
Sugar and Spice: Trade Comes to the Sugarscape

Like the other animals, we find and pick up what we can use, and appropriate territories. But unlike the other animals, we also trade and produce for trade.

*Jane Jacobs, **Systems of Survival** [1992: xi]*

In the previous chapters we have studied simple agents having local rules for movement, sex, cultural exchange, and combat. In this chapter, we explore another crucial social behavior: *trade*. So far, our methodology has been to postulate one or more agent rules and then study the society that unfolds. Sometimes we presented a “target” social outcome before providing any rules (for example, the “proto-history”), while at other times we argued that the rules themselves were of interest since they were in some sense simple or minimal (for example, the movement rule **M**).

In this chapter, we proceed somewhat differently. We draw on neoclassical microeconomic theory for rules governing agent trade behavior.¹ These rules mediate the interaction of infinitely lived agents who have unchanging, well-behaved preferences that they truthfully reveal to one another and who engage in trade only if it makes them better off (technically, trade must be Pareto-improving). However, instead of the neoclassical stipulation that the agents interact only with the price system—that is, all agents are price-takers—we implement trade as occurring between neighboring agents at prices determined locally by a simple bargaining rule.² Individual agents do not use any nonlocal price infor-

1. In some agent-based computer simulations the term “trade” is used loosely, to denote any interagent transfer of internal stocks, independent of whether the agents have any internal mechanism for computing the welfare associated with such transfers. This is not a usage of interest to economists.

2. Kirman [1994], in his review of the literature on economies with interacting agents, suggests that “models in which agents interact with each other directly rather than indirectly through the market price mechanism provide a rich and promising class of alternatives which may help us to overcome some of the difficulties of the standard models.”

mation in their decisionmaking. Because price formation is local, this is a model of *completely decentralized exchange* between neoclassical agents. Later we relax some of the least realistic aspects of the neoclassical setup; for example, giving the agents finite lives and nonfixed preferences.

The main issue we address is the extent to which interacting agents are capable of producing *socially optimal* outcomes, that is, allocations of resources having the property that no agent can be made better off through further trade. The artificial societies modeling approach allows us to explore such questions systematically and reproducibly. In particular, we compare the performance of distinct classes of agents—neoclassical agents and various non-neoclassical ones. We find that neoclassical agents trading bilaterally are able to approach, over time, a price close to that associated with an optimal allocation. However, when the agents are made progressively less neoclassical—when they are permitted to sexually reproduce or have culturally varying preferences—the markets that emerge generally have suboptimal performance for indefinite periods of time.

Such results have important implications. First and foremost, the putative case for laissez-faire economic policies is that, left to their own devices, market processes yield equilibrium prices. Individual (decentralized) utility maximization at these prices then induces Pareto optimal allocations of goods and services. But if no price equilibrium occurs, then the efficiency of the allocations achieved becomes an open question and the theoretical case for pure market solutions is weakened.

We also investigate the effect of trade on variables studied in previous chapters. We find that the *carrying capacity* of the resource-scape is increased by trade, but so is the *skewness of the wealth distribution*. More agents exist in a society that engages in trade, but the resulting society is more unequal. Furthermore, the markets that result from our local trade rule generate *horizontal inequality*—agents with identical endowments and preferences end up in different welfare states. Importantly, the welfare theorems of neoclassical economics do not hold in such markets.

When agents are allowed to enter into credit relationships with one another—for purposes of bearing children—interesting financial networks emerge. Some agents end up as pure lenders, others as pure borrowers, and many turn out to be both lenders and borrowers. Indeed, entire financial hierarchies emerge within the agent society.

It seems natural to think of market processes as a form of social computation, with the agents operating as distributed processing “nodes” and the flow of commodities serving as inter-node communication. Each

Figure IV-1. Sugar Mountains in the Northeast and Southwest, Spice in the Northwest and Southeast



node (agent) executes a local optimization algorithm (purposive behavior), attempting to maximize a local objective (utility) function through decentralized interactions with other nodes (agents). The market as a whole—the social computer—tends toward a globally optimal allocation of goods, as if it were “attempting” to *compute* such an allocation. In this chapter we study how the success of this social computation depends on agent specifications.

Spice: A Second Commodity

To begin, since trade involves an exchange of distinct items between individuals, the first task is to add a second commodity to the landscape. This second resource, “spice,” is arranged in two mountains opposite the original sugar mountains, as depicted in figure IV-1.³ At each position there is a sugar level and capacity, as before, as well as a spice level and capacity.

Each agent now keeps two separate accumulations, one of sugar and one of spice, and has two distinct metabolisms, one for each good. These

3. An infinite variety of other arrangements of the resources is possible, of course, and we have experimented with various topographies. However, the configuration depicted in figure IV-1 will be used exclusively here. While the details of particular model runs are intimately intertwined with the economic geography employed, the qualitative character of the results does not depend on any particular topography.

metabolic rates are heterogeneous over the agent population, just as in the single commodity case, and represent the amount of the commodities the agents must consume each period to stay alive. Agents die if *either* their sugar or their spice accumulation falls to zero.

The Agent Welfare Function

We now need a way for the agents to compare their needs for the two goods. A “rational” agent having, say, equal sugar and spice metabolisms but with a large accumulation of sugar and small holdings of spice should pursue sites having relatively more spice than sugar. One way to capture this is to have the agents compute how “close” they are to starving to death due to a lack of either sugar or spice. They then attempt to gather relatively more of the good whose absence most jeopardizes their survival. In particular, imagine that an agent with metabolisms (m_1, m_2) and accumulations (w_1, w_2) computed the “amount of time until death given no further resource gathering” for each resource; these durations are just $\tau_1 \equiv w_1/m_1$ and $\tau_2 \equiv w_2/m_2$. The relative size of these two quantities, the dimensionless number τ_1/τ_2 , is a measure of the relative importance of finding sugar to finding spice. A number less than one means that sugar is relatively more important, while a number greater than one means that spice is needed more than sugar.

An agent welfare function giving just these relative valuations at the margins is⁴

$$W(w_1, w_2) = w_1^{m_1/m_\tau} w_2^{m_2/m_\tau}, \quad (1)$$

where $m_\tau = m_1 + m_2$. Note that this is a Cobb-Douglas functional form. The metabolisms make an agent’s welfare dependent upon its biology in just the way we want; that is, if an agent has a higher metabolism for the first commodity (sugar) than for the second (spice), then it views a site having equal quantities of sugar and spice as if there were relatively less sugar present.

This welfare function is *state-dependent* insofar as the arguments (w_1, w_2) denote accumulated quantities of the two commodities, not instantaneous consumption. This gives the agents the behavioral characteristic that as they age, for example, and accumulate wealth, they

4. This will be made precise below in the discussion of “internal valuations.”

view the same resource site differently.⁵ This state-dependence, while a departure from the utility function usual in neoclassical economics, is a natural way to represent preferences for agents who do not consume their entire commodity bundle each period.

The Agent Movement Rule in the Presence of Two Commodities

Given this welfare function, the movement rule followed by the agents is identical to what it was in the simple one commodity case, namely, look around for the best position and move there. The only difference in the two commodity case is that establishing which location is “best” involves evaluating the welfare function at each prospective site. Let s denote a site, with x_1^s and x_2^s the sugar and spice levels at that site. Formally, the agents perform an optimization calculation over the sites in their vision-parameterized neighborhood, N_v , according to

$$\max_{s \in N_v} W(w_1 + x_1^s, w_2 + x_2^s). \quad (2)$$

In other words, given an agent with some sugar wealth w_1 and spice wealth w_2 , every position within the agent’s vision is inspected and the agent calculates what its welfare would be were it to go there and collect the sugar and spice. Expression (2) says simply that the agent selects the site producing maximum welfare.⁶ As in the case of a single commodity, if there are several sites that produce equal welfare then the first site encountered is selected. Overall, the new movement rule for each agent is as follows.

Multicommodity agent movement rule **M**:⁷

- Look out as far as vision permits in each of the four lattice directions, north, south, east, and west;

5. Derivations in Appendix C give formal conditions under which an agent facing identical (distributions of) resource levels at distinct times in its life will rank sites differently due solely to changes in its wealth.

6. It is possible to unify the one and two commodity cases conceptually by imagining that in the former case the agents are “optimizing” a welfare function that has just one argument; that is,

$$W(w;m) = w^m.$$

7. We use **M** to symbolize all variants of the movement rule. In the Sugarscape software system the number of commodities, n , is a user-adjustable parameter, and so **M** has actually been implemented as the n -dimensional analog of expression (2).

- Considering only unoccupied lattice positions, find the nearest position producing maximum welfare;
- Move to the new position;
- Collect all the resources at that location.

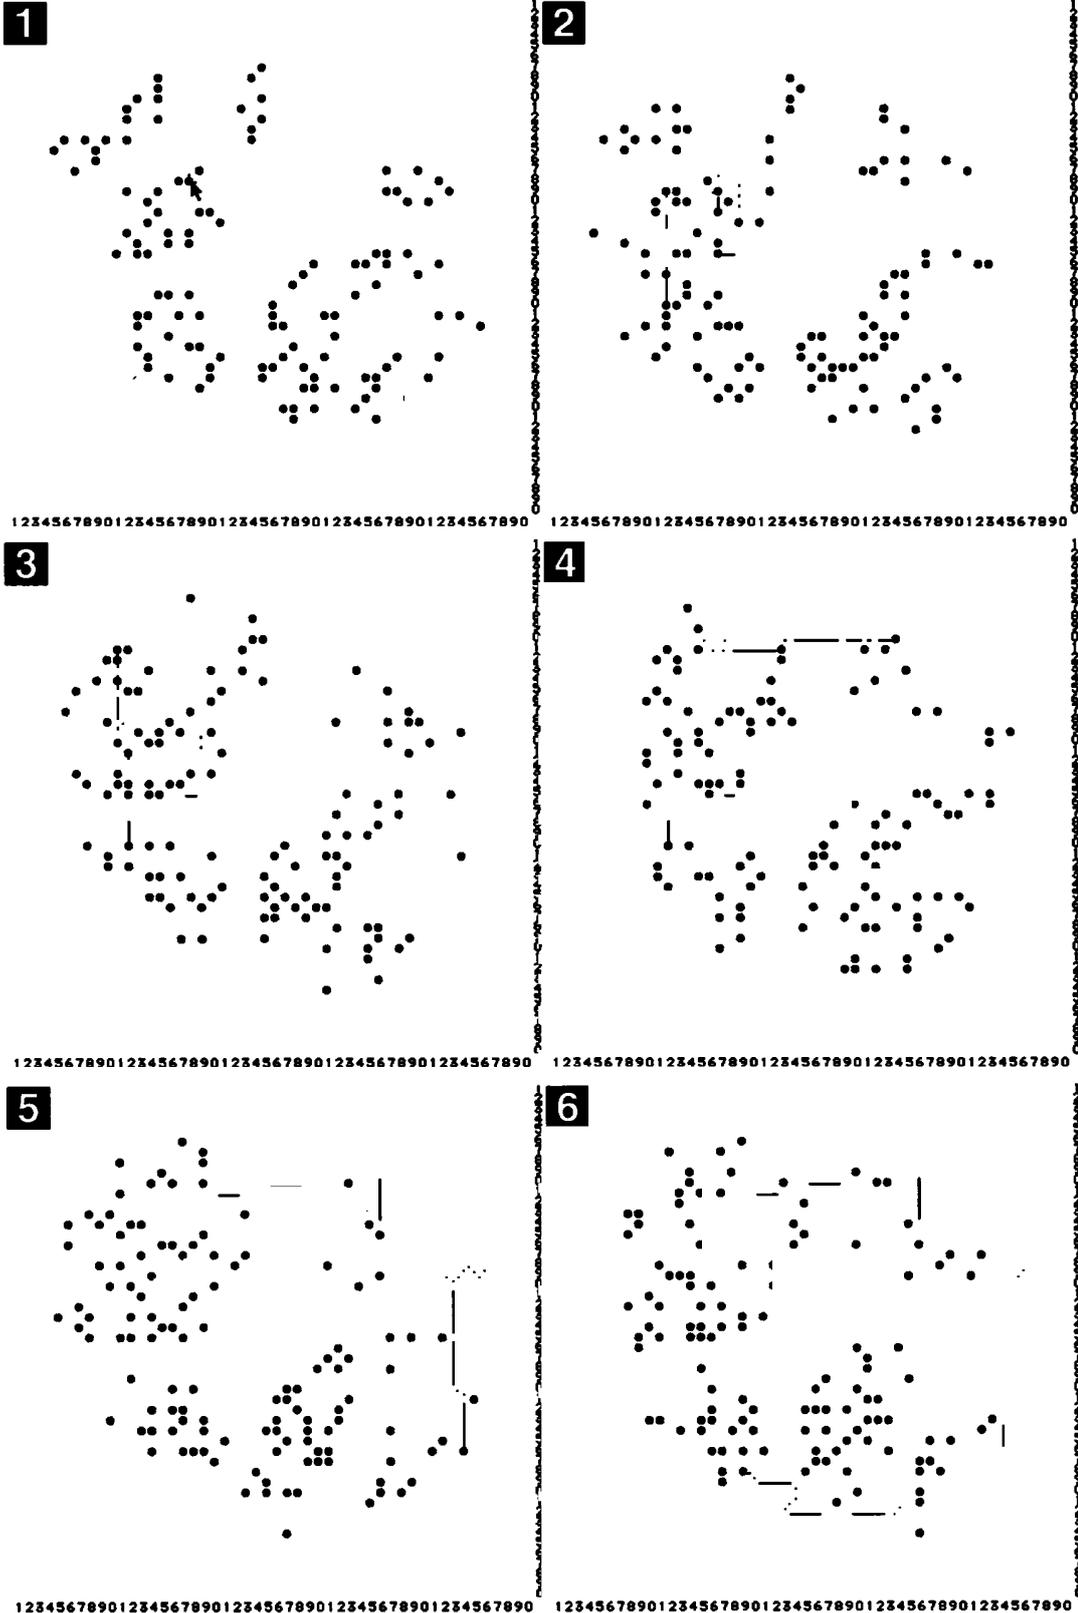
Now we study how the addition of the second commodity affects individual movement dynamics. To see that the effect is profound, one need only look at a particular run of the model, such as animation IV-1. Here vision is uniformly distributed in the agent population between 1 and 10, while metabolism for each of the two commodities is uniformly distributed between 1 and 5. A black tail is attached to one arbitrarily chosen agent, a so-called observational agent, to highlight the complexity of individual trajectories.

