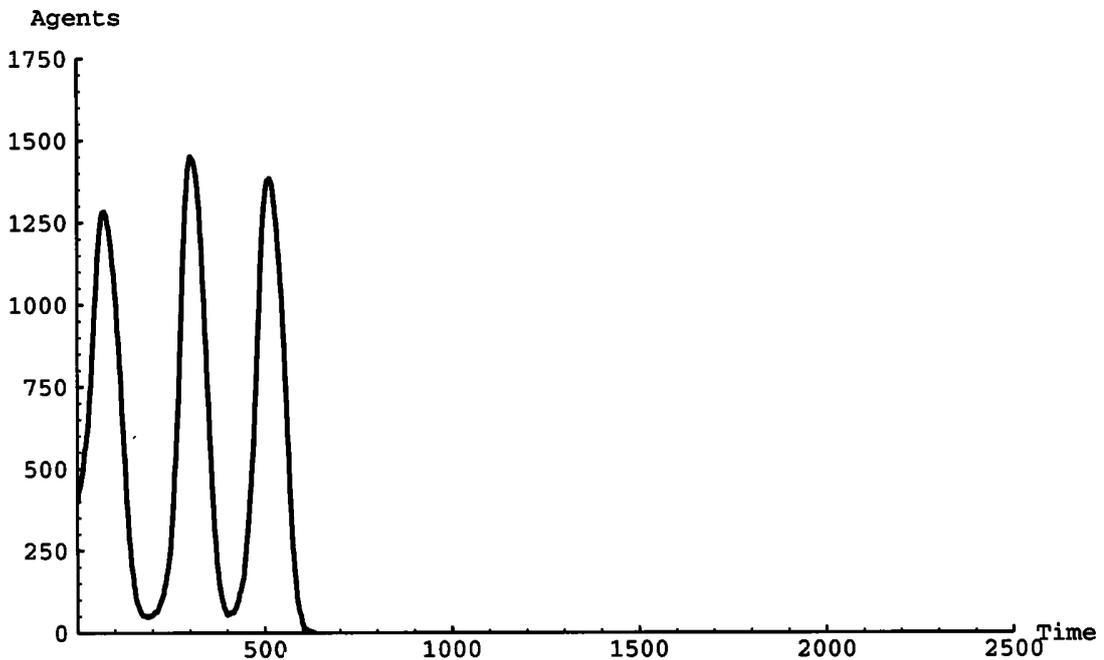
There are other ways to produce population oscillations on the sugarscape. If we restore the previous limits on the ages at which infertility sets in but reduce the amount of sugar necessary for reproduction to the range 10 to 40, we get the population trajectory depicted in figure III-4.

Now the amplitude of oscillation starts out at over 1000, gradually diminishing to approximately 800. The frequency of these oscillations is approximately 125 years. Compared with the previous oscillations these are much more energetic.

Finally, when the variations in S that produced figures III-3 and III-4 are combined—decreased duration of fertility and decreased wealth requirements—the population may oscillate so severely that, at some point, a minimum population occurs that has too little diversity or is spatially too thin (that is, permits too little mating) to re-initiate a cycle of growth. Such an outcome is shown in figure III-5, where the society suffers extinction after completing three growth surges.

10. Over the 2500 time periods represented in figure III-3, some 12,500 agents inhabit the sugarscape.

Figure III-5. Severe Population Swings Leading to Extinction under Rules $(\{G_1\}, \{M, S\})$



In many fields there are ongoing debates concerning the role of endogenous (internally generated) versus exogenous (externally imposed) factors in explaining important phenomena. In paleontology and related fields there are longstanding exogenous-endogenous debates about extinction events. Do they require external shocks, or can the internal dynamics of the system itself bring them about?¹¹ Our agent-based modeling suggests that internal dynamics alone are *sufficient to generate* cataclysmic events.

We summarize the various regimes of population dynamics described above in figure III-6.

Clearly, a great variety of macroscopic dynamics can be produced by this simple model. In subsequent chapters we will study ways in which population dynamics are coupled to other social processes.

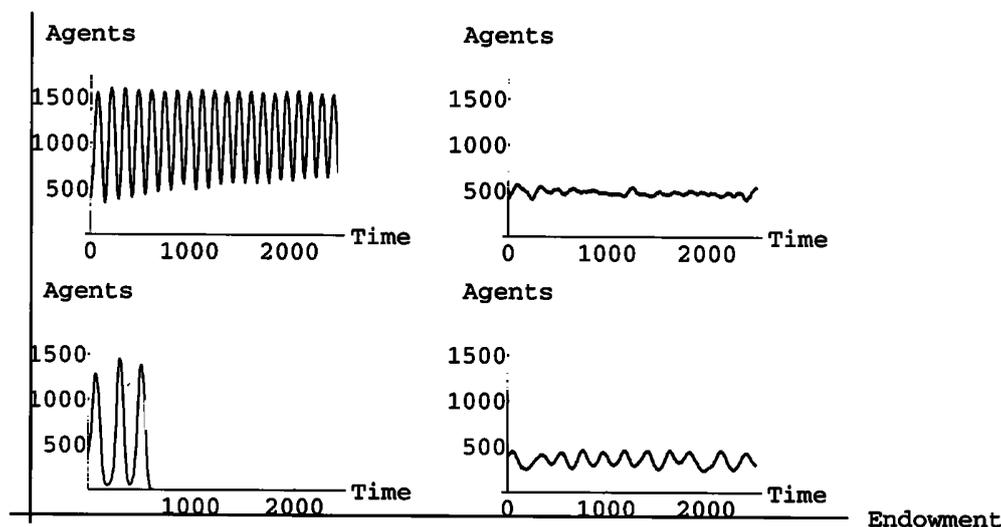
Nature and Nurture: The Genetic Effect of Inheritance

What is the relationship of social institutions—such as property rights—to processes of biological evolution? This is the type of interdisciplinary question that traditional fields do not normally address. Biology does not

11. Recently Hastings and Higgins [1994] have obtained very complex population dynamics in a spatially distributed ecological model.

Figure III-6. Regimes of Population Dynamics under Rules $(\{G_1\}, \{M, S\})$ for Various Parameters

Infertility Age



include economics and vice versa. But in an artificial society we can study biology and economics at once.

In the runs described above, when an agent dies all its wealth simply disappears. We saw evolution at work and watched as average metabolism fell and average vision rose over time. Now let us allow agents to pass their accumulated holdings of sugar on to their offspring when they die; that is, we permit inheritance. Formally, this involves defining a new rule of agent-agent (parent-child) interaction.

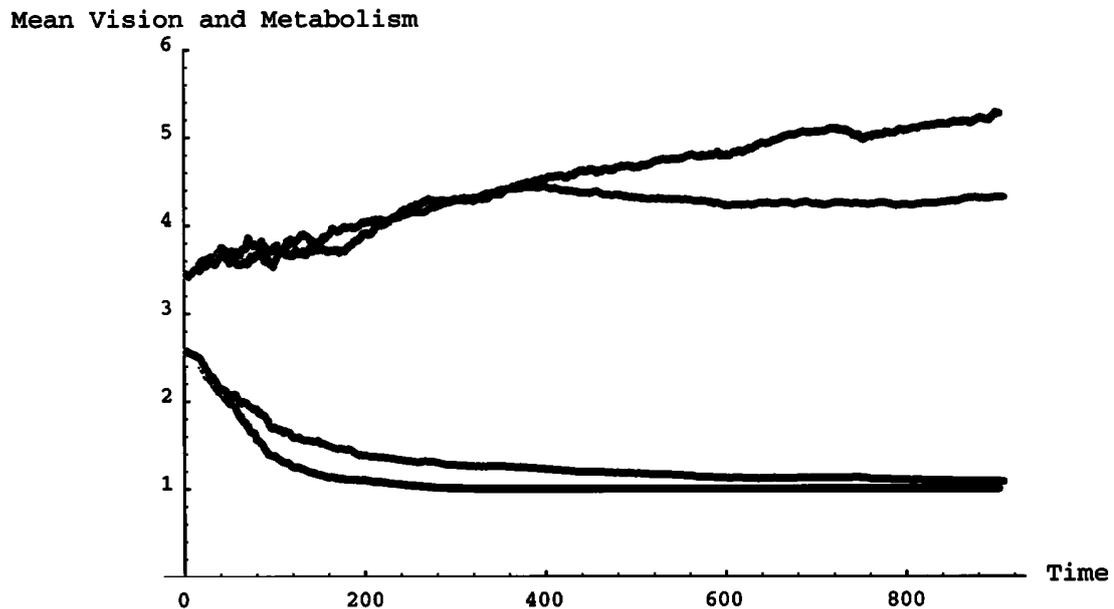
Agent inheritance rule I: When an agent dies its wealth is equally divided among all its living children.¹²

How does this *social* convention affect *biological* evolution—in particular, what happens to the trajectories of average metabolism and average vision over time? We have rerun the model under exactly the conditions that produced figure III-2, only now letting rule I be active. Time series for average vision and metabolism have been overlaid in red on those of figure III-2 in figure III-7.

In the case of vision the message is clear: *Inheritance retards selection.* Agents who might otherwise have been “weeded out” are given an

12. In the Sugarscape software system, other inheritance schemes may be selected by the user, including division of wealth among sons, daughters, or friends.

Figure III-7. Evolution of Mean Agent Vision and Metabolism under Rules $(\{G_1\}, \{M, S, I\})$



advantage through inheritance. However, it does not seem that inheritance has a comparable effect on metabolism, given how the two (lower) curves for this genetic attribute ultimately meet.

Interestingly, some “Social Darwinists” oppose wealth transfers *to the poor* on the ground that the undiluted operation of selective pressures is “best for the species.” Conveniently, they fail to mention that intergenerational transfers of wealth *from the rich to their offspring* dilute those very pressures.

Inheritance raises the Gini coefficient in society—inequality grows under inheritance. This is shown in animation III-4.