The agents search locally for the spot that makes them best off and they move there. However, because of the spatial separation of the two resources, agents move back and forth between the two types of mountains: Staying on one mountain for an extended period of time augments the agent's holdings of one commodity but dissipates its holdings of the other, forcing it to migrate. Note that if one were to average the observational agent's location over time its mean position would fall somewhere between the two types of mountains, despite the fact that the agent spends precious little time at such locations. That is, spatio-temporal averaging gives us little understanding of actual agent behavior.

The *carrying capacity* of this landscape is lower than in the single commodity (sugar-only) case, because there are now *two* ways to die, namely, by running out of *either* resource. A common route to death on the two-resource landscape is for an agent to run low on one of the resources while "stocking up" on the other and then find itself in a region of the resource-scape where there is a low density (and a flat gradient) of the needed good: having eaten its way deep into a high sugar (low spice) zone, the agent dies of spice deprivation, for example. Most agents in animation IV-1 never suffer this fate. When spice depletion threatens they have sufficient vision to find spice rich zones and replenish their stocks. Of course, there is another way agents might obtain commodities they need: through trade.⁸

8. Below we show that the general effect of trade is indeed to *augment* the carrying capacity.

Animation IV-1. Elaborate Trajectory of an “Observational” Agent in the Case of Both Sugar and Spice Present under Rule System $(\{G_1\}, \{M\})$



Trade Rules

Permitting agents to trade requires a rule system for the exchange of sugar and spice between agents.⁹ When will agents trade? How much will they trade? And at what price will exchange occur? There are a variety of ways in which to proceed.

The neoclassical theory of general equilibrium describes how a single centralized market run by a so-called auctioneer can arrive at an equilibrium price vector for the entire economy—a set of prices at which all markets clear. The image of an auctioneer announcing prices to the entire economy is quite unrealistic; no individual or institution could ever possess either complete knowledge of agent preferences and endowments or sufficient computational power to determine the appropriate prices. And even if market-clearing prices were somehow identified, why would all agents use them, why would all agents be price-takers?¹⁰

A more recognizable image is presented by Kreps [1990: 196], under the heading “Why (not) Believe in Walrasian Equilibrium?” He writes:

. . . we can imagine consumers wandering around a large market square, with all their possessions on their backs. They have chance meetings with each other, and when two consumers meet, they examine what each has to offer, to see if they can arrange a mutually agreeable trade. . . . If an exchange is made, the two swap goods and wander around in search of more advantageous trades made at chance meetings.

We implement trade in precisely this fashion, as welfare-improving (that is, mutually agreeable) bilateral barter between agents. No use is made of an auctioneer or any similar artifice. Agents move around the resource-scape following **M**, but are now permitted to trade with the agents they land next to, that is, their von Neumann neighbors. When an agent-neighbor pair interacts to trade, the process begins by having each agent compute its internal valuations of sugar and spice. Then a

9. Since there is no money in our artificial society, it is perhaps more accurate to describe interagent trade as barter. Throughout this chapter we shall use the terms “trade,” “exchange,” and “barter” interchangeably. On the emergence of money in an agent-based model see Marimon, McGrattan, and Sargent [1990].

10. That is, if certain groups of agents can engage in welfare-improving trade *between themselves* at prices other than the market-clearing ones, why would they not do so? Such advantageous reallocations of endowments have been studied by Guesnerie and Laffont [1978].

bargaining process is conducted and a price is agreed to. Finally an exchange of goods between agents occurs if both agents are made better off by the exchange. This process is repeated until no further gains from trade are possible.¹¹ We now present the details of this process.

Internal Valuations

According to microeconomic theory, an agent's internal valuations of economic commodities are given by its so-called marginal rate of substitution (*MRS*) of one commodity for another. An agent's *MRS* of spice for sugar is the *amount* of spice the agent considers to be as valuable as one unit of sugar, that is, the *value* of sugar in units of spice.¹² For the welfare function (1) above, the *MRS* can be shown to be

$$MRS \equiv \frac{dw_2}{dw_1} = \frac{\frac{\partial W(w_1, w_2)}{\partial w_1}}{\frac{\partial W(w_1, w_2)}{\partial w_2}} = \frac{\frac{m_1 w_1^{(m_1 - m_T)/m_T} w_2^{m_2/m_T}}{m_T}}{\frac{m_2 w_1^{m_1/m_T} w_2^{(m_2 - m_T)/m_T}}{m_T}} = \frac{m_1 w_2}{m_2 w_1} = \frac{\frac{w_2}{m_2}}{\frac{w_1}{m_1}} = \frac{\tau_2}{\tau_1}. \quad (3)$$

Note from (3) that an agent's *MRS* depends in an essential way on its metabolisms, that is, its biology. Earlier we noted that the quantities τ_1 and τ_2 represented the times to death by sugar and spice starvation, respectively, assuming no further resource gathering. These quantities are also measures of the *relative internal scarcity* of the two resources, in

11. Because our agents trade at nonequilibrium prices this is a non-Walrasian model. In particular, it is a kind of Edgeworth barter process (Negishi [1961], Uzawa [1962], Hahn [1962], and Mukherji [1974]; see Arrow and Hahn [1971: Chapter 13], Hahn [1982], and Fisher [1983: 29-31] for reviews). However, the bilateral nature of our model makes it more completely decentralized than the usual Edgeworth process since prices will generally be heterogeneous during each round of trading. The model closest to ours is Albin and Foley [1990], in which agents maintain fixed positions on a circle and trade with their neighbors. Other non-tâtonnement models include Aubin [1981], Benninga [1992], Feldman [1973], Hey [1974], Lengwiler [1994], Smale [1976], Stacchetti [1985], Walker [1984], and the simulation study of Takayasu *et al.* [1992]. Models of decentralized exchange in which the role of money is studied include Eckalbar [1984, 1986], Friedman [1979], Kiyotaki and Wright [1989, 1991], Madden [1976], Marimon, McGrattan, and Sargent [1990], Menger [1892], Norman [1987], Ostroy and Starr [1974, 1990], and Starr [1976]. Stochastic models of exchange include Bhattacharya and Majumdar [1973], Föllmer [1974], Garman [1976], Keisler [1986, 1992, 1995, 1996], Hurwicz, Radner, and Reiter [1975a, 1975b], and Mendelson [1985]. There is a growing literature of models in which economic agents interact directly with neighbors; for example, see An and Kiefer [1992], Anderlini and Ianni [1993a, 1993b], Ellison [1992], Kiefer, Ye, and An [1993], and Herz [1993].

12. Technically, the *MRS* is the local slope of the sugar-spice indifference curve.

the sense that an agent whose $MRS < 1$, for example, thinks of itself as being relatively poor in spice.

When two agents, A and B, encounter one another—that is, when one moves into the other’s neighborhood—the MRS of each agent is computed. Here we treat these internal valuations as common knowledge; that is, the agents truthfully reveal their preferences to one another. If $MRS_A > MRS_B$ then agent A considers sugar to be relatively more valuable than does agent B, and so A is a sugar buyer and a spice seller while agent B is the opposite.¹³ The general conditions are summarized in table IV-1. As long as the MRS s are not the same there is potential for trade; that is, one or both of the agents may be made better off through exchange.

Table IV-1. Relative MRS s and the Directions of Resource Exchange

Action	$MRS_A > MRS_B$		$MRS_A < MRS_B$	
	A	B	A	B
Buys	sugar	spice	spice	sugar
Sells	spice	sugar	sugar	spice

The Bargaining Rule and Local Price Formation

Having established the *direction* in which resources will be exchanged, it remains to specify a rule for establishing the *quantities* to be exchanged. The ratio of the spice to sugar quantities exchanged is simply the *price*. This price must, of necessity, fall in the range $[MRS_A, MRS_B]$. To see this, consider the case of two agents, A and B, for whom $MRS_A > MRS_B$. Since A will acquire sugar from B in exchange for spice (see table IV-1), its MRS will, according to (3), decrease as a result of the exchange, while B’s MRS will increase. But A will not give up spice for sugar at just any price. Rather, the most spice it is willing to give up for a unit of sugar is precisely its MRS ; for one unit of sugar it is willing to trade any amount of spice *below* the quantity given by the MRS . Analogously for B: it is willing to trade at any price *above* its MRS . Thus the range of feasible prices is $[MRS_A, MRS_B]$.

A rule for specifying exchange quantities, and therefore price, might

13. Note that whether a particular agent is a sugar buyer or seller is completely endogenous—it depends on the MRS of the other agent with whom the exchange interaction occurs.

be called a *bargaining rule* since it can be interpreted as the (adaptive) way in which two goal-seeking agents instantiate a price from the range of feasible prices.¹⁴ While all prices within the feasible range are “agreeable” to the agents, not all prices appear to be equally “fair.” Prices near either end of the range would seem to be a better deal for one of the agents, particularly when the price range is *very* large. Following Albin and Foley [1990], we use as the exchange price the *geometric mean* of the endpoints of the feasible price range. That is, the trading price, p , is determined according to

$$p(MRS_A, MRS_B) = \sqrt{MRS_A MRS_B}. \quad (4)$$

The primary result of this rule is to moderate the effect of two agents having vastly different *MRSs*.¹⁵ It turns out that it is more natural to work with $\pi = \ln(p)$, and in describing our artificial economy below we shall compute statistics for π .¹⁶

Finally, with the price determined, we need to specify the actual quantities of sugar and spice to be exchanged. Here we add the element of indivisibility by stipulating that each exchange involve unit quantity of one of the commodities. In particular, for $p > 1$, p units of spice are exchanged for 1 unit of sugar. If $p < 1$, then 1 unit of spice is exchanged for $1/p$ units of sugar.