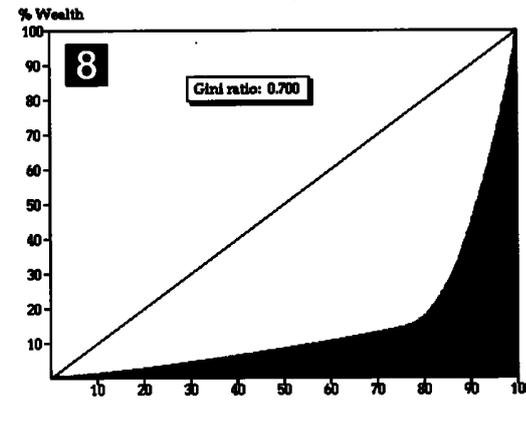
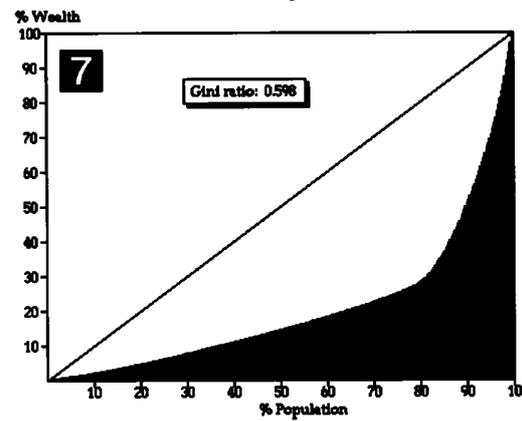
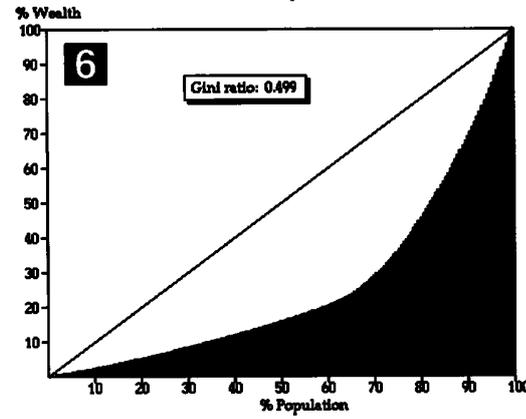
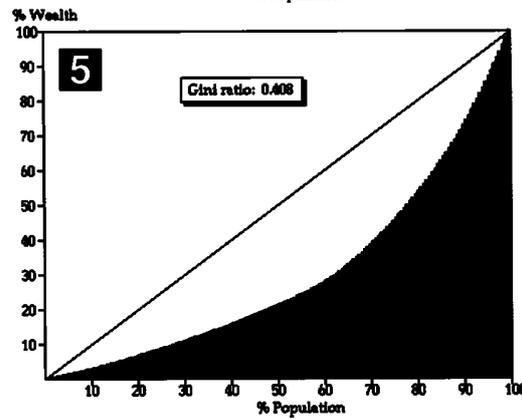
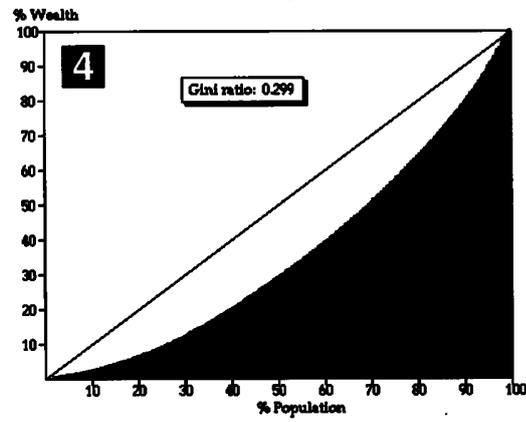
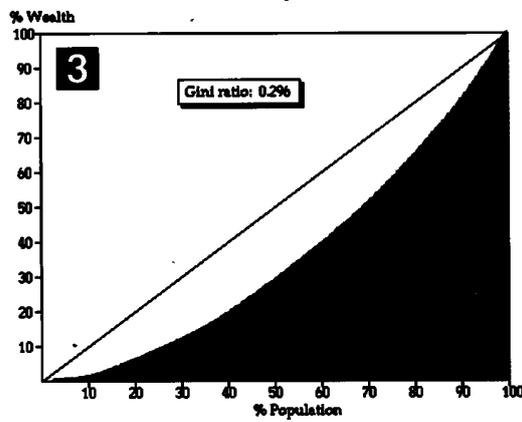
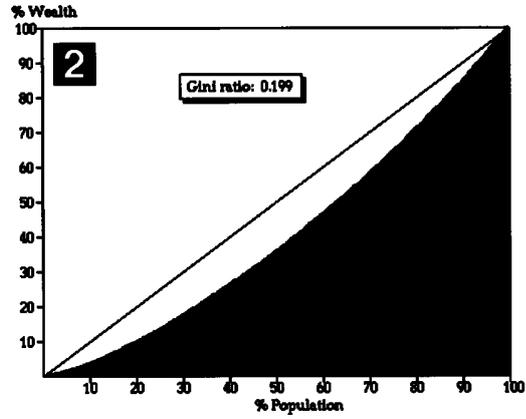
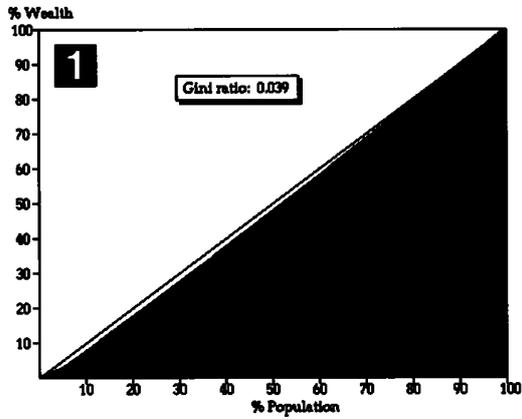
In the previous chapter we noted that the Gini ratio produced by the rule system $(\{G_1\}, \{M, R_{[60,100]}\})$ was small in comparison to that of real economies. With inheritance the Gini ratio is far higher, reaching as high as 0.743 here.¹³

Genealogical Networks

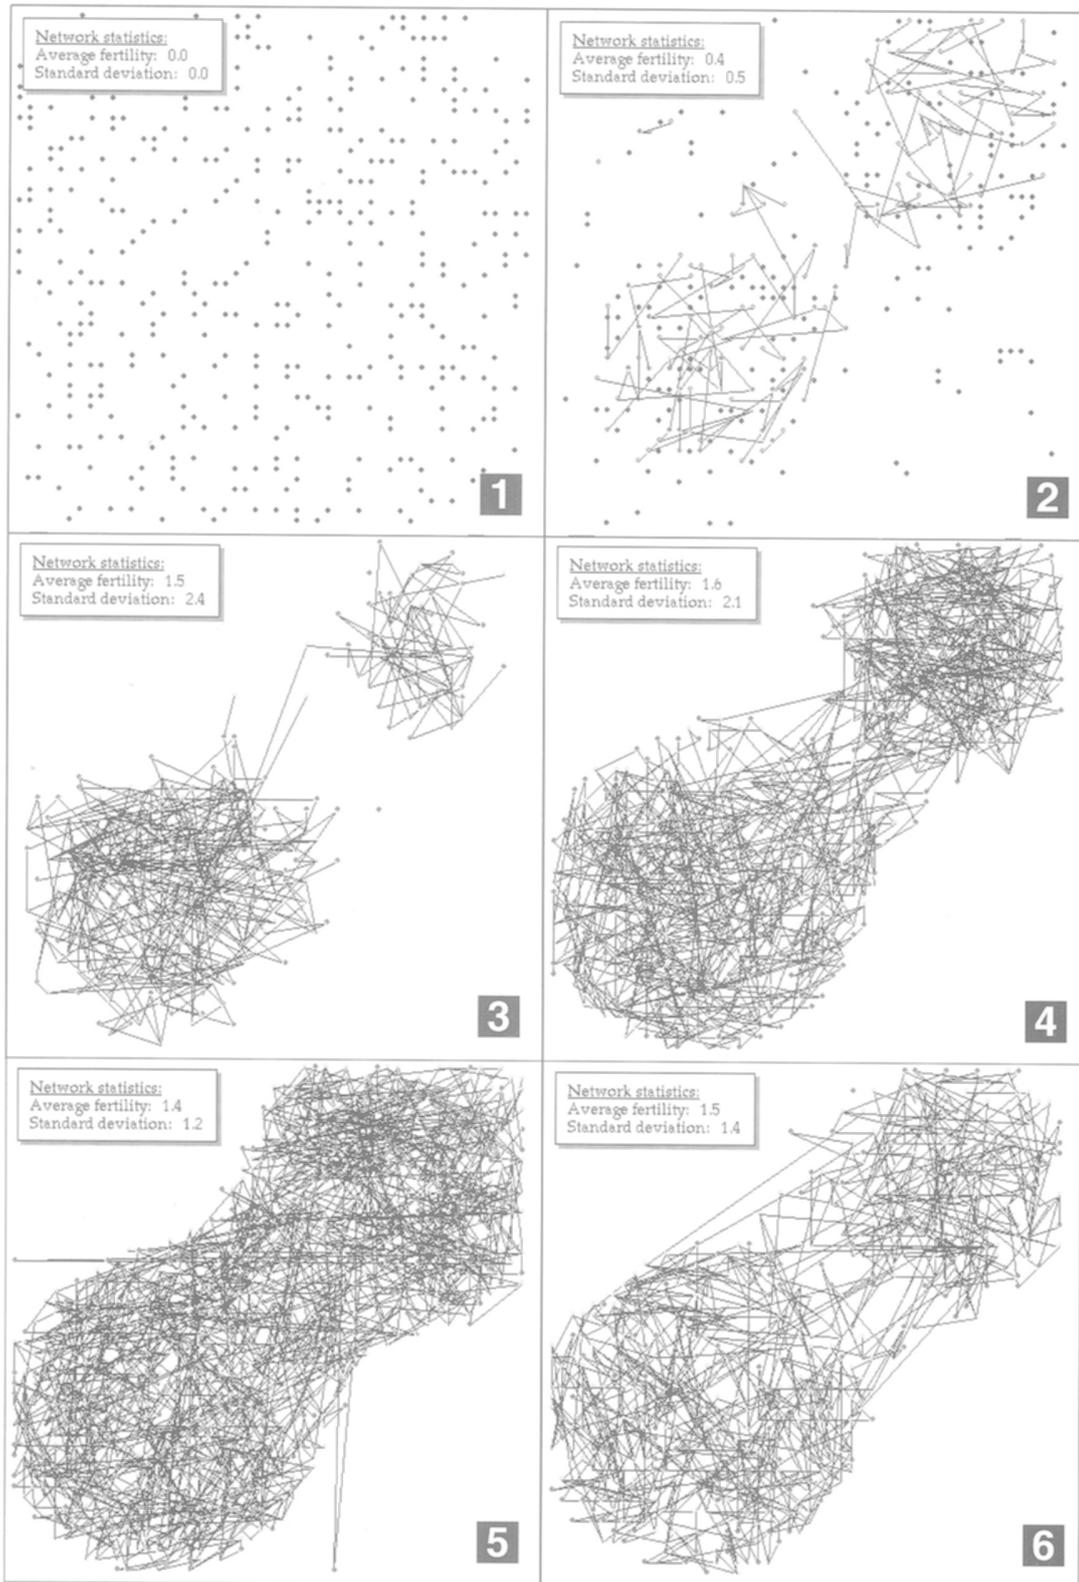
In Chapter II a neighborhood network was defined. Here the sex rule provides a natural basis for a well-known social network, the “family

13. Brittain [1977, 1978] analyzes inheritance data for the United States.

Animation III-4. Evolution of the Lorenz Curve and Gini Coefficient under Rules ($\{G_1\}$, $\{M, S, I\}$)



Animation III-5. Evolution of Genealogical Networks under Rules $(\{G_1\}, \{M, S\})$



tree.” This is depicted in animation III-5 by drawing a line from every parent to each of its children.¹⁴ The initial population is colored black. When a member of this population has a child, the new parent is colored red, the child green. Agents who are both parents and children are colored yellow. (Note that members of the initial population can never be yellow.)

It is interesting to watch the evolution of such genealogical networks. At the outset there are no connections since, in the initial population, no agents are related. However, after several generations, when none of the initial agents remains, all the agents on the sugarscape have some definite genealogical lineage. Notice that average fertility and the standard deviation in fertility both vary substantially in the course of this run.

We have made an initial foray into agent-based demography. The range of phenomena obtained is heartening. Clearly, this is a rich area and our efforts barely scratch the surface. With this first ingredient of the proto-history—sexual reproduction and endogenous population dynamics—in hand, we now proceed to the second, the formation of cultural groups.

Cultural Processes

Our simple agents may not yet seem quite human since all they do is move, eat, and procreate. In this section we give our agents internal states representing cultural factors and augment their behavioral repertoire with simple local rules for cultural interchange. This proves sufficient to produce agent populations having dynamic, heterogeneous cultures. Then, given that any two agents may be either similar or different culturally, it makes sense to talk about distinct cultural formations or tribes of agents. Indeed, we will “grow” such tribes here “from the bottom up.”¹⁵

14. This is implemented by, among other things, having each agent keep pointers to all of its children.

15. Axelrod [1995] studies a bottom-up model of culture in which the agents maintain fixed positions on a two-dimensional lattice. Axtell *et al.* [1996] discuss an implementation of Axelrod’s culture model in Sugarscape and the usefulness of such “docking” experiments for agent-based social science.