The Trade Algorithm

Given that two agents have “bargained to” a price, and thereby specified the quantities to be exchanged, the trade only goes forward if it makes both agents better off. That is, trade must improve the welfare of both agents. Furthermore, since discrete quantities are being traded, and

14. There exists an enormous literature on bilateral bargaining when agents have incomplete information. A good introduction is Osborne and Rubinstein [1990], while important papers are reprinted in Linhart, Radner, and Satterthwaite [1992]; see also Gale [1986a, 1986b] and Binmore and Dasgupta [1987]. Since our agents truthfully reveal their preferences we do not make use of these ideas here. Clearly this is an important topic for future work.

15. We have also experimented with a bargaining rule that simply picks a random number from the interval $[MRS_A, MRS_B]$. The qualitative character of the results reported below is insensitive to this change.

16. To see this, note that trading 10 units of spice for one sugar ($p = 10$) should be treated as equally distant from $p = 1$ as trading 10 sugars for one spice ($p = 1/10$). With $p = \ln(p)$, this requirement is met since $\ln(10) - \ln(1) = \ln(1) - \ln(1/10)$.

therefore repeated exchange may never lead to identical agent *MRS*s, special care must be taken to avoid infinite loops in which a pair of agents alternates between being buyers and sellers of the same resource upon successive application of the trade rule. This is accomplished by forbidding the *MRS*s to cross over one another.¹⁷ Putting all this together we have:¹⁸

Agent trade rule T:

- Agent and neighbor compute their *MRS*s; if these are equal then end, else continue;
- The direction of exchange is as follows: spice flows from the agent with the higher *MRS* to the agent with the lower *MRS* while sugar goes in the opposite direction;
- The geometric mean of the two *MRS*s is calculated—this will serve as the price, p ;
- The quantities to be exchanged are as follows: if $p > 1$ then p units of spice for 1 unit of sugar; if $p < 1$ then $1/p$ units of sugar for 1 unit of spice;
- If this trade will (a) make both agents better off (increases the welfare of both agents), and (b) not cause the agents' *MRS*s to cross over one another, then the trade is made and return to start, else end.

Note that the bargaining rule constitutes step 3 of the algorithm.¹⁹

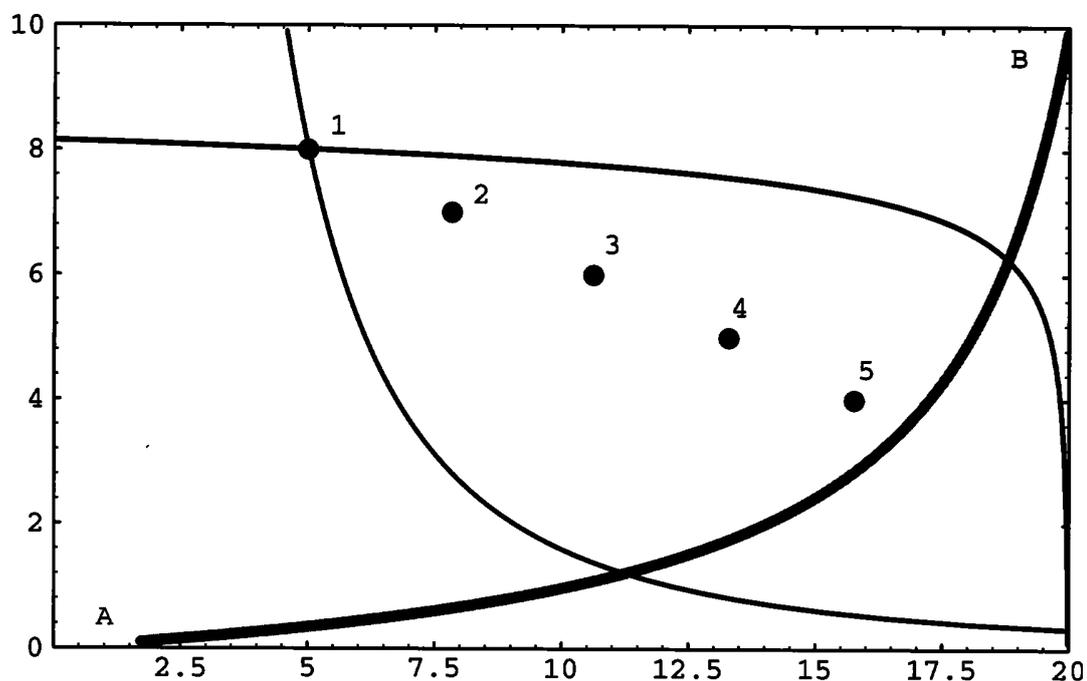
For a graphical interpretation of **T**, consider the so-called Edgeworth box shown in figure IV-2. Here agent A has sugar-spice endowment of (5, 8), while agent B possesses (15, 2). The red line intersects A's endowment and is A's line of constant utility; that is, A is indifferent between its endowment and all other sugar-spice allocations on the red line. Any allocation below this line is unacceptable to A since such an allocation would yield less welfare than A enjoys at its present position. All allocations above the isoutility line are preferred by A to its current allocation.

17. That is, given $MRS_A >(<) MRS_B$, agents stop trading if one additional trade will make $MRS_A <(>) MRS_B$.

18. Heretofore, we have described all rules for the agents and the sugarscape as "simple local rules." We would like it if the trade algorithm **T** could also be described in this way, but realize that one can reasonably say that this rule, although completely local, is hardly simple (requiring, for example, partial differentiation, computation of square roots, and so on). Perhaps it is better described as being the simplest local rule in the neoclassical spirit.

19. It is possible to substitute other bargaining rules simply by replacing this step.

Figure IV-2. Edgeworth Box Representation of Two Agents Trading according to Rule **T**



Analogously, B prefers allocations that are below the blue line, its isocost curve. From any initial endowments we can draw A's and B's isocost curves. For some endowments the area between the curves will be larger than that shown in figure IV-2, while for others it will be smaller. When initial endowments fall on the gray line, the so-called contract curve, the agents' isocost curves are exactly tangent—the *MRSs* of the two agents coincide. At these positions there is zero area between the agents' isocost curves and, as a result, there are no potential gains from trade.²⁰

From point 1, each application of rule **T** moves the agents to progressively higher welfare states, first to position 2, then to 3, and so on until finally they reach position 5. Additional trading, beyond 5, would cause the agents' *MRSs* to cross over, and so is not allowed.²¹ When **T** results in allocations for agents A and B that are on the contract curve, we say

20. For a detailed discussion of the Edgeworth box, see Kreps [1990: 152–53, 155–56].

21. One might reasonably wonder why we have built **T** to take only incremental steps toward the contract curve, instead of jumping directly to it. One rationale is to limit the complexity of our agents; in making small welfare improving trades they use only the local shape of their welfare function in the vicinity of their endowment. They then make a relatively small trade and recompute their marginal valuations with respect to their new holdings.

that *local Pareto optimality* has been achieved.²² When the allocations produced by **T** are just off the contract curve, as in figure IV-2, we say that near Pareto optimality has been attained locally.

Rule **T** specifies *how* two agents interact to trade. It remains to specify *which* agents interact through **T**. All the rules of agent interaction that we have described so far—rules for sexual reproduction, for cultural interchange—involve local interaction, and here we shall not deviate from this pure bottom-up approach. When an agent following **M** moves to a new location it has from 0 to 4 (von Neumann) neighbors. It interacts through **T** exactly once with each of its neighbors, selected in random order.²³

The Sugarscape interagent trade rule can be summarized as follows: If neighboring agents have different marginal rates of substitution then they attempt to arrange an exchange that makes them both better off. Bargaining proceeds and a trade price is “agreed” to. Quantities of sugar and spice in proportion to the trade price are specified for exchange. If exchange of the commodities will not cause the agents’ *MRSs* to cross over then the transaction occurs, the agents recompute their *MRSs*, and bargaining begins anew. In this way nearly Pareto optimal allocations are produced locally.

With these micro-rules in place we are now in a position to study the aggregate or market behavior of neoclassical agents engaged in bilateral trade. How will prices evolve in such markets? Will trade volumes vary regularly or erratically? Rule **T** stipulates that individual agents are made better off through trade, but will the society of agents *as a whole* be able to extract the full welfare benefits of trade? How sensitive will market performance be to neoclassical assumptions about agents? These are the questions to which we now turn.

22. In this usage local Pareto optimality is synonymous with pairwise or bilateral Pareto optimality; see Feldman [1973] and Goldman and Starr [1982].

23. A variant of this would let an agent engage in **T** with a neighboring agent multiple times during a single move. For example, say an agent has 2 neighbors and trades with one of them according to **T**, that is, until they have approximately equal *MRSs*. Then the agent turns to the other neighbor and interacts with it following **T**. After the second set of trades is complete the agent’s *MRS* will be different from what it was at the termination of trade with the first neighbor, and therefore it may be feasible to trade further with this first neighbor. The agent is permitted to do so, and it switches back and forth between its neighbors until no more gains from trade are possible. In this variant, the active agent would act as a kind of arbitrageur between its two neighbors.

Markets of Bilateral Traders

General equilibrium theory describes how a centralized market run by an idealized auctioneer can arrive at an equilibrium price. The immediate question for us—having banished the auctioneer and all other types of nonlocal information—is *whether our population of spatially distributed neoclassical agents can produce anything like an equilibrium price through local interactions alone*. It turns out that there is a definite sense in which they can! However, the character of the equilibrium achieved by our agents is rather different from that of general equilibrium theory, for the markets which result produce less than optimal agent welfare—the potential gains from trade are not fully extracted—despite essential convergence to the general equilibrium price. Furthermore, when we relax certain neoclassical assumptions (infinitely lived agents, fixed preferences) overall market performance is further degraded.

Neoclassical Agents and Statistical Price Equilibrium

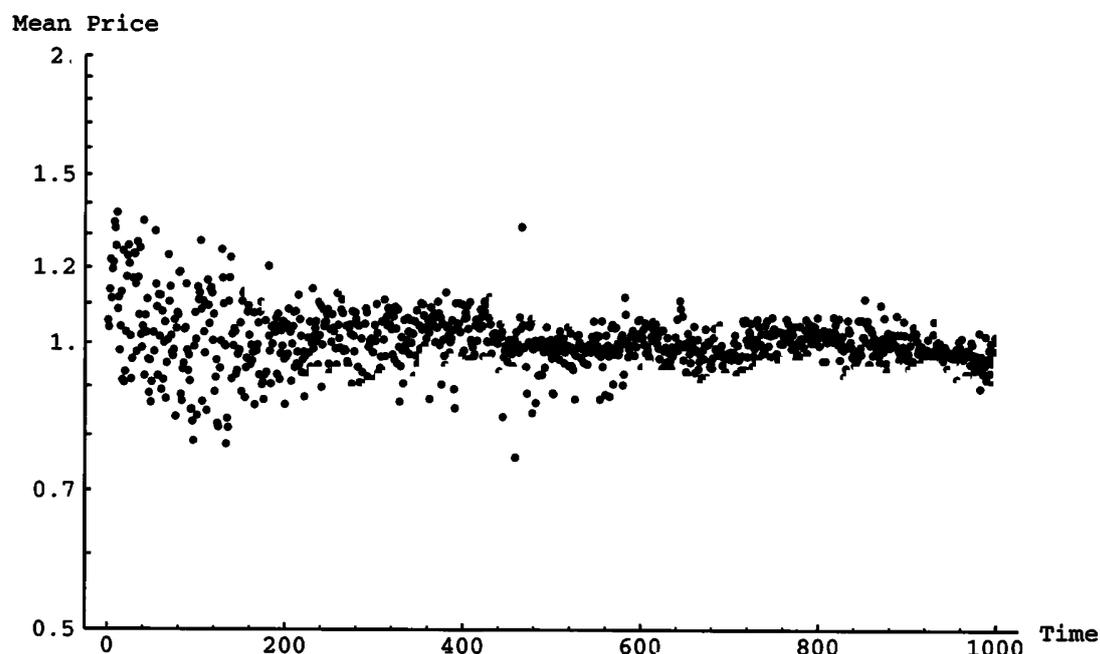
On the sugar-spice landscape we randomly place a population of 200 infinitely lived agents, having Cobb-Douglas utility functions given by (1), with behavioral rules **M** and **T**. Metabolisms for sugar and spice are uniformly distributed in the agent population between 1 and 5. This has the effect of making preferences symmetrical, that is, there are as many agents who prefer sugar to spice as there are who prefer the reverse. Vision is also uniformly distributed between 1 and 5. Initial endowments are randomly distributed between 25 and 50 for both sugar and spice and are thus also symmetrical with respect to the two resources. Therefore, since there is approximately the same amount of sugar and spice present on the landscape, the symmetry of preferences and endowments implies that the general equilibrium price of sugar to spice will be about 1, varying somewhat from time period to time period.²⁴

To display the economic behavior of our artificial market, it would not do to simply “look down from above” on the landscape of agents, as in past animations, since this fails to depict either the formation of prices or the exchange of goods. Instead, we track the time series of average trade price per period.²⁵ Such a plot is shown in figure IV-3.