Cultural Tags

Recall that every agent is born with a genetic endowment: a metabolism, a vision, a sex, and so forth. Although the distribution of these genetic attributes changes from generation to generation, the genetic makeup of any particular agent is fixed over its lifetime. Of course, in reality, important attributes (for example, tastes) *do* change in the course of one's life.¹⁶ We wish to capture processes of this sort. So beyond its fixed genetic endowment, each agent is born with a structure that represents its cultural attributes. This is a string of zeros and ones.¹⁷ The length of this nongenetic string is the same for all agents.¹⁸ For example, an agent might have a cultural string consisting of 10011010011. We will refer to each element of the string as a tag and will often call the entire structure a "tag string," or simply the agent's "tags." Agents can change one another's tags, which causes the distribution of tags in society to change over time.

Cultural Transmission

Consider an agent who has just landed at some site on the sugarscape. That agent—let us call her Rose—has up to four von Neumann neighbors (as discussed in Chapter II). For illustration, imagine she has two; call them A and B. Cultural transmission might proceed in a great variety of ways. We will adopt the following tag-flipping scheme. First, a neighbor is selected, say neighbor A. Then, one of Rose's tag positions is selected at random. Suppose it is position six and suppose Rose has a 1 at that position—a cultural tag of 1. Then, if neighbor A has a tag of 0 at that position (its position six), it gets flipped to Rose's value of 1. If, at that position, neighbor A already matches Rose, no flip occurs. Now Rose moves on to neighbor B. Again, one of Rose's tag positions is

16. In fact, there is a longstanding debate in economics as to whether or not preferences for commodities change during one's life. This is a topic to which we will return in Chapter IV, where we use the cultural exchange apparatus described here to model preferences that vary.

17. The idea that cultural attributes might be profitably modeled as if they were alleles on a cultural chromosome—called "memes" by Dawkins [1976: 206]—has been studied systematically by Cavalli-Sforza and Feldman [1981] and applied to problems of gene-culture coevolution, such as the lactose absorption problem [Feldman and Cavalli-Sforza, 1989]. Related work includes Boyd and Richerson [1985].

18. In the Sugarscape software system, the string length is a user-specified parameter. We have experimented with lengths from 1 to 1000.

selected at random. If, at that position, neighbor B already matches Rose, no change is made. Otherwise, neighbor B's tag is flipped to agree with Rose's tag at that position. Rose's turn is then over, and it is the next agent's turn to flip its neighbor's tags. A summary statement follows.

Cultural transmission rule (tag-flipping):

- For each neighbor, a tag is randomly selected;
- If the neighbor agrees with the agent at that tag position, no change is made; if they disagree, the neighbor's tag is flipped to agree with the agent's tag.¹⁹

Now, imagine that we start with a primordial soup of agents with random genetics, random tag strings, and random initial positions on the sugarscape. In the course of an agent's life, its movement, based on the sugar drive, brings it into the neighborhoods of all sorts of other agents, who may flip its tags, just as we, in the course of our lives, may be influenced—in our tastes or beliefs—by contact with other individuals.

Cultural Groups

Having fixed on a tag transformation rule, a separate issue is how to define groups. As usual, we choose to do it in a simple fashion.

Group membership rule (tag majority): Agents are defined to be members of the Blue group when 0s outnumber 1s on their tag strings, and members of the Red group in the opposite case.²⁰

19. Many other cultural transmission rules are possible. An agent might flip n of its neighbors' tags, not just one, as above. Or, reversing roles, it could be neighbors who flip the agent's tags, k at a time. Or, agents and neighbors could *swap* tags, and so on. Eigen and Winkler [1981] have considered a variety of rules in the guise of "statistical bead games."

20. Many other rules for group membership are possible. One might identify particular positions—or sequences thereof—with certain groups. Tag position five might encode an agent's religion (0 for Muslim, 1 for Catholic). Group membership could require tag unanimity, with one tribe having all 0s and the other having all 1s. For tag strings of length 11, one three-group scheme is given below.

<u>Agent Group</u>	<u>Number of Zeros on String</u>
Blue	0 – 3
Green	4 – 7
Red	8 – 11

By increasing the string length and introducing considerations of tag ordering, very refined schemes become possible.

So, an agent with tag string 01010001010 would be a Blue, while one with 01001110101 a Red.²¹ Since tag order is irrelevant here we might call this a “voting rule.”

Notation

We have been denoting all rules with bold-faced letters (for example, **M** for movement). In principle, we could allot separate symbols for our cultural transmission (tag-flipping) and group membership (tag majority) rules. But, since we will only employ these rules together, we collapse them into a single symbol, **K**, which denotes this combination.

Cultural Dynamics

Recall that one component of the “proto-history” is the formation of spatially segregated, culturally distinct groups. Are the simple rules elaborated above *sufficient to generate* such outcomes? Returning to the familiar sugarscape, let us begin with a population of agents with random genetics, random tag strings each of length eleven, and random initial locations. The sugarscape grows back at unit rate. Agent movement is governed by rule **M** (each agent moves to the nearest unoccupied site having largest sugar within its vision and gathers the sugar) and sex is turned off. In the animations that follow agents are colored according to their group, with Blues colored blue and Reds colored red. A typical cultural evolution of this artificial society is shown in animation III-6.

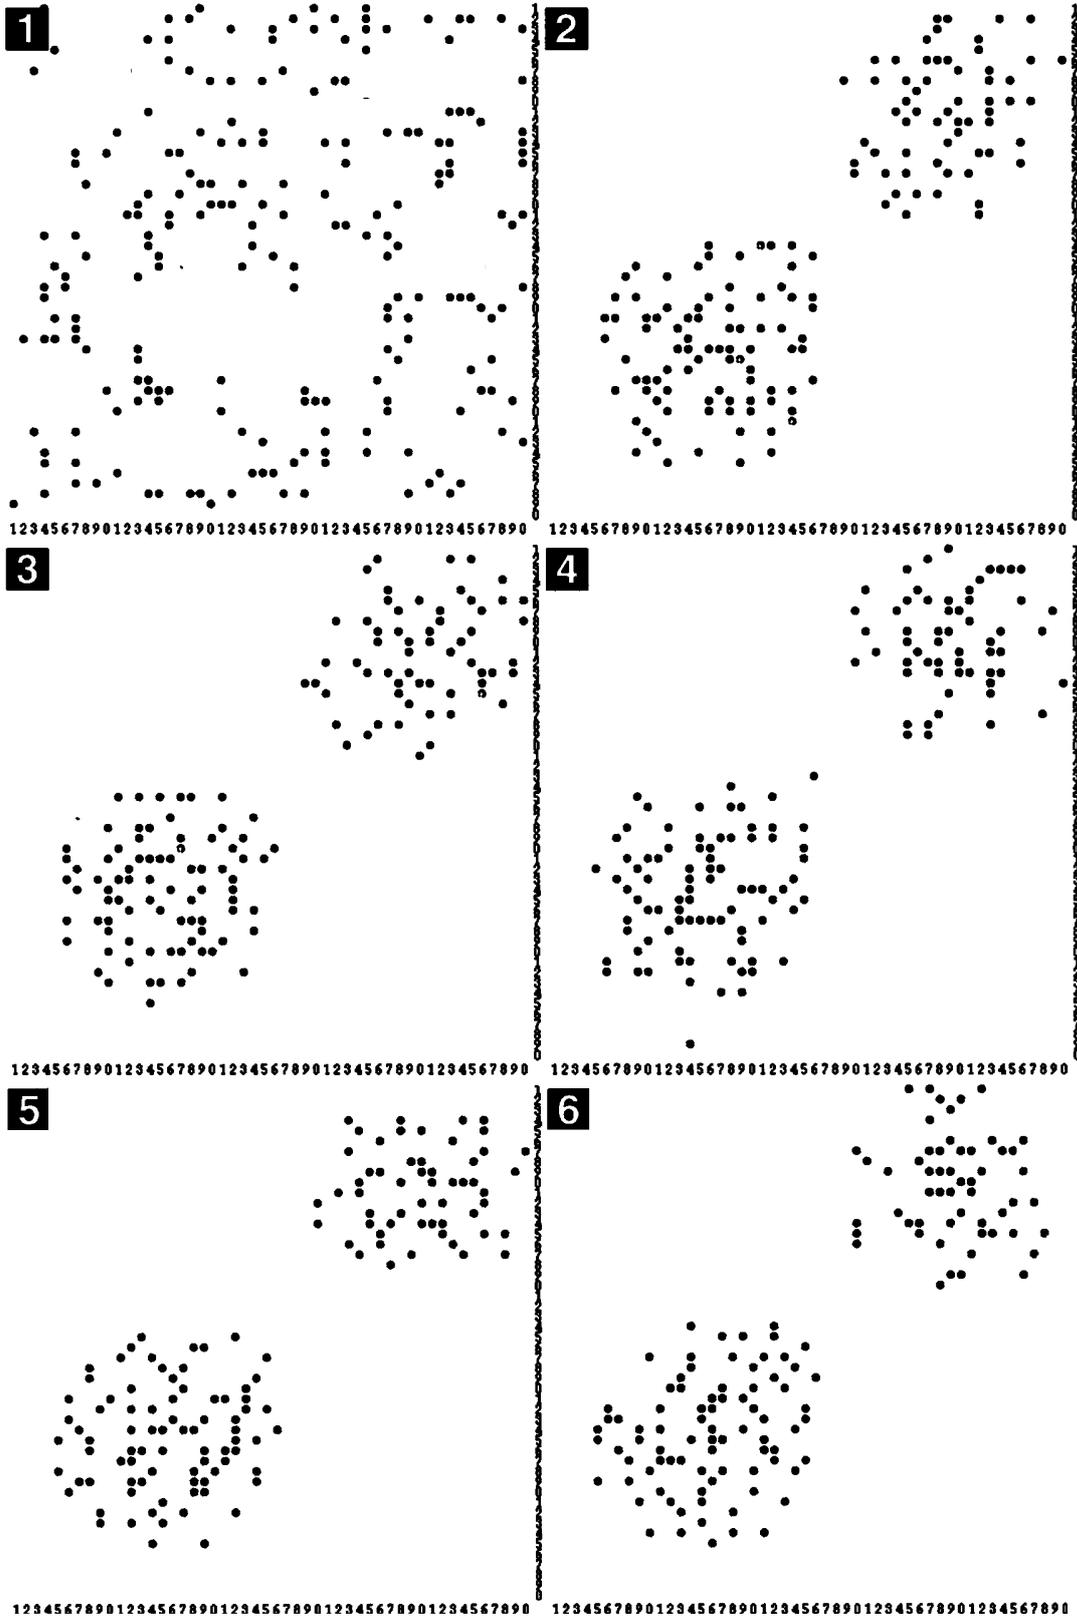
The animation terminates with all agents Blue after some 2700 time periods. If, as in this run, the initial population segregates spatially—with separate subpopulations hiving separate sugar heights—then each such subpopulation will ultimately converge to pure Blue or pure Red.²² Thus, **K** is *sufficient to generate* cultural groups.