24. Below we investigate how the general equilibrium price varies, and plot the dynamic supply and demand curves for our artificial economy.

25. We use the phrases “average trade price” and “mean price” to denote the geometric mean of all trade prices that occur in a given period.

Figure IV-3. Typical Time Series for Average Trade Price under Rule System $(\{G_1\}, \{M, T\})$



Note that initially there is significant variation in prices but that over time prices tend to bunch around the “market-clearing” level of 1. The total volume of trade is quite large, with nearly 150,000 trades occurring over the time shown in figure IV-3. There is extensive variation in trade volume per period, as shown in figure IV-4. Trade volumes are distributed approximately lognormally, with a few big trade periods and lots of smaller ones.²⁶

Another way to look at how prices converge toward the general equilibrium level is to plot the standard deviation in the logarithm of the average trade price per period. For the previous run, this is shown in figure IV-5. Here, and in all subsequent plots of price standard deviation time series, raw data are shown in black with smoothed data in red.

While the standard deviation in price never vanishes, it does tend to stabilize at a relatively small value, averaging about 0.05 by $t = 1000$.²⁷ In this case it would seem unobjectionable to say that a price equilibrium is essentially attained by this market. *Economic equilibrium emerges from the bottom up.*

26. Actually, the distribution of trade volume is nonstationary when agent lifetimes are infinite.

27. Because agents are infinitely lived in this run of the model, the standard deviation in price will never reach a stationary value but will continue to fall.

Figure IV-4. Typical Time Series of Trade Volume under Rule System $(\{G_1\}, \{M, T\})$

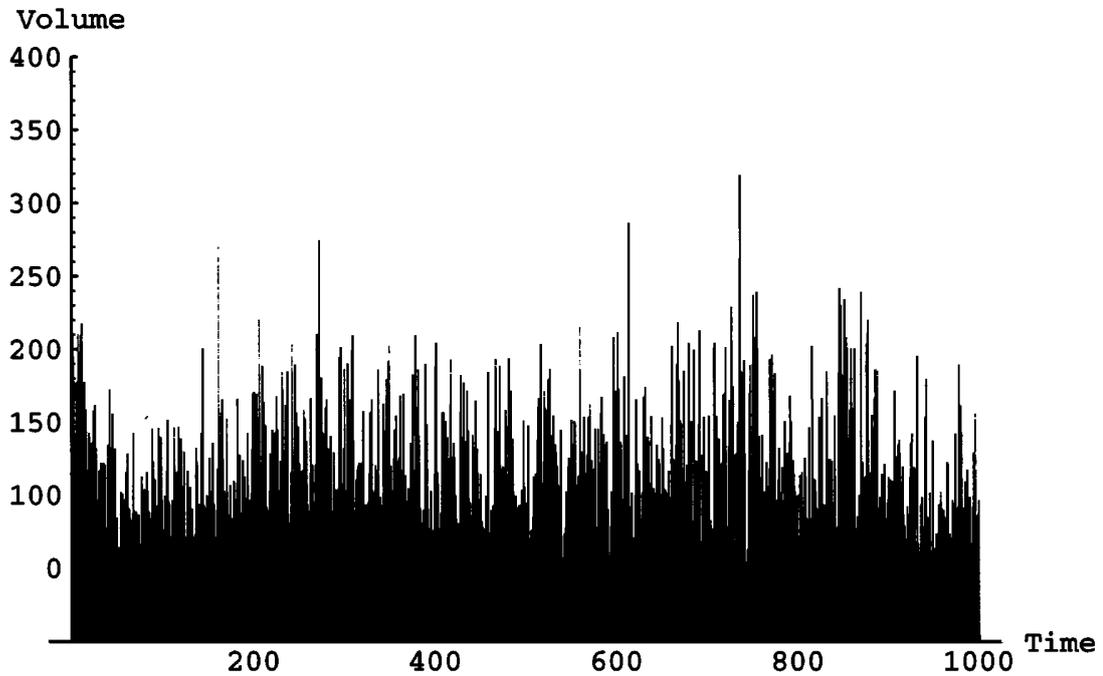
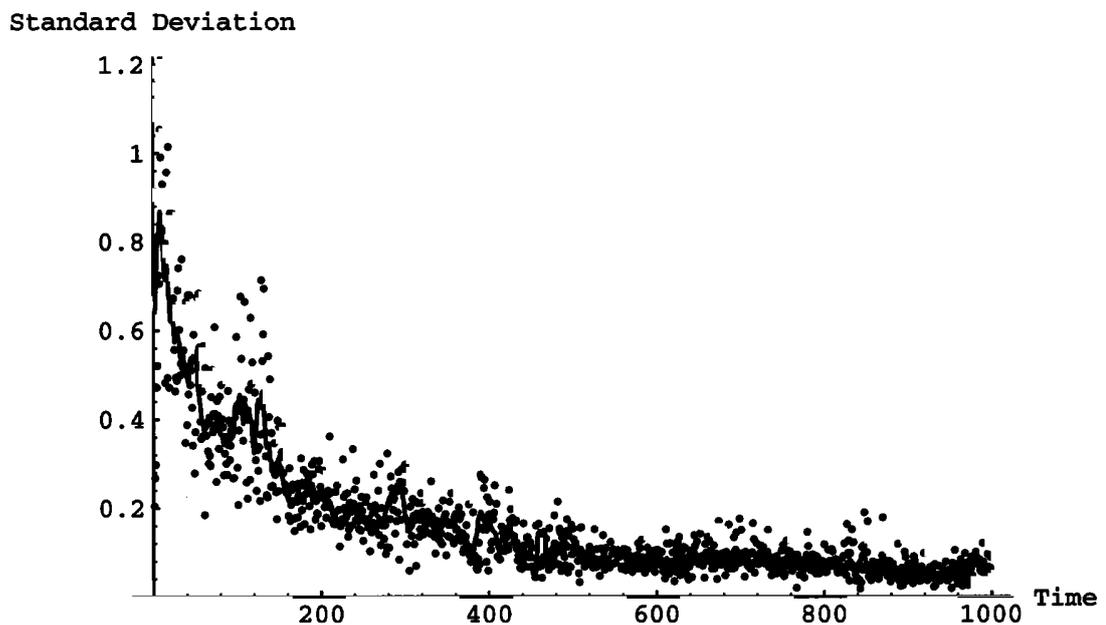


Figure IV-5. Typical Time Series for the Standard Deviation in the Logarithm of Average Trade Price under Rule System $(\{G_1\}, \{M, T\})$



The Invisible Hand

There is a sense in which this completely decentralized, distributed achievement of economic equilibrium is a *more* powerful result than is offered by general equilibrium theory, since *dynamics* of price formation are fully accounted for, and there is no recourse to a mythical auctioneer. This result harks back to Adam Smith and the classical economists whose image of markets involved no such entity.²⁸

Performance of Markets Produced by Neoclassical Traders

Having seen typical price-volume time series for markets of neoclassical agents engaged in bilateral trade, we now investigate the nature of these markets.

Carrying Capacity Is Increased by Trade

In Chapter II we found that the notion of *carrying capacity* emerged naturally on the sugarscape.²⁹ Here we study the effect of trade on the carrying capacity. We do this by noting the number of agents who survive in the long run, first with trade turned off, then with it turned on. Figure IV-6 is a plot of the dependence of carrying capacity on average agent vision, the lower line representing the no-trade case, the upper line the with-trade case.

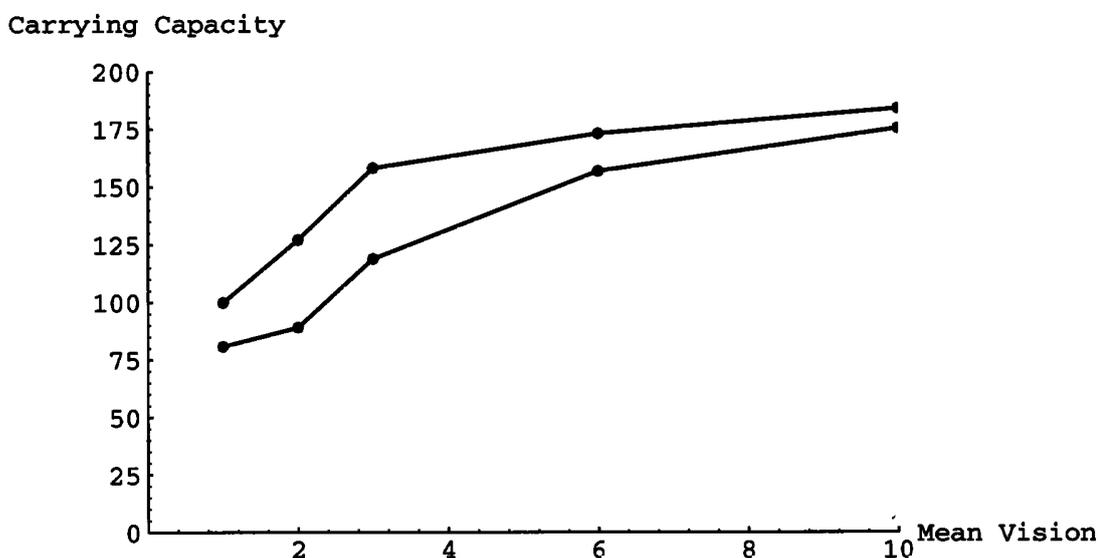
Clearly, trade *increases* the carrying capacity. This result is in accord with intuition. It was argued earlier that trade was a way for agents to avoid death due to a deficiency in one commodity. To see how this is so, imagine a pair of neighboring agents. Agent 1 has an abundance of sugar but is close to death by spice deprivation; Agent 2 has a surfeit of spice but is on the verge of death through sugar deprivation. If trade is forbidden then each will die. Clearly, however, an exchange of Agent 1's sugar for Agent 2's spice will keep both alive. This is how trade increases the carrying capacity of the sugar-spice scape.

Let us now discuss the nature of the equilibrium produced by com-

28. It is usual in economics to associate the Smithian invisible hand with welfare properties of markets, and we do this below. Our usage here has more in common with what Nozick [1974, 1994] calls an "invisible-hand process." For an excellent discussion of Smith's varied usage of the term "invisible hand," see Rothschild [1994].

29. In figure II-5 we presented the dependence of carrying capacity on vision and metabolism distributions in the agent population. As average vision increased and mean metabolism decreased, the carrying capacity increased.

Figure IV-6. Carrying Capacity as a Function of Mean Agent Vision, with and without Trade, under Rule System $(\{G_1\}, \{M, T\})$



pletely decentralized trade. It is of a profoundly different character than the Walrasian general equilibrium.

Statistical Equilibrium

The equilibrium concept used in general equilibrium theory is a deterministic one. That is, once the auctioneer announces the market-clearing price vector, all agents trade at exactly these prices. Each agent ends up with an allocation that cannot be improved upon. That is, a Pareto-optimal set of allocations obtains. Because these allocations are optimal, no further trading occurs and the economy is said to be in equilibrium. Overall, equilibrium happens in a single trade step.³⁰

In the model of bilateral exchange described above, each agent trades not at the general equilibrium price but rather at a locally negotiated one. Imagine that it is some particular agent's turn to move, and you must predict the exact price at which its next trade will occur. This price depends not only on that agent's own internal valuation (*MRS*) but also on that of its trading partner. Predicting the actual trade price involves predicting who this neighbor is likely to be, that agent's *MRS*, and so on. With anything less than a complete description of the entire state space of the artificial society, this calculation can only be made probabilistically.

Recently, Foley [1994] has advanced a novel theory of statistical eco-

30. For a classic exposition, see Arrow and Hahn [1971].

conomic equilibrium that has much in common with economic behavior observed in our model.³¹ He has argued that general equilibrium theory is “methodologically too ambitious” in that it attempts to compute the allocation for each agent exactly. Indeed, such computations would seem intractable in the relatively simple case of our artificial economy, to say nothing of the real world.

This brings us to the so-called First Welfare Theorem of neoclassical economics.³² This result is the foundation for economists’ claims that markets allocate goods to their optimal social uses. The theorem states that Walrasian equilibria are Pareto-efficient. They are states in which *no reallocation exists such that an agent can be made better off without making at least one other agent worse off*. But in statistical equilibrium

the First Welfare Theorem should be revised to say that a market equilibrium approximates but cannot achieve a Pareto-efficient allocation. How close a given market comes to Pareto-efficiency can be measured by the price dispersion in transactions. [Foley 1994: 343]

It is exactly this price dispersion that we studied above and will investigate further below in the context of non-neoclassical agents. Thus the philosophical underpinning for laissez-faire policies appears to be weak for markets that display statistical equilibrium.³³

Horizontal Inequality

Foley [1994] has introduced the term *horizontal inequality* to describe the fact that agents having identical abilities (vision in our model), preferences (parameterized by metabolism in Sugarscape), and endowments will generally have different welfare levels in statistical equilibrium, a phenomenon that is strictly prohibited in Walrasian general equilibrium.

Differences in final consumption and welfare in Walrasian competitive equilibrium always correspond to differences in initial endowments. But trading at different price ratios leads agents with the same initial endowments to different consumption and utility levels. [Foley 1994: 342]

31. This is to be distinguished from the theory of stochastic general equilibrium under incomplete information; for a review see Radner [1982].

32. See Varian [1984: 198–203] for the welfare properties of Walrasian equilibria.

33. The First Welfare Theorem is commonly referred to as the “invisible hand theorem” [Stokey and Lucas 1989: 451–54]. This suggests that decentralized trade must arrive—as if “led by an invisible hand” [Smith 1976: 456]—at Walrasian equilibrium. Our market of decentralized trade certainly does not arrive there.

In other words, the welfare properties of neoclassical general equilibrium markets are *not* preserved in statistical equilibrium, due to the production of horizontal inequality. So once again the *character* of the equilibrium in our model turns out to differ markedly from that in the orthodox theory of general equilibrium.³⁴ In fact, we expect the production of horizontal inequality to occur in proportion to the variance or dispersion in price in statistical equilibrium. Later it will be shown that such dispersion can be very large indeed.

Local Efficiency, Global Inefficiency

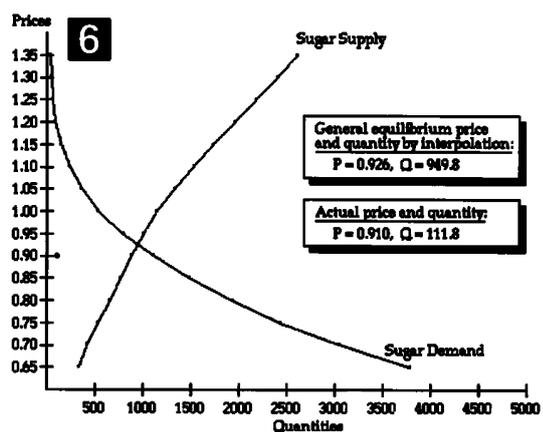
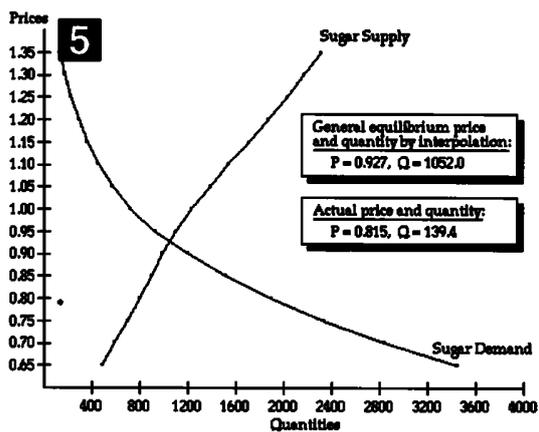
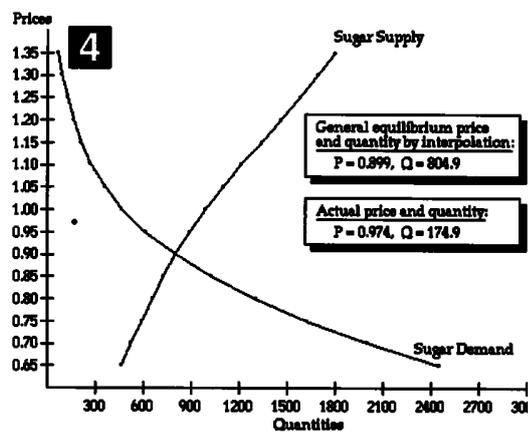
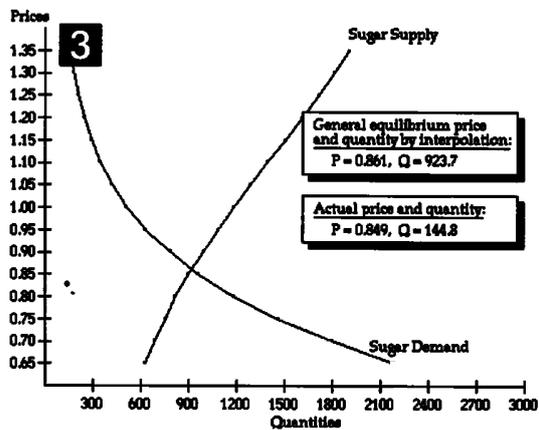
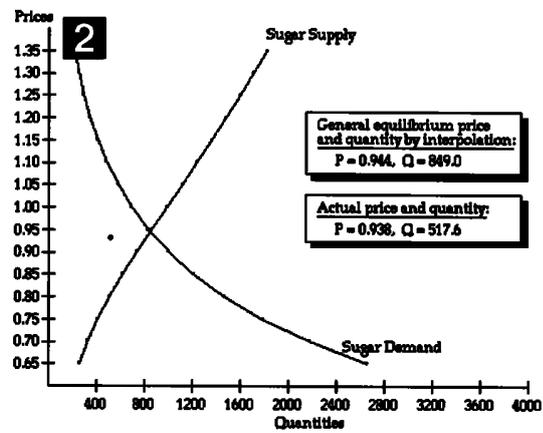
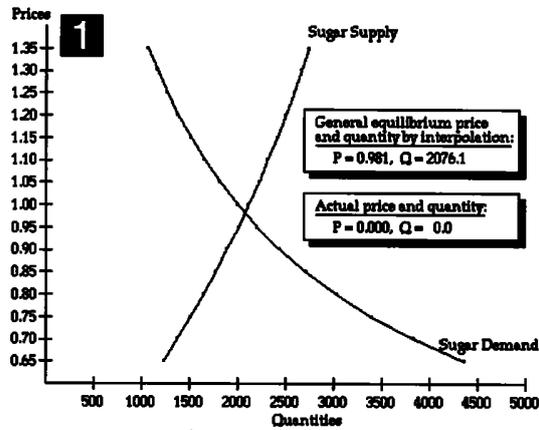
The statistical character of the *price* equilibrium produced by bilateral trade algorithm **T** is very different from the usual general equilibrium notion. It is also true that the *quantities* traded are always different from those that would obtain were the system in general equilibrium. To see this we can make a supply-demand plot for our artificial economy. This is done by querying individual agents as to the quantity of sugar each is willing to supply or demand at a given price. Summing these quantities yields the aggregate supply and demand schedules. The general equilibrium price and quantity can then be computed by interpolation and compared (noiselessly) to the actual (average) trade price and (total) quantity exchanged. Furthermore, these computations can be repeated each period and animated. This is done in animation IV-2 for an artificial economy like the one described in figures IV-3, IV-4, and IV-5.

Notice that while the actual price moves around the general equilibrium price, the actual quantity traded is *always less* than what is necessary to “clear the market.” Since agents are unable to trade with anyone other than their neighbors, there is always some “pent-up” demand that goes unfulfilled. That is, if the agents were perfectly mixed, they would engage in additional trades beyond what they achieve through **T**. Over time, as the agents move around, they do meet and interact with these other agents. However, as they move they are accumulating additional goods that they are willing to trade, thus shifting the equilibrium fur-

34. The Second Welfare Theorem of neoclassical economics, like the first one, needs to be modified in statistical equilibrium. It states that any Pareto-efficient allocation can be achieved by a Walrasian equilibrium price vector given an appropriate reallocation of endowments. However, in statistical equilibrium

unless the [initial] endowment can be redistributed directly to the Pareto-efficient allocation, in which case there is nothing for the market to do, the generation of endogenous horizontal inequality among agents appears to be an inescapable by-product of the allocation of resources through decentralized markets. [Foley 1994: 343]

Animation IV-2. Evolution of Supply and Demand under Rule System $(\{G_1\}, \{M, T\})$



ther. The decentralized economy is always far from general equilibrium in this sense.³⁵

This result is of prime significance. For whenever the actual trade volumes are less than the general equilibrium ones, agent society is not extracting all the welfare from trade that it might. If the agents could coordinate their activities beyond their local neighborhoods they could all be made better off. Here we see that *even though T produces exchanges that are nearly Pareto-optimal locally, the resulting market has far from optimal welfare properties globally.*³⁶

Far from Equilibrium Economics

A very general principle is lurking here. At each time step agents engage in production (resource gathering according to **M**) and consumption activities as well as pure exchange with their neighbors according to **T**. Because the exchange rule requires time to reach an equilibrium societal allocation, it essentially gives production and consumption time to alter the equilibrium to which **T** is converging. That is, as production and consumption modify endowments they also modify the target the exchange process is trying to achieve. The result is that the economy is far from equilibrium in a very definite sense. These circumstances—exchange taking time to converge while production and consumption constantly shift the equilibrium—are *sufficient* conditions for the existence of a nonequilibrium economy.³⁷

35. In animation IV-2 there is an increase in both the actual trade volume and the general equilibrium volume over time (the former always lagging the latter). This nonstationarity is due to the infinite livedness of the agents.

36. The far-from-equilibrium character of this spatially distributed market is an interesting result from the perspective of prices as signals appropriate for decentralizing decision-making. Although the market has not reached general equilibrium it is essentially generating the general equilibrium price (though our agents, following **T**, do not use this signal). There are at least two implications of this. First, “getting the price right” is not sufficient to guarantee allocative efficiency. The second conclusion is of a different character. In certain markets it may be that agents use local information exclusively in their economic decisionmaking. In such markets aggregate data such as average prices, a primary focus of economists’ attentions, are simply emergent statistically from micro-heterogeneity and of no particular interest to the agents.