One way to monitor tag-flipping dynamics is to use a histogram displaying—at each time—the percentage of all agents having 0s at each

21. In order to keep this rule unambiguous the number of tags should be odd.

22. For a spatially segregated population engaged in cultural transmission according to rule **K** it can be shown that a monochromatic state is an absorbing state of the process. If some (small) rate of cultural tag mutation is introduced, then the system will hover near one of the monochromatic states, occasionally changing colors completely. Similar dynamics arise in a variety of contexts; see Arthur [1988, 1990], Arthur, Ermoliev, and Kaniovski [1987], and Kaniovski [1994].

Animation III-6. Tag-Flipping Dynamics under Rules $(\{G_1\}, \{M, K\})$



Animation III-7. Tag Histogram Evolution under Rules $(\{G_1\}, \{M, K\})$

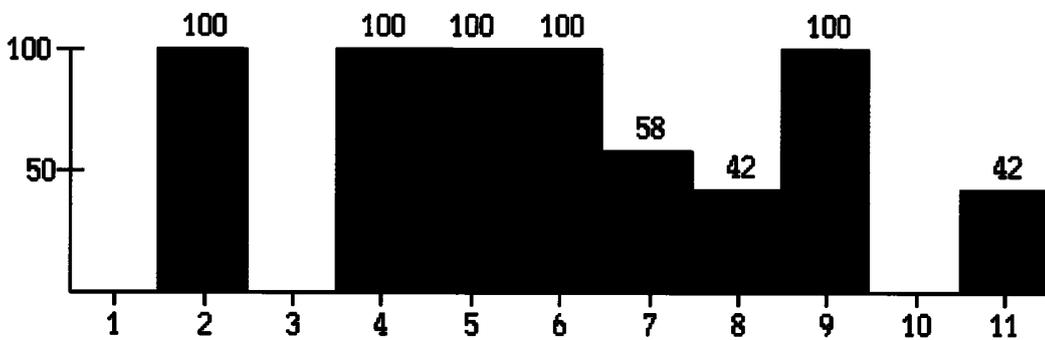
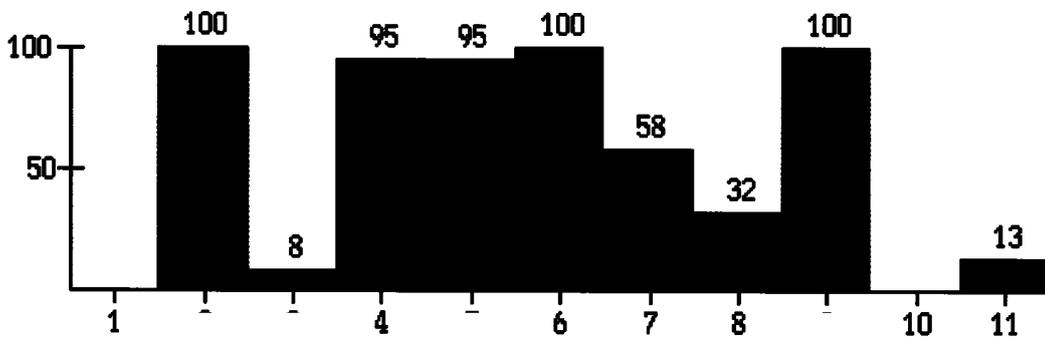
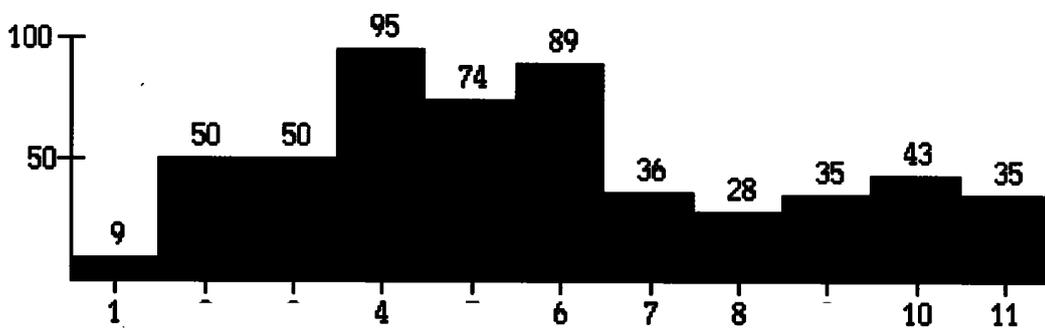
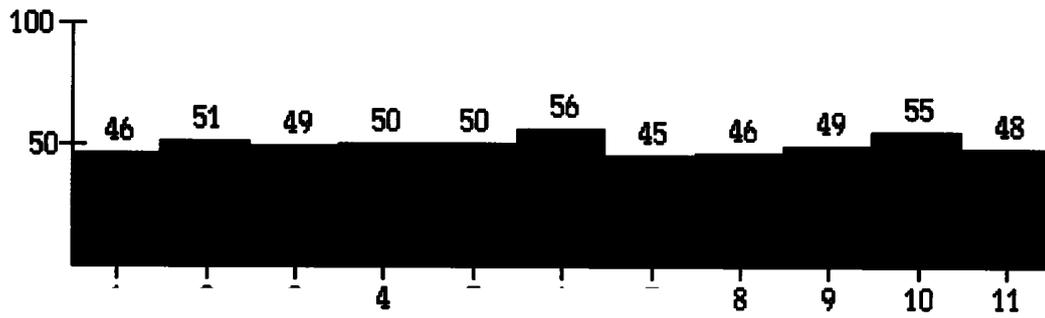
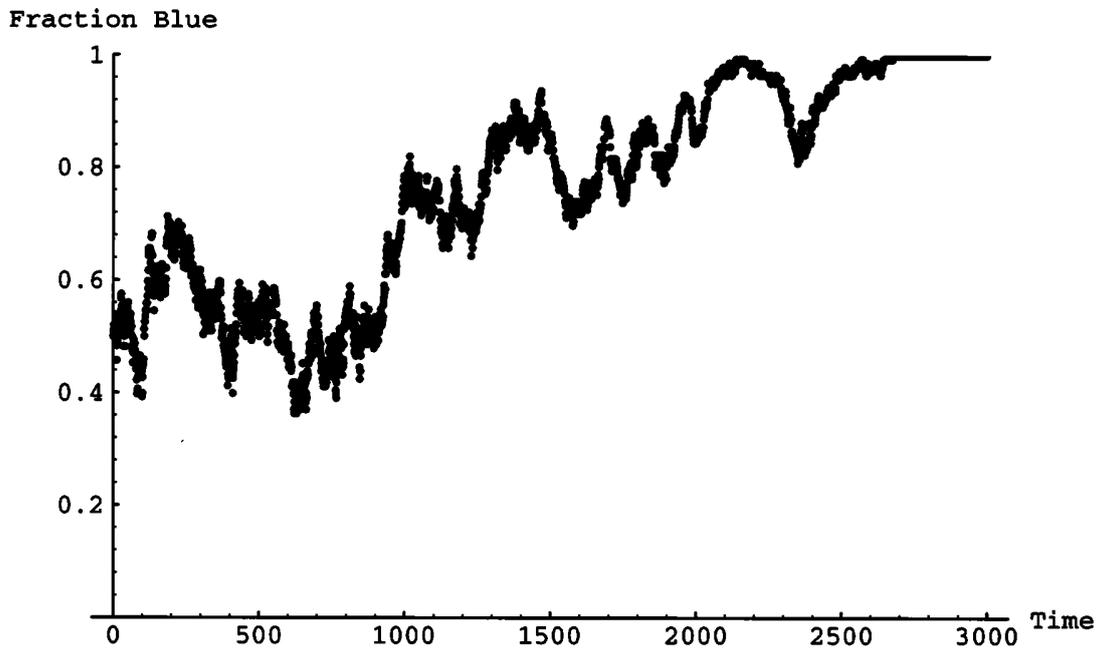


Figure III-8. Typical Cultural Tag Time Series Realization under Rules $(\{G_1\}, \{M, K\})$



position on the tag string. Suppose we freeze such a histogram at some instant in the tag-flipping process, as shown in the first frame of animation III-7. The horizontal axis is divided into eleven bins, one for each tag position. The height of the bin gives the percentage of agents having a 0 at that position. So, in this example 49 percent initially have a 0 as the third tag of their string, 55 percent have a 0 as the tenth tag, and so forth. Of course, tag-flipping unfolds in time so the histogram is not frozen but evolves as the reader will see by running animation III-7. This dynamic histogram gives tag statistics obtained from the previous animation.

For some tag positions the percentage of all agents having a zero ultimately converges to zero (all agents have a 1 there) or one hundred (all agents have a zero there). Of course, once either of these “unanimous” states is reached, there can be no further tag-flipping since there is no tag diversity at that position. In short, there is “lock-in.”

Figure III-8 shows a plot of the fraction of Blue agents over time. Note that long-run convergence to a single group need not be monotonic; wild fluctuations may occur en route to equilibrium. As the length of the cultural chromosome increases, so does the time required for convergence. Similarly, adding agents increases convergence time.

Now, two individuals might each consider themselves “American” culturally, while differing politically, religiously, or in other respects. An

interesting feature of this agent group membership rule is that agents can be very different culturally, measured position-by-position, and yet be members of the same group. To see this consider two agents having tag length five. Suppose the first has tag string 00011 and the second has tags 11000. These two agents have but a single tag in common, the third one, and yet they are both Blue since 0s predominate. All agents can be the same color (as in figure III-8) *without* being culturally identical. A corollary of this is that a pair of Blues can produce a Red agent. How can this happen? Imagine a Blue agent (who is the flipper) and a Blue neighbor (the “flippee,” as it were) with the following tag strings of length five:

Agent’s tag string	10100	Blue
Neighbor’s pre-flip tag string	01010	Blue

Each is Blue since 0s outnumber 1s. But suppose “God” (the random number generator) picks tag position three. Since the agent has a 1 there, it flips the neighbor’s tag to 1 at that position, resulting in the neighbor’s new tag string: 01110. But now 1s outnumber 0s, so the neighbor turns Red! Once more, a simple rule—here the tag-flipping rule—produces interesting results.

Ultimately, we want to have cultural transmission operational at the same time the sex rule is active, so we need some way to specify the state of a newborn child’s cultural tags. The transmission of cultural attributes from parents to children is termed *vertical*, as against the *horizontal* transmission we have been discussing.

Vertical Transmission of Culture

When sex rule **S** is active, a child’s tag string is formed by comparing the parents’ tags at each position and applying the following rule: If the parents have the same tag (both have 0 or both have 1), the child is assigned that tag. If, however, one parent has tag 0 and the other has tag 1, then we “toss a fair coin.” If it comes up heads, the child’s tag is 1, if tails, it is 0. All of this is summarized in table III-2.²³

23. Those with a background in population genetics will notice that this is strictly analogous to a random mating table for one locus with two alleles.

Table III-2. Probability That a Child Receives a 0 or 1 Tag When Born, Based on the Parents' Tags

Parents' tags	Probability that a child's tag is	
	0	1
Mother 0, father 0	1	0
Mother 0, father 1	1/2	1/2
Mother 1, father 0	1/2	1/2
Mother 1, father 1	0	1

This procedure is applied at each position, resulting in a cultural endowment—a tag string—for every newborn child. Of course, once the child is out on its own all agent behavioral rules apply, including **K**. Thus horizontal transmission will soon modify the child's initial, vertically transmitted, tags.

Networks of Friends

In Chapter II we defined agent neighbor networks and showed how these change over time. Earlier in this chapter we displayed genealogical networks. Here, given that the agents are flipping tags with their neighbors as they move around the sugarscape, a natural notion of "friendship" arises. Agents who at some point are neighbors and are close culturally are defined to be friends.²⁴ When an agent is born it has no friends. However, in moving around the landscape it meets many agents—as neighbors—and interacts with them culturally. Those agents with whom it interacts and who are closest to it culturally are ones it remembers as its friends.²⁵ Then, if one draws lines between friends, one has a friendship network.