37. Fisher [1983: 14] makes a similar point: “In a real economy . . . trading, as well as production and consumption, goes on out of equilibrium. It follows that, in the course of convergence to equilibrium (assuming that occurs), endowments change. In turn this changes the set of equilibria. Put more succinctly, the set of equilibria is path dependent—it depends not merely on the initial state but on the dynamic adjustment process”

Effect of the Distribution of Agent Vision on Price Variance

In figure IV-5 the variance in trade price decays to a relatively small value. Initially, the agents' endowments may have little to do with their preferences (since both endowments and preferences are randomly assigned). Hence, when they encounter one another they may trade at prices far from the general equilibrium level. But exchange serves to bring their internal valuations (*MRSs*) closer together. Over time, the dispersion in *MRSs* decreases as agents increasingly encounter others with *MRSs* similar to their own. However, as described above, the processes of production and consumption make complete convergence impossible, and so some price variance persists indefinitely.

One can get significantly larger amounts of price variance by making the market "thinner." For example, when agent interactions are restricted, less trade occurs, price convergence slows, and there results a broader distribution of *MRSs* in the economy. There are a variety of ways to produce such thin markets on the sugarscape. Here we investigate the effect of agent vision on the speed of price convergence.

In the run of the model described in figures IV-3, IV-4, and IV-5, agent vision was uniformly distributed between 1 and 5. If we reduce vision to 1 across the entire agent society, then the agents will move around much less and there will be more price heterogeneity. This is depicted in figure IV-7 where the annual mean price is displayed.

The average price over the roughly 100,000 trades that occur during this period is 1.0, quite close to the general equilibrium level. But nothing like the "law of one price" obtains. This is displayed more clearly by a plot of the standard deviation in the natural logarithm of per period mean prices (see figure IV-8). While the standard deviation trends downward, there is significantly more variation in the price than encountered in figure IV-5. In short, nothing like general equilibrium obtains here.

Price variance is a feature of real-world markets. The amount of price dispersion in any particular market is, of course, an empirical question. While we do not purport to be modeling any particular market here, the degree of price heterogeneity displayed in figure IV-8 is of the same magnitude as that observed in econometric studies of price dispersion.³⁸

What matters is the equilibrium that the economy will reach from given initial conditions, not the equilibrium that it would have been in, given initial endowments, had the prices happened to be just right." See also Negishi [1961] and Hicks [1946: 127-29].

38. These include Carlson and Pescatrice [1980] and Pratt, Wise, and Zeckhauser [1979]. Economists seek to explain persistent price dispersion in terms of imperfectly

Figure IV-7. Typical Time Series for Average Trade Price under Rule System $(\{G_1\}, \{M, T\})$, with Agent Vision Set at 1

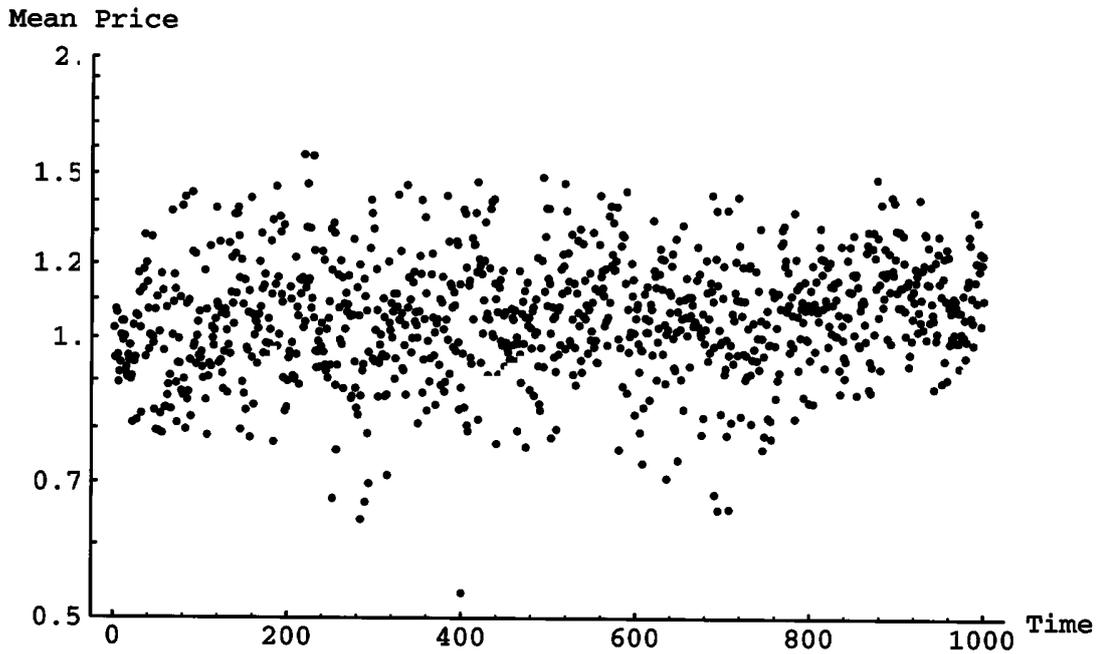


Figure IV-8. Typical Time Series for the Standard Deviation in the Logarithm of Average Trade Price under Rule System $(\{G_1\}, \{M, T\})$, with Agent Vision set at 1

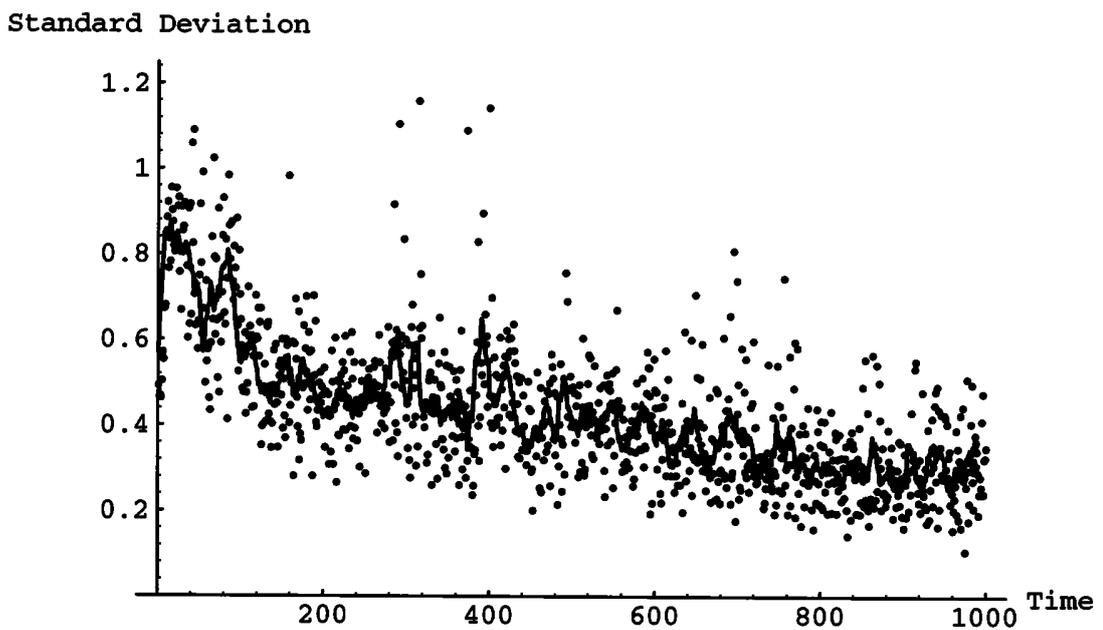
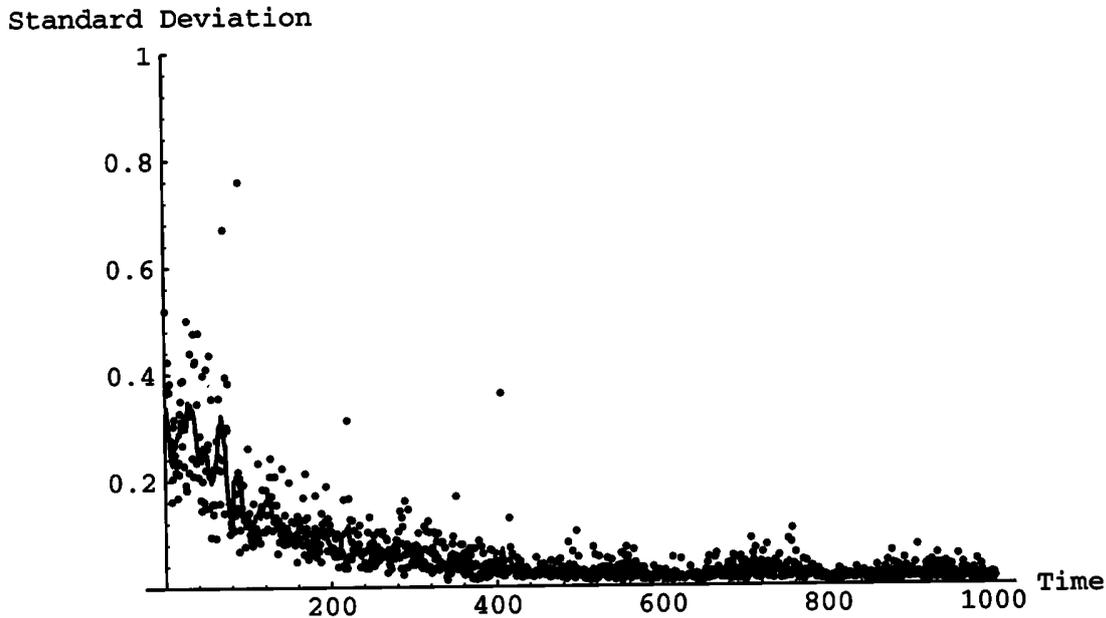


Figure IV-9. Typical Time Series for the Standard Deviation in the Logarithm of Average Trade Price under Rule System $(\{G_1\}, \{M, T\})$, with Agent Vision Uniformly Distributed between 1 and 15



It is also possible to create markets on the sugarscape that have much *less* price variance than that shown in figure IV-5. When average agent vision is *large*, price heterogeneity *decreases*. Figure IV-9 gives a time series for the standard deviation in the logarithm of mean prices in a population in which vision is uniformly distributed between 1 and 15.

Here, due to higher mean vision, there is much more intense interaction—more perfect mixing—of the agent population and therefore equilibrium is approached quickly. By contrast to the preceding case (low agent vision, high price variance), the artificial market of figure IV-9 more closely resembles the information-rich environment of, for example, financial markets.

Non-Neoclassical Agents and Further Departures from Equilibrium

Up to now our agents, endowed with fixed preferences and infinite lives, have been basically neoclassical. In agent-based models like Sugarscape it is not difficult to relax these assumptions. In what follows we make

informed consumers who engage in (costly) search for the best prices [Ioannides 1975, Reinganum 1979]. Our model is not a search model, yet it also yields price dispersion.

our agents more human, first, by giving them finite lives and, second, by permitting their preferences to evolve. We shall see that the effect of these new rules is to add variance to the distribution of prices and to modify the price itself. In fact, the mean price will follow a kind of “random walk.”

Finite Lives: Replacement

In Chapter II we introduced finite death ages into the agent population for purposes of studying the wealth distributions that emerged under rule **M**. The replacement rule $\mathbf{R}_{[a,b]}$ denotes that the maximum agent age is uniformly distributed over interval $[a,b]$.