To implement this in Sugarscape, we employ the Hamming distance to measure the closeness of cultural tag strings.²⁶ Each agent keeps track of

24. We offer this definition of "friendship" as a simple local rule that can be implemented efficiently, not as a faithful representation of current thinking about the basis for human friendship.

25. In the agent object this is implemented as a pointer to the friend agent (see Appendix A for more on the object-oriented implementation of the agents).

26. The Hamming distance between two (equal length) binary strings is obtained by comparing the strings position-by-position and totaling the number of positions at which they are different. Therefore, two strings having a Hamming distance equal to zero are identical.

the five agents it has encountered who are nearest it culturally; these are its friends. Each time an agent encounters a new neighbor the agent determines how close they are culturally and, if the neighbor is closer than any of the agent's five friends, the neighbor displaces one of them. Drawing connections between friends yields the network shown in animation III-8.²⁷

Many variants on this general idea are possible.²⁸ For instance, instead of connecting all friends one could draw lines only between *mutual friends*; that is, a line would connect agents A and B only if A considers B as a friend and *vice versa*. Another variation would be to connect only *best friends*; that is, A must consider B to be its best friend—closest culturally—and B must think the same of A. Finally, note that whether two agents are friends or not has no effect on their behavior. In this sense the network of friends is *external* to either the cultural exchange process or the friend assignment rule. A natural extension of the “network-of-agents” concept would be to permit regular agent-agent interaction over such networks, reinforcing positive interactions and perhaps breaking connections as a result of negative interactions, a kind of Hebbian picture.²⁹ In this way the networks take on a feedback flavor; interagent cultural transmission begets networks of friends, which in turn modify the transmission dynamics.

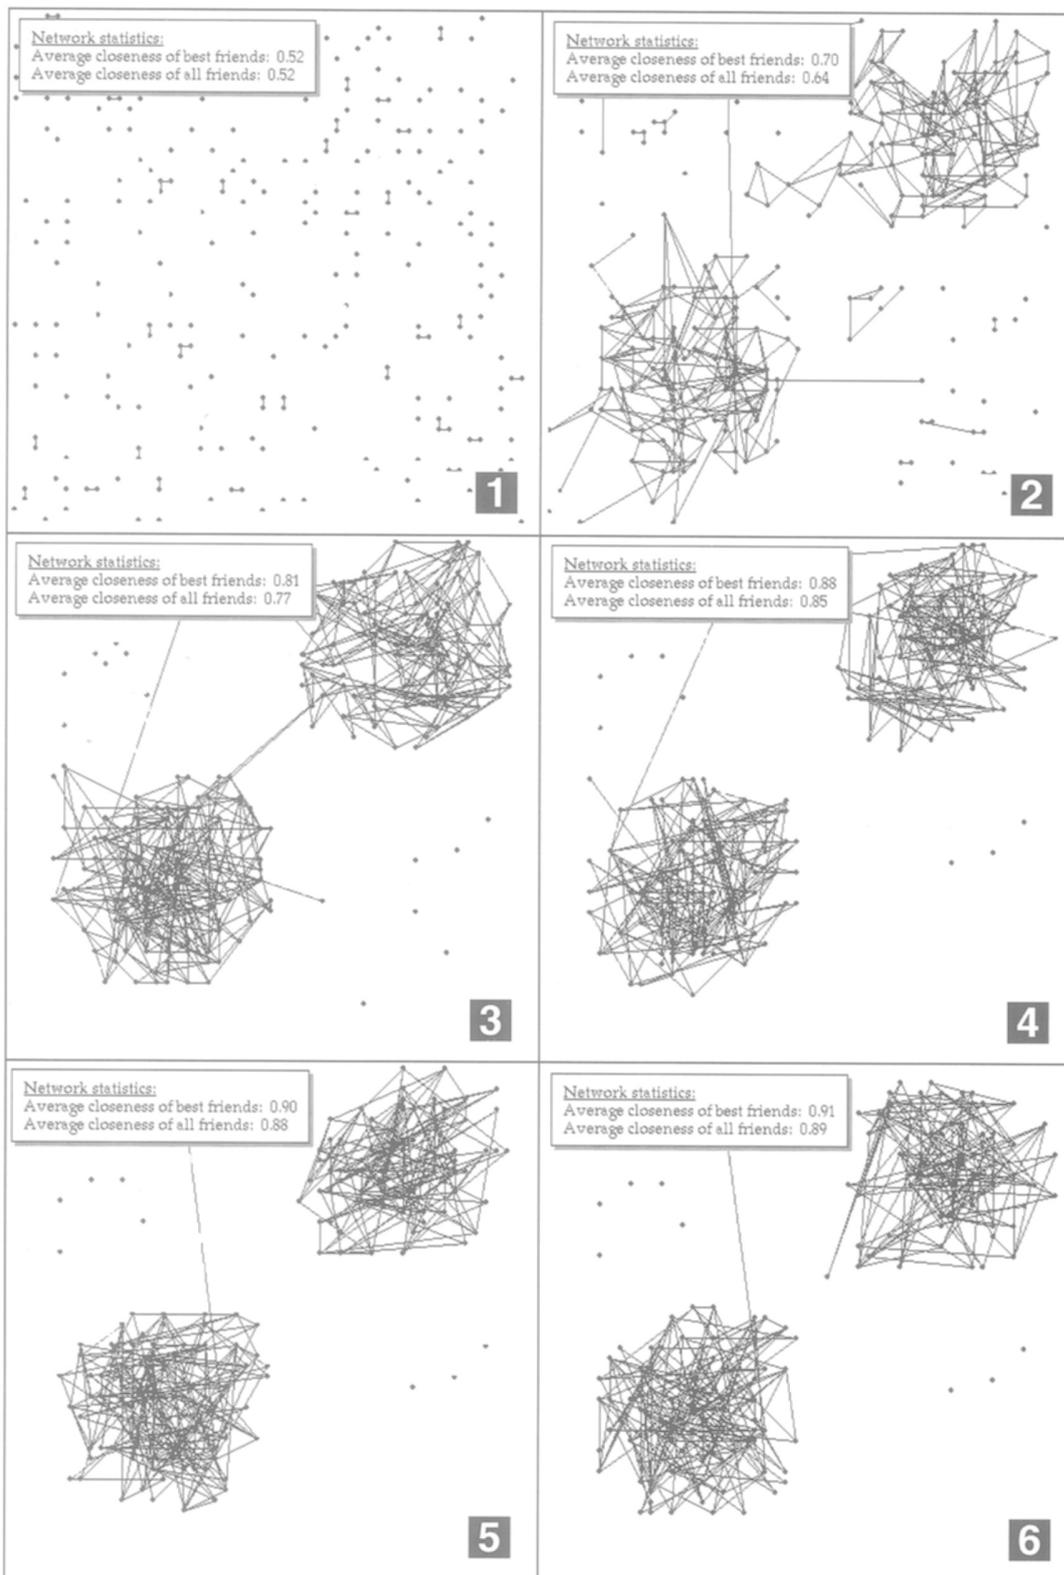
Networks such as those we have described manifest themselves in the real world in many important ways. Politically, restrictions on freedom of assembly, freedom of speech (press censorship), and freedom of movement (internal passport requirements) are standard tactics of repressive governments. The main aim of these measures is to keep individual dissenters—of which there may be a great many—isolated from one another, to keep them from *connecting* with other dissenters, and so

27. Occasionally a line across the entire lattice is observed. Since the sugarscape is a torus, an agent at the extreme left of the lattice may be a friend of an agent at the extreme right, yielding a friend connection line that spans the entire lattice.

28. There are actually two distinct ways to keep track of friends, producing somewhat different pictures. In animation III-8, once an agent stores another as its friend it never checks to see whether or not the agent continues to be close culturally once the two agents cease to be neighbors. That is, the agent's list of friends can become highly anachronistic as both it and its friends engage in cultural exchange over time. An alternative way to implement friends would be to keep updating the cultural closeness of friends each time period, although this would involve spatially nonlocal communication.

29. Recently Holland [1993] has studied the effects of tags on social interactions in an agent-based model.

Animation III-8. Evolution of a Network of Friends under Rules $(\{G_1\}, \{M, K\})$



to thwart the emergence of an organized *community* of dissenters, conscious of their numbers. How do changing political borders and “information revolutions” (for example, the Internet) affect the emergence of groups? How does “samizdata” spread across a landscape? Artificial societies allow us to study such questions systematically.

Now that rules for sexual reproduction and cultural exchange have been elaborated, let us turn to combat.

Combat

The cultural processes described above have proved *sufficient to generate* tribes—distinct cultural formations of agents. In this section we permit combat between agents from different tribes.³⁰ We do this by modifying the movement rule.

Specifically, imagine being a Blue agent. And suppose that within your range of vision there is a lattice position of sugar height 3, and that there is a Red agent sitting at that position. Then, if you take over that position, you take in the 3 sugar units *plus some additional reward from preying upon the Red agent*. One possibility is that you get the total accumulated sugar wealth of the agent. Or you might get a flat reward of, say, 2 sugar units. In the latter case, the full value of taking over the position would be $3 + 2 = 5$ sugar units. We will examine both types of reward rules. First, however, we need to establish reasonable conditions under which agents can prey on members of the opposite tribe.

To begin, it does not seem plausible that a “tiny” agent (one with little accumulated sugar) should be able to prey on a “huge” agent (one with vast accumulated sugar). At a minimum, then, we require that the predator be bigger than the prey in terms of accumulated sugar. It turns out that, to produce interesting dynamics, something more is required. In particular, if you are a Blue agent then you can plunder a Red agent—call him Rollo—only on two conditions. First, you must be bigger than Rollo. But second, there must be no other Red agent within your vision bigger than you will be after you defeat Rollo. In that case, we *define* the attack site as being *invulnerable to retaliation*. This second requirement provides

30. For an interesting discussion of tribal warfare from an anthropological perspective, see Ferguson [1992].

an element of deterrence, which enriches the dynamics significantly. A formal statement of the combat rule, for any reward α , is as follows:

Agent combat rule C_α

Look out as far as vision permits in the four principal lattice directions;

- Throw out all sites occupied by members of the agent's own tribe;
- Throw out all sites occupied by members of different tribes who are wealthier than the agent;
- The reward of each remaining site is given by the resource level at the site plus, if it is occupied, the minimum of α and the occupant's wealth;
- Throw out all sites that are vulnerable to retaliation;
- Select the nearest position having maximum reward and go there;
- Gather the resources at the site plus the minimum of α and the occupant's wealth, if the site was occupied;
- If the site was occupied, then the former occupant is considered "killed"—permanently removed from play.