In the context of trade, the effect of the replacement rule is to add agents to the population who initially have random internal valuations, that is, *MRSs* quite distant from the price levels that prevail. A new agent is born into the world with an initial endowment uncorrelated with its wants. It seeks, through trade, to improve its welfare by bringing its endowments into line with its needs. That is, an agent with high sugar metabolism and low spice metabolism wants to accumulate much larger stocks of sugar than of spice. When agent lifetimes are relatively short in comparison with the time required for the distribution of *MRSs* to homogenize, high price variance will result. An example of this is illustrated in figure IV-10, a plot of the standard deviation in the logarithm of annual average trade prices in the case of maximum age distributed uniformly between 60 and 100, and vision returned to its earlier distribution (uniform between 1 and 5). Clearly, this straightforward departure from the neoclassical agent produces market performance at considerable variance with Walrasian general equilibrium.³⁹

As the average agent lifetime grows there is more time for young agents to have their internal valuations brought into line with the overall market.⁴⁰ So the price dispersion decreases as mean agent lifetime increases. This effect is shown in figure IV-11 where agent maximum ages are distributed uniformly between 980 and 1020.⁴¹

39. At any instant in this finitely lived agent economy it is certainly the case that equilibria *exist*. Figure IV-10 demonstrates that such equilibria will not generally be *achieved*.

40 One might argue that in the real world the issue is not agent lifetimes per se, but rather the duration of agents' participation in markets. Of course, in Sugarscape all agents who are alive participate in the market through rule **T**.

41. Note that the variance in agent lifetimes is identical in figures IV-10 and IV-11.

Figure IV-10. Typical Time Series for the Standard Deviation in the Logarithm of Average Trade Price under Rule System $(\{G_1\}, \{M, R_{[60,100]}, T\})$

Standard Deviation

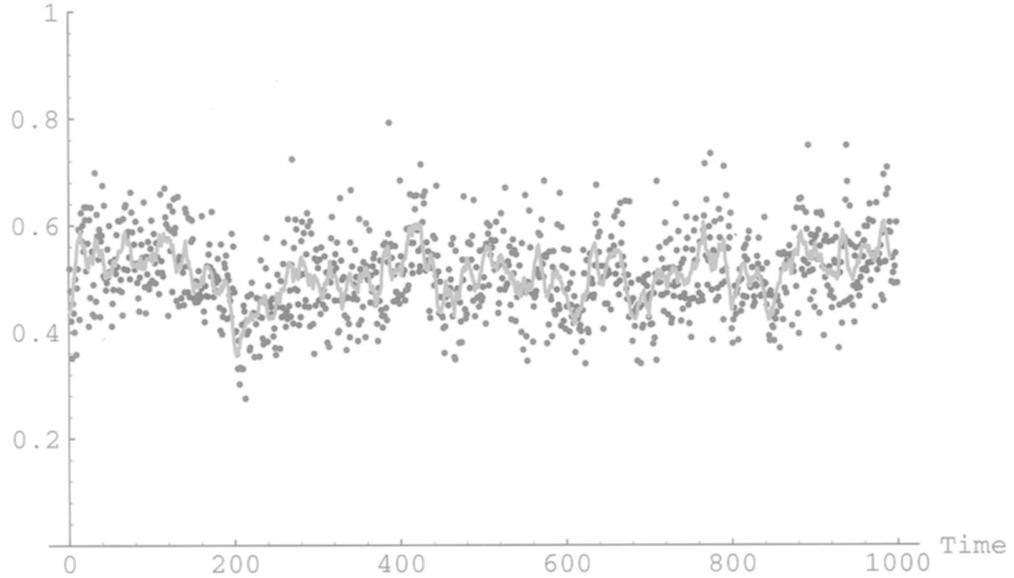


Figure IV-11. Typical Time Series for the Standard Deviation in the Logarithm of Average Trade Price under Rule System $(\{G_1\}, \{M, R_{[960,1000]}, T\})$

Standard Deviation

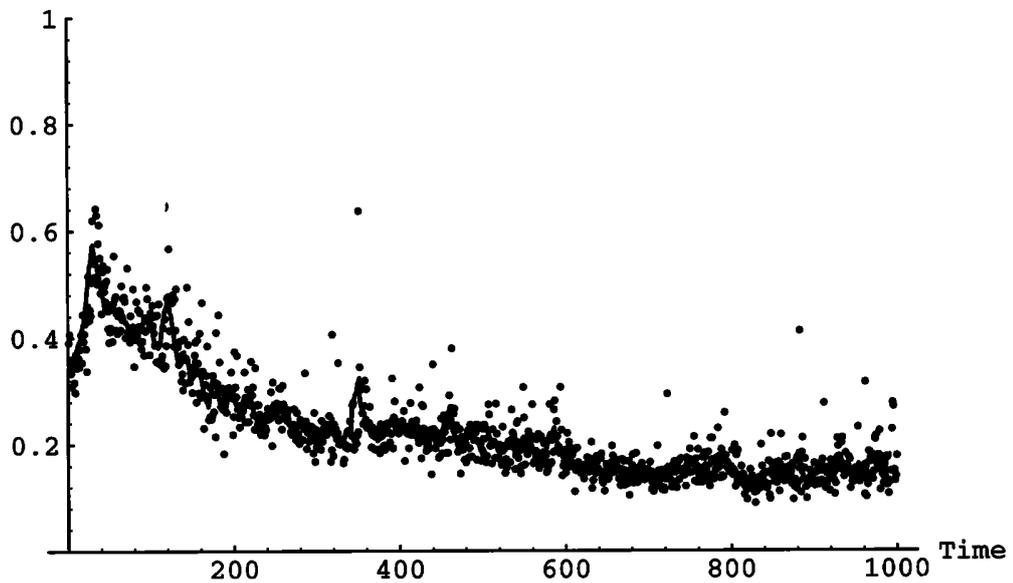
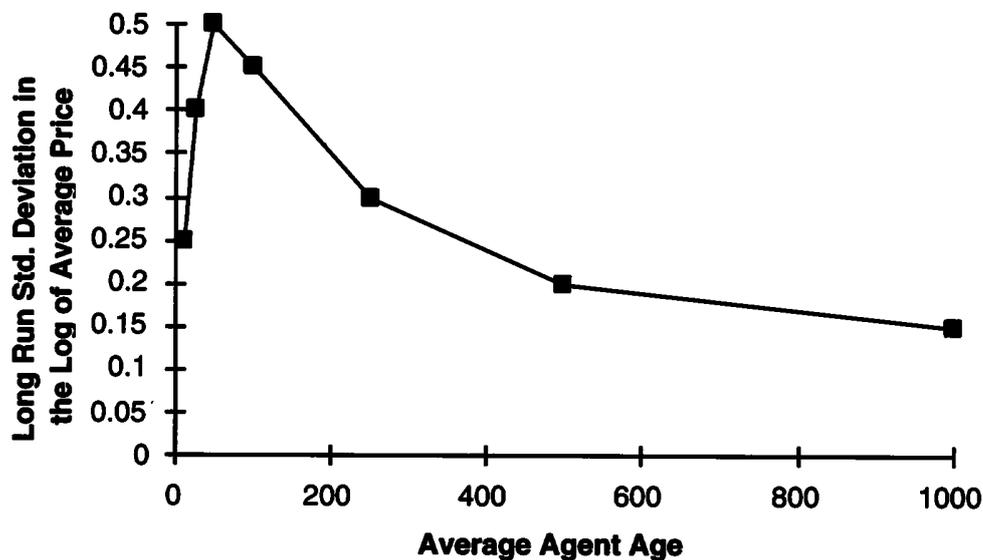


Figure IV-12. Dependence of the Long-Run Standard Deviation in the Logarithm of Average Trade Price on Average Lifetime under Rule System $(\{G_1\}, \{M, R_{[a,b]}, T\})$



We have studied this effect for a variety of agent lifetime specifications and summarize the results in figure IV-12.

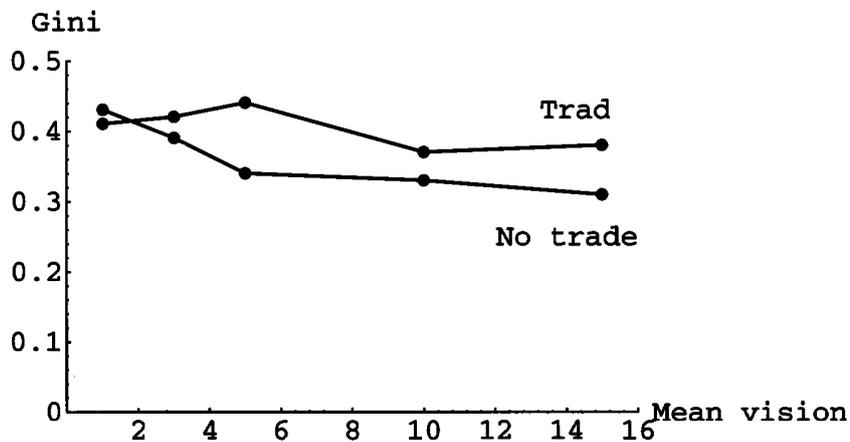
Equity

In Chapter II, highly skewed distributions of wealth were observed for agents following movement rule *M*. How is the *distribution* of wealth altered by trade? In particular, is society made more or less equitable by trade? Now that finite lives and agent replacement have been reintroduced this question is conveniently studied.⁴² One familiar measure of equity is the Gini coefficient, *G*, illustrated in animations II-4 and III-4, where it was displayed along with Lorenz curves. In figure IV-13 the dependence of *G* on trade is displayed.

Overall, the effect of trade is to *further skew* the distribution of wealth in society. So, while trade increases the carrying capacity, allowing more agents to survive, it also increases the inequality of the wealth distribution. In this sense, there is a tradeoff between economic equality and economic performance.

42. It is not possible to study wealth distributions in the context of infinitely lived agents since such distributions are nonstationary.

Figure IV-13. Dependence of the Gini Coefficient on Trade, Parameterized by Mean Vision and Mean Metabolism, under Rule System $(\{G_1\}, \{M, R_{[60,100]}, T\})$



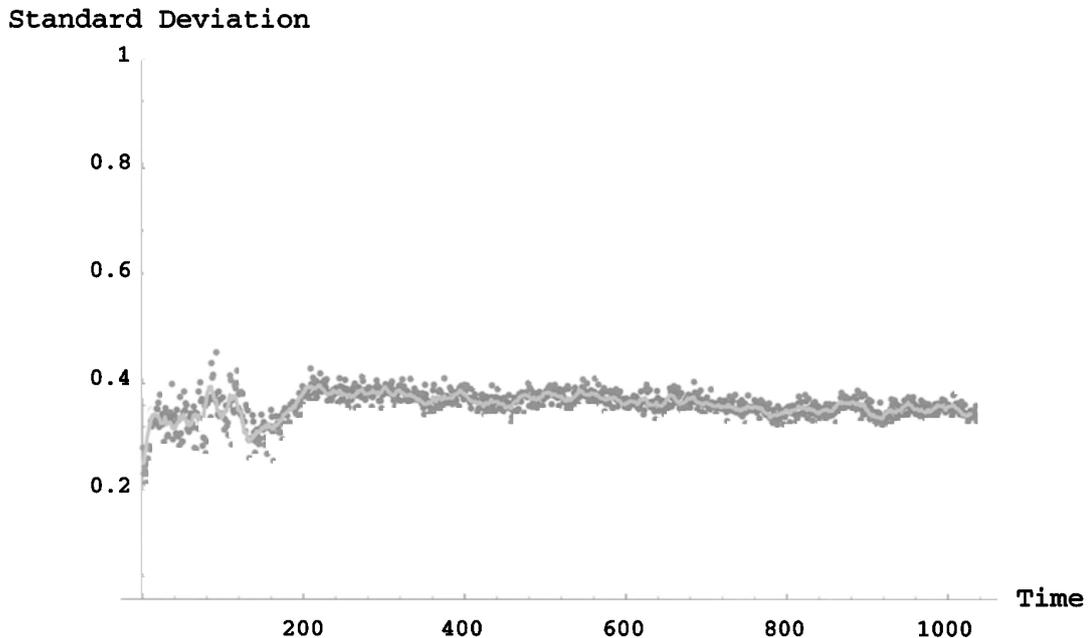
Finite Lives: Sexual Reproduction

When agents reproduce sexually via rule **S**, described in Chapter III, each new agent's preferences are the result of "cross-over" of its parents' preferences. With **S** turned on we expect—as in the case of agent replacement through $R_{[a,b]}$ —an increase in the price variance due to the continual introduction of novel agents (with random internal valuations) into the society. In figure IV-14 the onset of puberty is a random variable in the interval $[12, 15]$, the range of ages at which childbearing ends is $[35, 45]$ for women and $[45, 55]$ for men, and the maximum agent age is selected from $[60, 100]$.