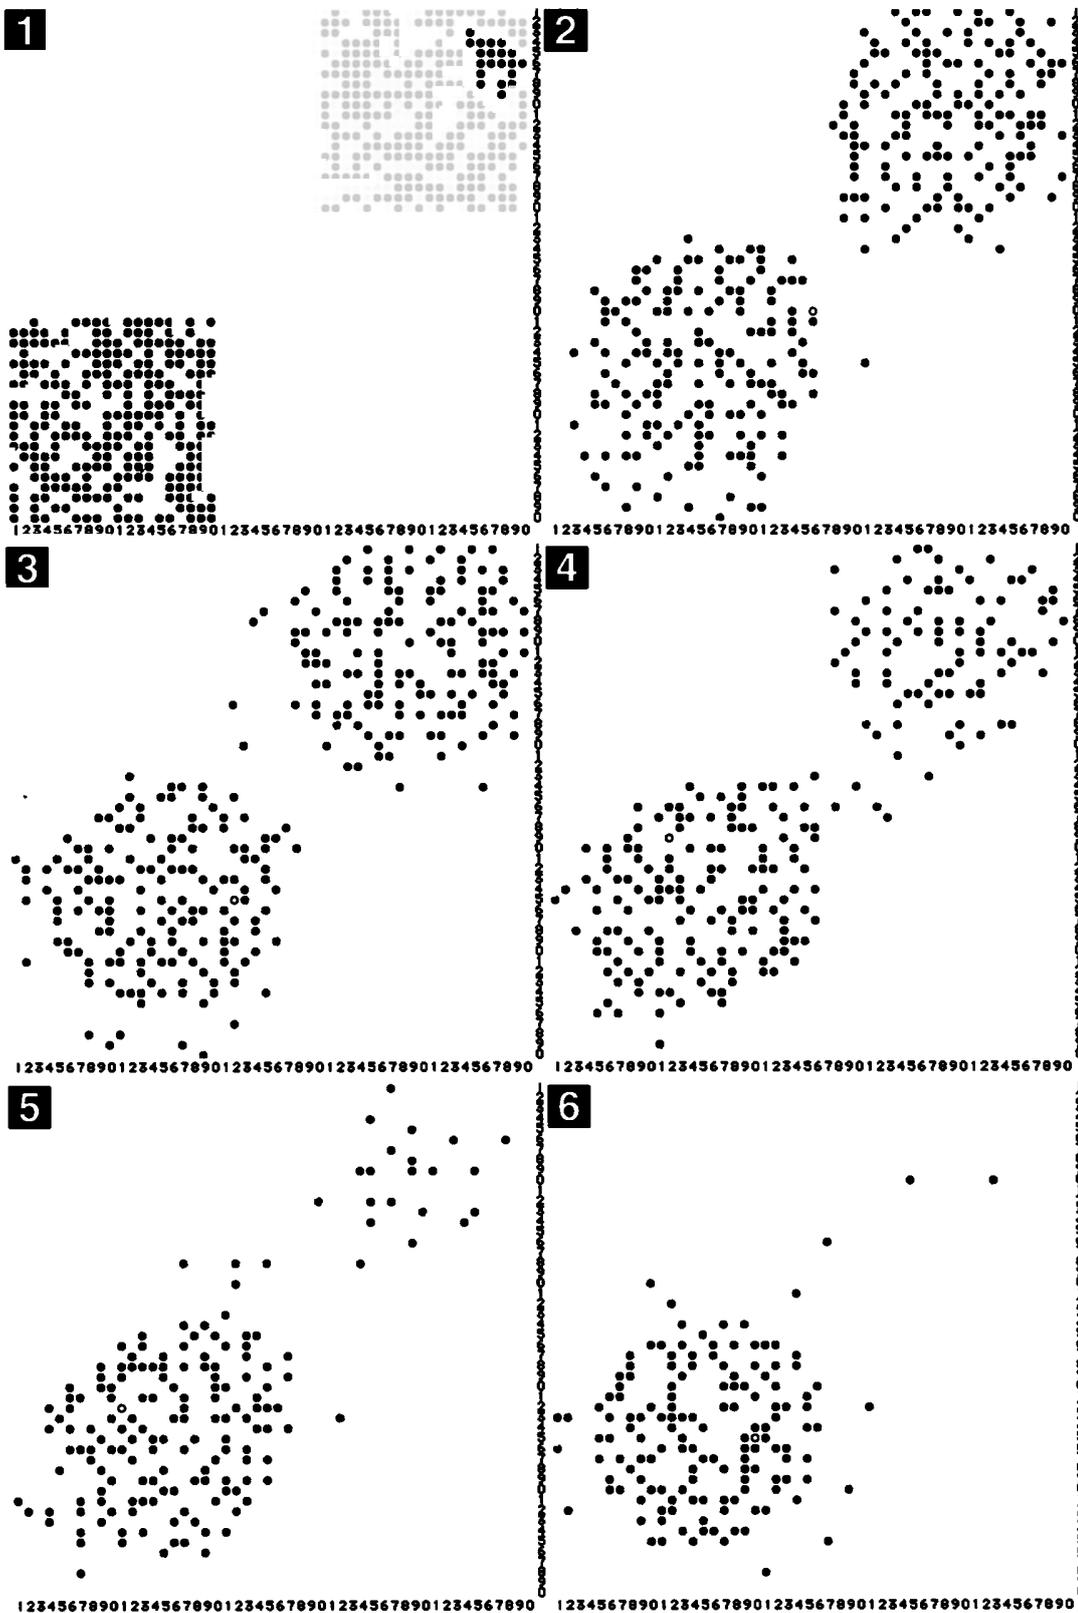
Note that the rule C_∞ implies that the aggressor receives the full accumulation of the defeated agent.

Reward Equal to Accumulated Wealth

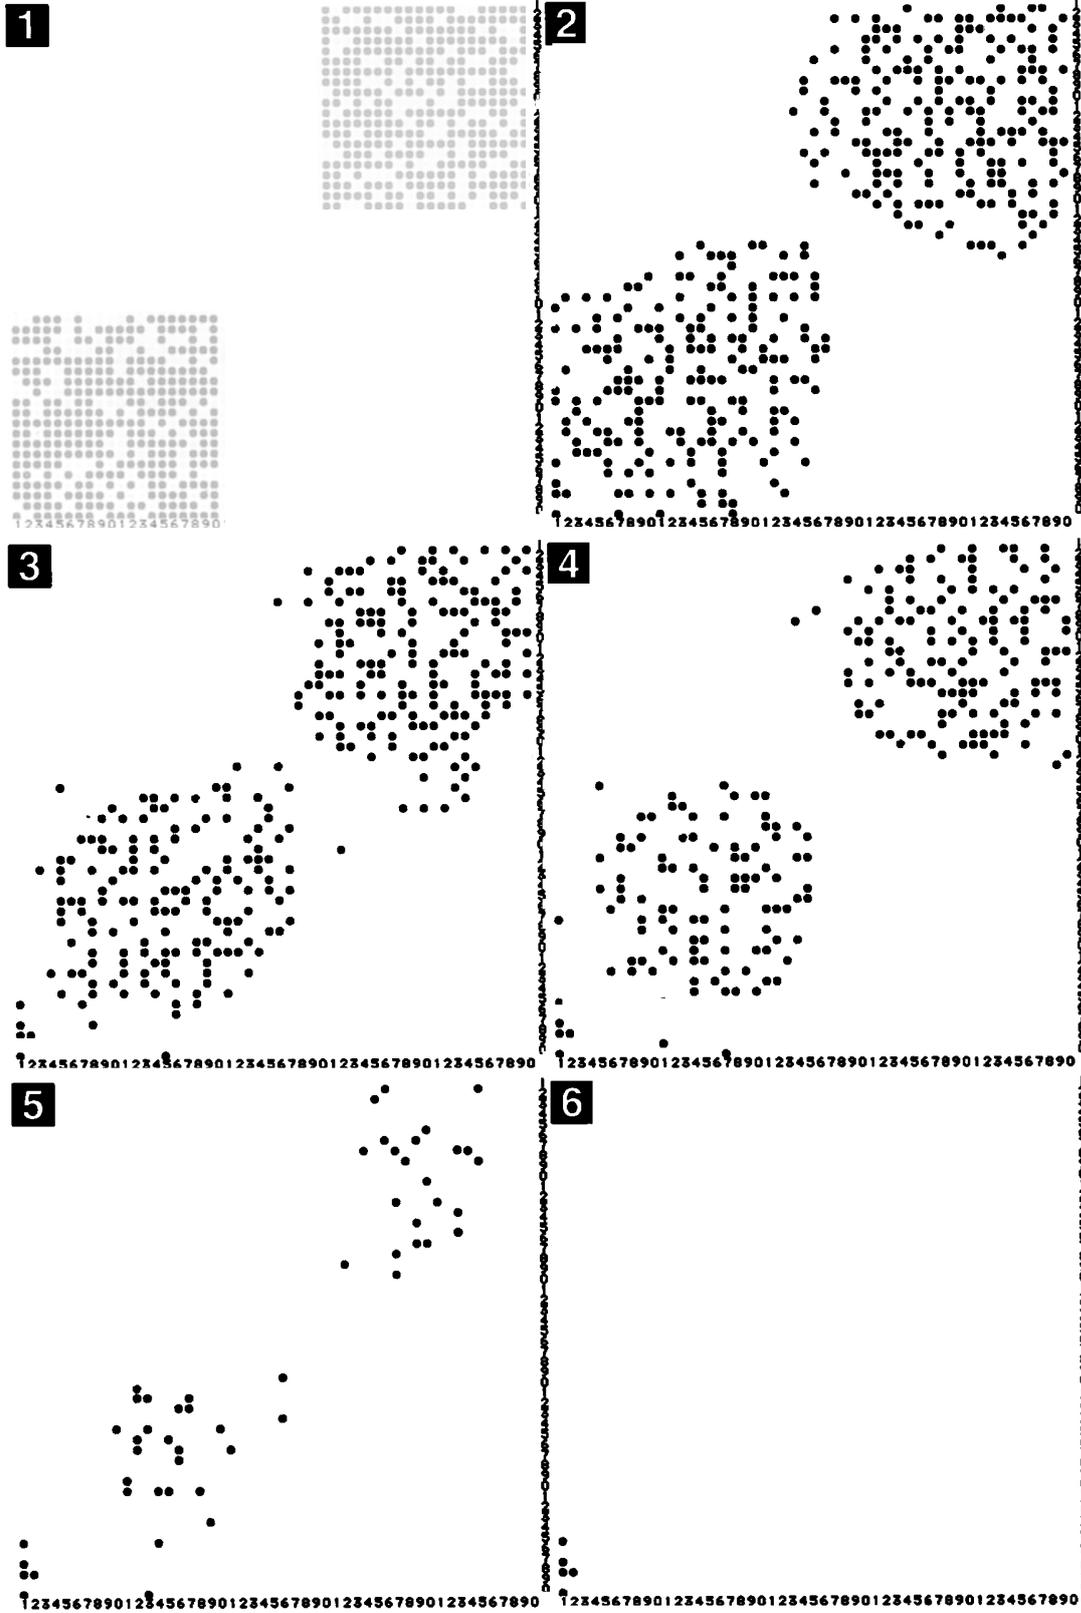
Indeed, in our first combat run we specify that the entire accumulated wealth of the prey goes to the predator; C_∞ is in effect. Allowing infinite lifetimes, predators can accumulate tremendous power under this rule. A type of increasing returns is evident—the bigger you are the faster you grow. Beginning with two separate tribes, social evolution under this reward rule always goes to one of three patterns: Blue eradicates Red, Red eradicates Blue, or small colonies (as small as one agent) of Reds and Blues coexist, each on its own mountain peak. Interestingly, Reds and Blues usually switch mountains in the course of the run. Animation III-9 shows an evolution to Blue dominance.

Starting from different initial conditions, animation III-10 shows an evolution to spatially segregated Blue and Red "colonies," each on its own mountain peak, with a small cluster of low vision Blues subsisting in the lowlands of the southwest.

Animation III-9. Blue Takeover under Rule System $(\{G_1\}, \{C_\infty\})$



Animation III-10. Red and Blue Coexistence under Rules $(\{G_1\}, \{C_\infty\})$



Notice that the mountains have, indeed, changed hands in these runs. In other words, no stable “battle front” emerges under this reward rule. If, however, the reward rule is modified, qualitatively different patterns of conflict become possible.

Reward Equal to a Fixed Value

Here, rather than the entire accumulated wealth of the victim, the conquering agent receives a *fixed sugar reward of 2 sugar units per kill*; that is, the combat rule is C_2 . To generate sustained combat between tribes, we fix the population at 400 agents by reintroducing the replacement rule, $R_{[60,100]}$, from Chapter II. Recall that this institutes a maximum lifetime uniformly distributed between 60 and 100 years, and when an agent dies it is replaced by a random agent of the same tribe (one with randomly chosen vision, metabolism, and initial sugar endowment). Animation III-11 shows the result.

Now we *do* get coherent battle fronts. Penetration is minimal and a prolonged “war of attrition” ensues, not the stunning *blitzkrieg* of animation III-9.

Effect of Rule Changes on Emergent Structures

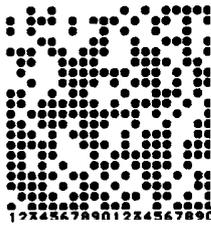
Clearly, *individual* behavior under the combat rule is different from *individual* behavior in the no-combat case. But, how does the prospect of combat affect emergent *collective* structures? Recall, for example, the *collective* waves of animation II-6. An initial block of Blue agents with maximum vision of ten (a relatively high value) propagated in a sequence of northeasterly waves. If we *turn combat off* and begin with opposed blocks of Blue and Red agents—in the southwest and northeast corners of the sugarscape—each population will propagate, again on a diagonal, toward the center in *collective waves that collide and interpenetrate* as shown in animation III-12.

Now, suppose we turn combat on, substituting rule C_α for rule M . Do waves still result? Animation III-13, with $\alpha = 2$, gives the answer.

Agents are deterred from racing forward to attack smaller agents of the opposed tribe by the presence of larger opposing agents within their vision. Precisely the factor—relatively high average vision—that produced the waves in the no-combat world now accounts for their absence. To complexity scientists, the moral is clear: When you change local rules, you may change emergent collective structures. For policymakers, there

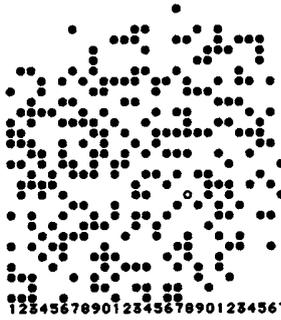
Animation III-11. "Trench War" under Rules $(\{G_1\}, \{C_2, R_{[60,100]}\})$

1



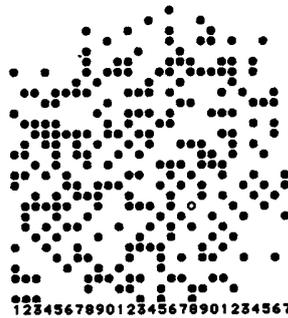
1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0

2



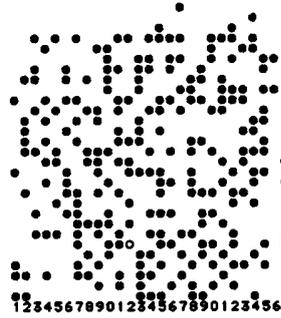
1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0

3



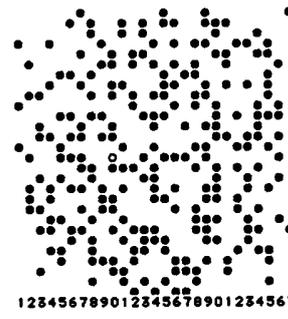
1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0

4



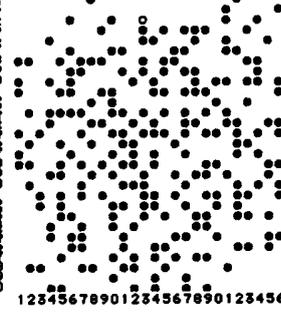
1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0

5



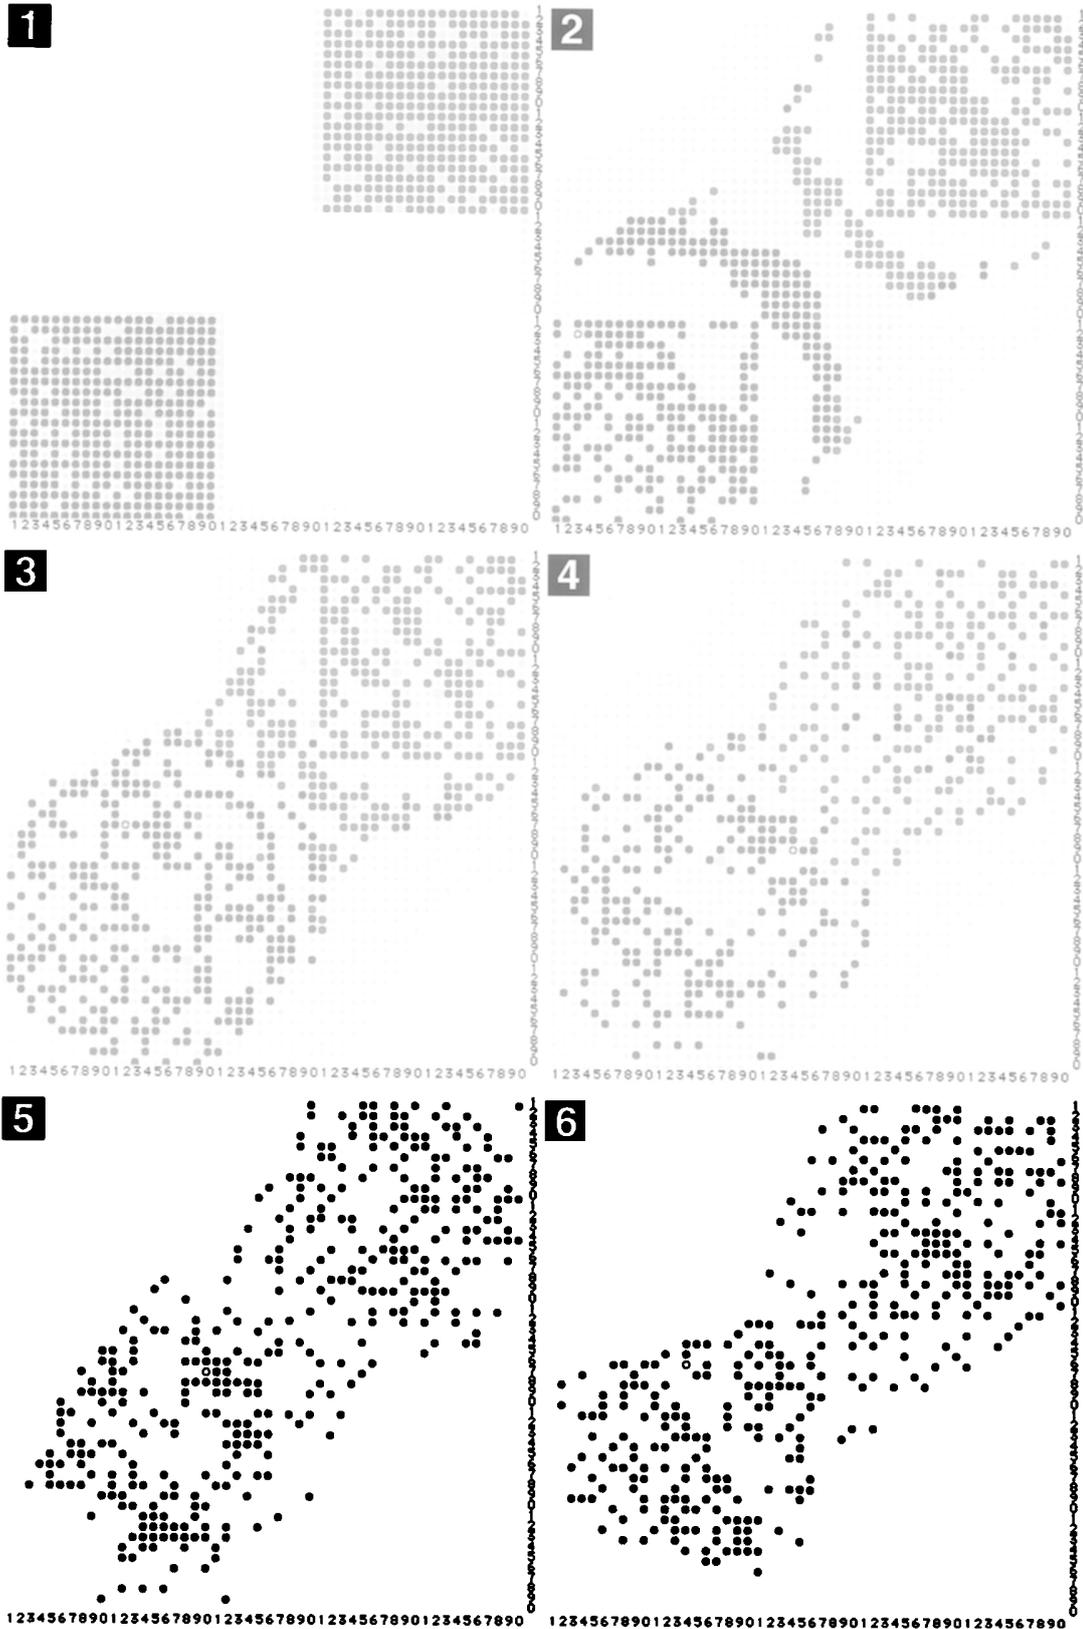
1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0

6



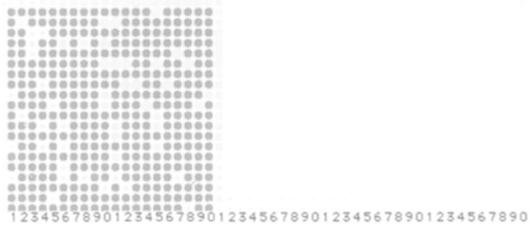
1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0

Animation III-12. Colliding Waves under Rules $\{G_1\}, \{M\}$



Animation III-13. Combat Eliminates Waves, Rule System $(\{G_1\}, \{C_2\})$

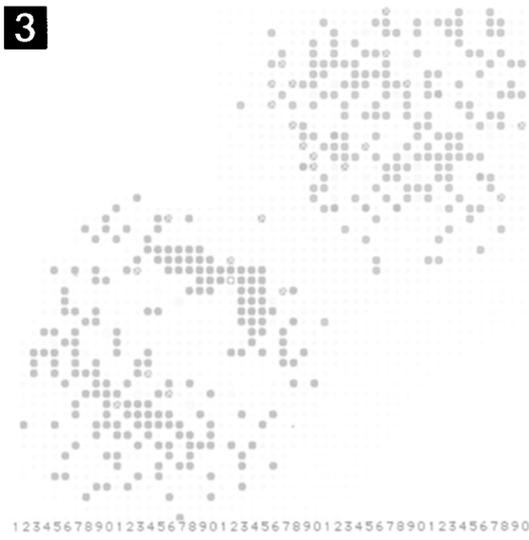
1



2



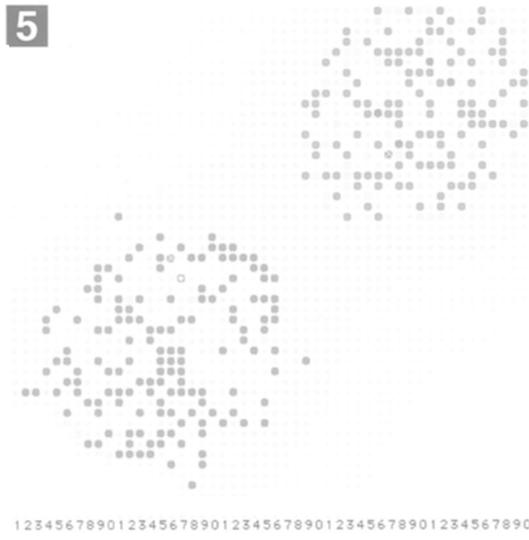
3



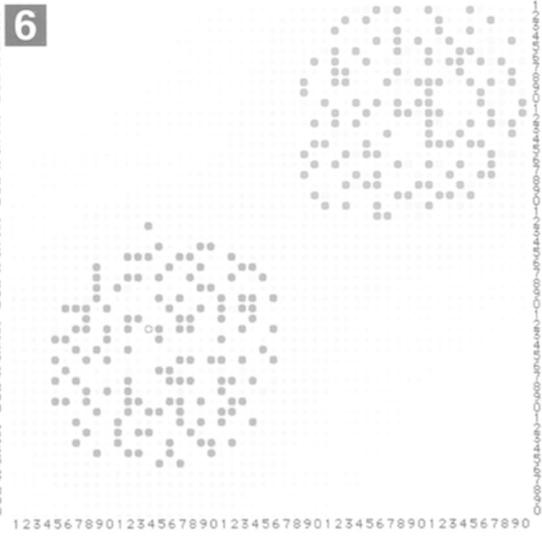
4



5



6



is a corollary: The most effective way to alter collective patterns of behavior may be from the bottom up, by modifying local rules.

Combat and Assimilation: Two Modes of Group Defense

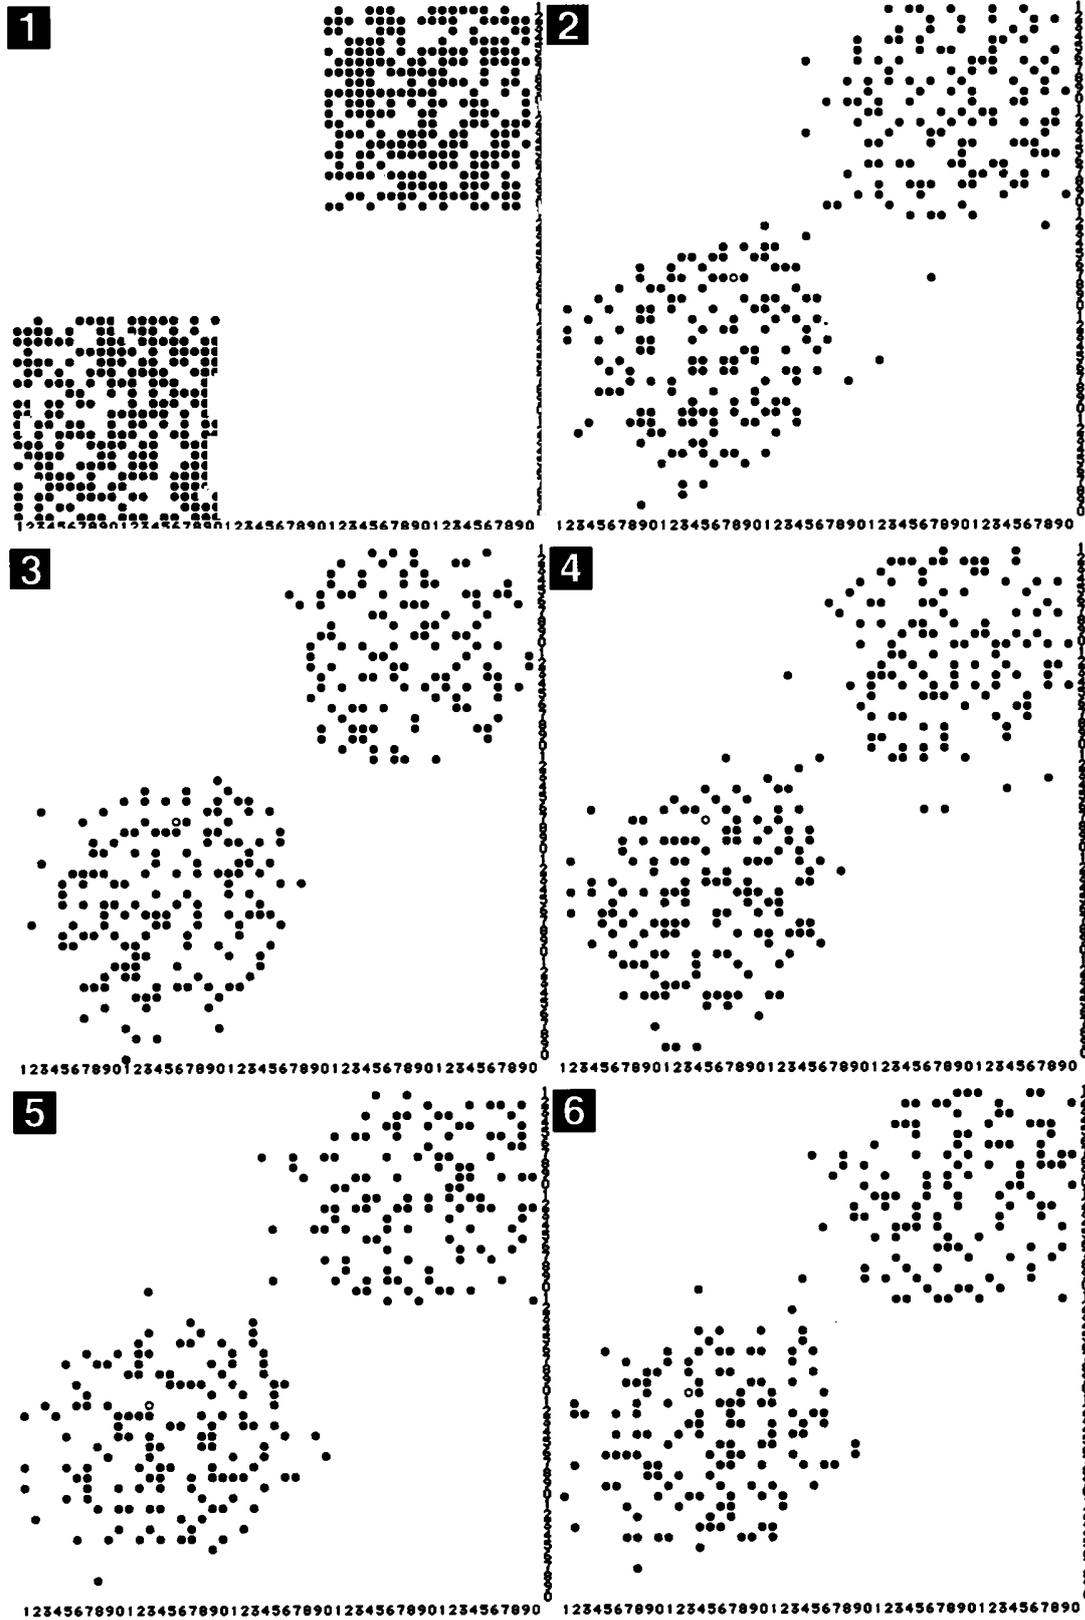
To examine how combat and cultural assimilation (tag-flipping) can interact, recall animation III-9. This was a pure combat simulation, beginning with segregated Red and Blue populations, one on each sugar mountain. The combat reward rule was that a victor acquires *the entire accumulated wealth* of the vanquished agent. In animation III-9 the result was Blue eradication of Red. In particular, we saw that once a sufficiently big Blue agent penetrated Red society, there was no way of stopping it, no way for any individual Red to defeat it in combat. But, what if that Blue invader's tags were being flipped while it was rampaging through Red society? Could the Reds convert it to a Red before it ravaged their society? Let us see. In animation III-14, everything is exactly as in animation III-9 (Blue takeover), except that cultural processes are unfolding.

Now, with cultural exchange processes active we do *not* see the same runaway to hegemony. Although a Blue agent certainly penetrates Red society and (through early combat victories) quickly acquires an insurmountable combat edge over any Red, the Blue invader's cultural tags are all the while being flipped—indeed, precisely because the Blue intruder is surrounded by Reds, the latter will get many opportunities to flip the invader's tags. The defending Reds are in fact able to convert, or assimilate, their attackers before being conquered. Later in this run, the Blue tribe defends itself the same way.

Once a huge invading Blue is converted to Red, it contributes a substantial measure of deterrence to the Red tribe. So, for deterrence, it is best if the big new convert is deployed far forward, close to the threatening Blue hordes. However, since the new convert might be “just barely Red”—by only one or two tags—it is not yet terribly trustworthy; it could easily be flipped back by a patch of neighboring Blues. Of course, before its conversion, the big former Blue might well have penetrated deep into the very center of Red society, where, as a new recruit, its complete ideological conversion—re-education—can proceed without “distraction,” as it were. The dynamics, ultimately, depend on the interplay between tag length and the combat reward.

Clearly, the longer the tag string, the longer it will take to convert an intruder. A total fanatic (one with tag string of “infinite” length) can never be converted. By the same token if the combat reward is very modest—say one sugar per kill instead of the entire wealth of the

Animation III-14. Combat and Cultural Transmission. Rule System $(\{G_1\}, \{C_\infty, K\})$



victim—there is much more time for tag-flipping before Blue takeover. In short, these two fundamental modes of group defense—combat and assimilation—trade off in interesting ways.³¹

While each rule of agent behavior discussed thus far—movement, combat, cultural transmission, and reproduction—deserves much more analysis, we have now assembled everything needed to “grow” the promised proto-history.

The Proto-History

The question posed at the start of this chapter was whether we can grow a crude caricature of early social history “from the bottom up” with these ingredients. Recall that the proto-history’s main components are the formation of spatially segregated tribes, tribe growth, and tribal interaction. We submit animation III-15 as a realization of the proto-history using the movement (**M**), sexual reproduction (**S**), and cultural (**K**) rules we have now accumulated.

At first there is a small, low density, primordial “soup” of agents with random genetics (vision, metabolism, and so on), random cultural tag strings, and random initial positions on the familiar twin-peaked sugar-scape. The fundamental drive for sugar produces migration to one or another of the sugar peaks, and thus spatial segregation into two sub-populations in which mating and cultural transmission occur. And each sub-population converges (culturally) to pure Red or pure Blue; the tribes are formed. Sexual reproduction now increases each tribe’s population, forcing Reds and Blues down from their sugar highland origins into the lowlands between. There the tribes interact perpetually, with collisions, penetrations, and conversions producing complex social histories. There are “expansionist” phases in which it looks as if one tribe will achieve hegemony; and there are “epochs” of stalemate, where scattered border contacts and “assimilations” are the rule.³²

31. Formal top-down models of group defense include Freedman and Wolkowicz [1986] and Freedman and Hongshun [1988].

32. This outcome—Red and Blue tribes on opposite mountains—is realized in somewhat less than 1/2 of the runs of the model with rules $\{(\mathbf{G}_1), \{\mathbf{M}, \mathbf{S}, \mathbf{K}\}\}$ active. In about 1/4 of the runs, Red cultural groups come to dominate both mountains, while Blue domination occurs with the same frequency. (As an aside, recall from animation III-6 that two non-communicating groups may both become Red without having *exactly* the same culture.) It is also the case that extinction on one or both of the mountains occurs with some small frequency.

Animation III-15. A Realization of the Proto-History, Produced by Rules $\{G_1\}$, $\{M, S, K\}$