Again, a persistent high level of price dispersion is observed. Overall, the effect of finite lives—with replacement or sexual reproduction—is to push the market away from anything like general equilibrium.

As shown in Chapter III, evolutionary processes are at work whenever the agents engage in sexual reproduction, modifying the distribution of vision and metabolism in the agent population. Therefore economic preferences are systematically varying on evolutionary time scales when **S** is operational. This is so since the distributions of metabolisms in the agent population are changing, as in figure III-2, and these metabolisms enter directly into the agent welfare functions. This is a kind of "vertical transmission" of preferences. We now consider the "horizontal transmission" of preferences.

Figure IV-14. Typical Time Series for the Standard Deviation in the Logarithm of Average Trade Price under Rule System $(\{G_1\}, \{M, S, T\})$



Effect of Culturally Varying Preferences

It is usual in neoclassical economics to assume fixed, exogenously given, agent preferences. The preferences of our agents, as manifested in the welfare function, although state-dependent, are fixed in the sense of being dependent on each agent's unchanging biological needs (metabolisms). It seems clear that, in fact, preferences do evolve over the course of an agent's life, as a function of contacts with other agents.⁴³ Imagine that the only foods are peanuts and sushi. Though born into a family of pure peanut eaters in Georgia, one might acquire a taste for sushi on a trip to Japan.

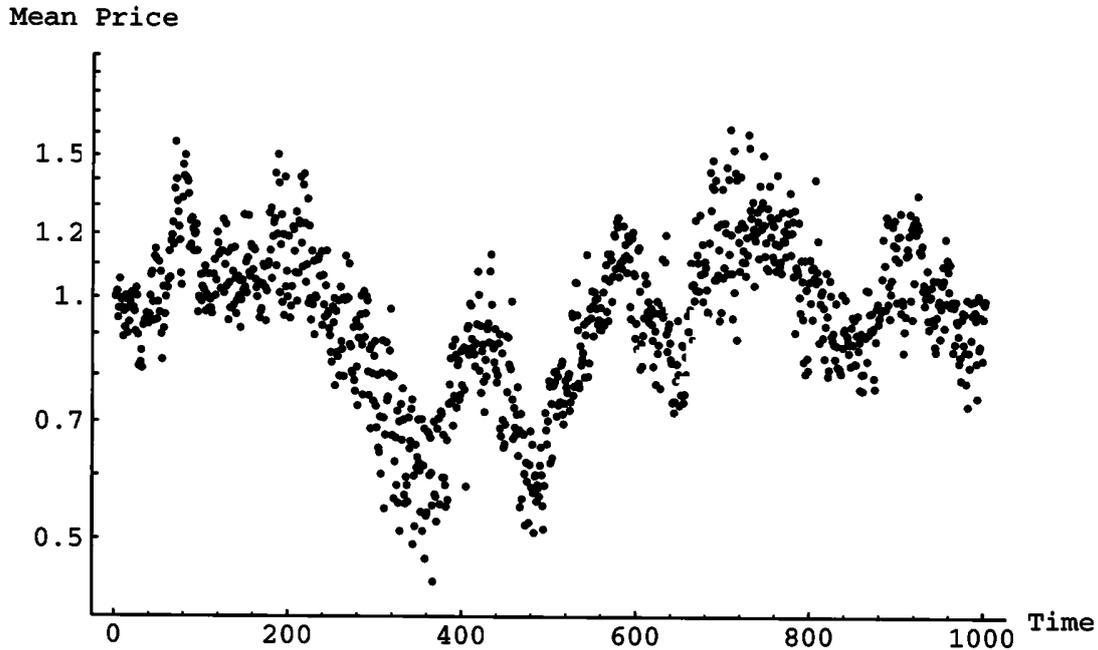
Here we let economic preferences vary according to the state of an agent's cultural tags.⁴⁴ By making agents' preferences depend on cultural variables, welfare functions evolve endogenously. In particular, call f the fraction of an agent's tags that are 0s; then $(1 - f)$ is the fraction of 1s.⁴⁵ We let these enter the welfare function according to

43. There is a large literature on preference formation and change, including Peleg and Yaari [1973], Stigler and Becker [1977], Cowen [1989, 1993], Karni and Schmeidler [1989], and Goodin [1990].

44. The cultural interchange machinery was introduced in Chapter III.

45. Note that the definition of group membership given in Chapter III can be stated as follows: if $f < 1/2$ then the agent belongs to the Red tribe; if $f > 1/2$, the Blue tribe.

Figure IV-15. Typical Time Series for Average Trade Price under Rule System $(\{G_1\}, \{M, K, T\})$



$$W(w_1, w_2) = w_1^{\frac{m_1 f}{\mu}} w_2^{\frac{m_2 (1-f)}{\mu}}, \quad (5)$$

where $\mu = m_1 f + m_2(1 - f)$. Thus, when cultural transmission processes are active, preferences evolve over time, yet at each instant the Cobb-Douglas algebraic form is preserved.⁴⁶

Figure IV-15 gives a typical annual average price time series when infinitely lived agents are governed by (5) and cultural transmission rule **K** is operational.

Note that now the mean price follows a kind of random walk. This occurs because culture is continuously evolving and therefore preferences are constantly changing.⁴⁷ There is also significant price dispersion, as shown in figure IV-16.

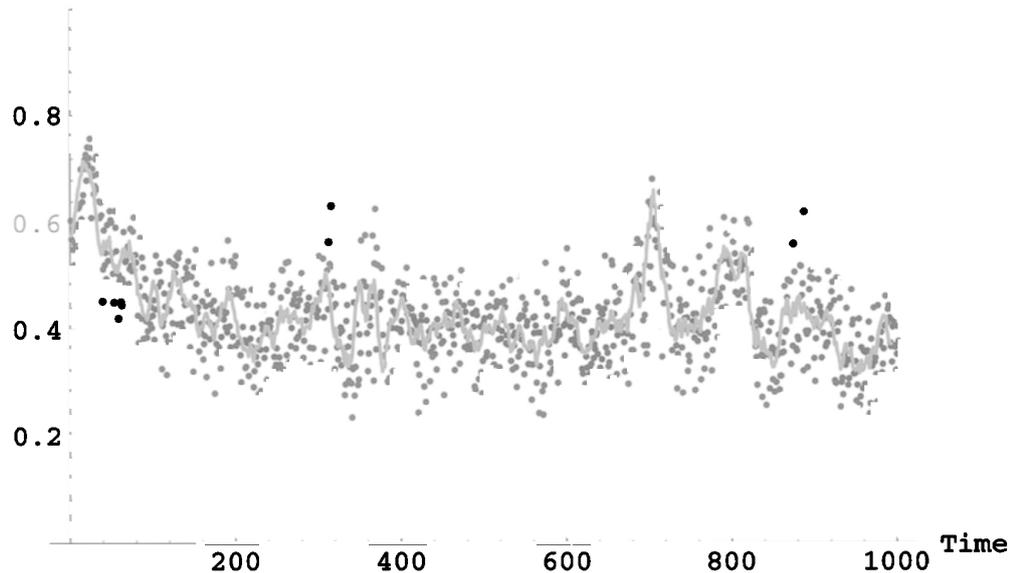
Note that the variance in price never settles down. Also, at 10^6 transactions, the volume of trade in this run is larger (by a factor of roughly 5)

46. For another use of binary strings to model evolving preferences, see Lindgren and Nordahl [1994: 93–94].

47. Note that through culturally varying preferences, an agent's biological (metabolic) requirements can be eclipsed by cultural forces. For example, in the case of f near 0, an agent virtually neglects its need for sugar and, unless f increases later, the agent may die from sugar starvation.

Figure IV-16. Typical Time Series for the Standard Deviation in the Logarithm of Average Trade Price under Rule System $(\{G_1\}, \{M, K, T\})$

Standard Deviation



than in figure IV-4. This is because as an agent's preferences change it finds itself holding goods that it no longer values highly. Or, as Shakespeare's Benedick asks, ". . . but doth not the appetite alter? A man loves the meat in his youth that he cannot endure in his age."⁴⁸

Let us now turn sexual reproduction (**S**) on as well, so that preferences change both "vertically" and "horizontally." A typical time series for the price standard deviation is shown in figure IV-17.

The combined effect of finite lives and evolving preferences is to produce so much variation in price that equilibrium seems lost forever.

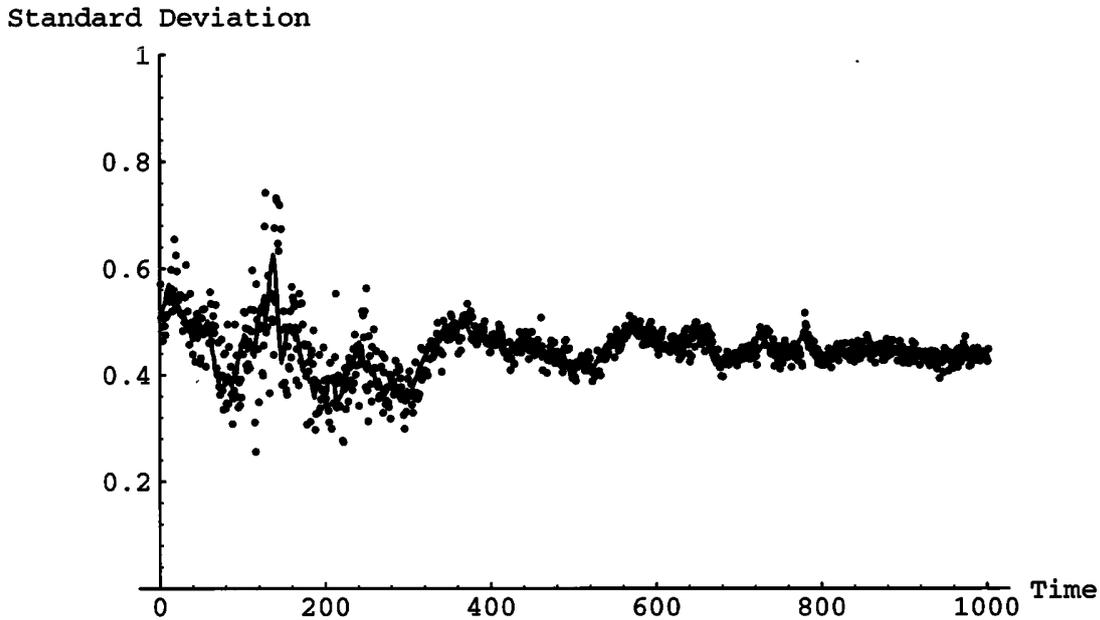
In summary, our quite realistic departures from the neoclassical model of individual behavior produce dramatic departures from the textbook picture of overall market performance. Now we turn to another topic, restoring the neoclassical assumptions.

Externalities and Price Disequilibrium: The Effect of Pollution

In Chapter II we introduced pollution onto the sugarscape. There we were concerned with the effect of pollution on agent movement. When we turned pollution on and allowed it to accumulate (no diffusion), agents

48. From *Much Ado About Nothing*, Act II, Scene III.

Figure IV-17. Typical Time Series for the Standard Deviation in the Logarithm of Average Trade Price under Rule System $(\{G_1\}, \{M, S, K, T\})$



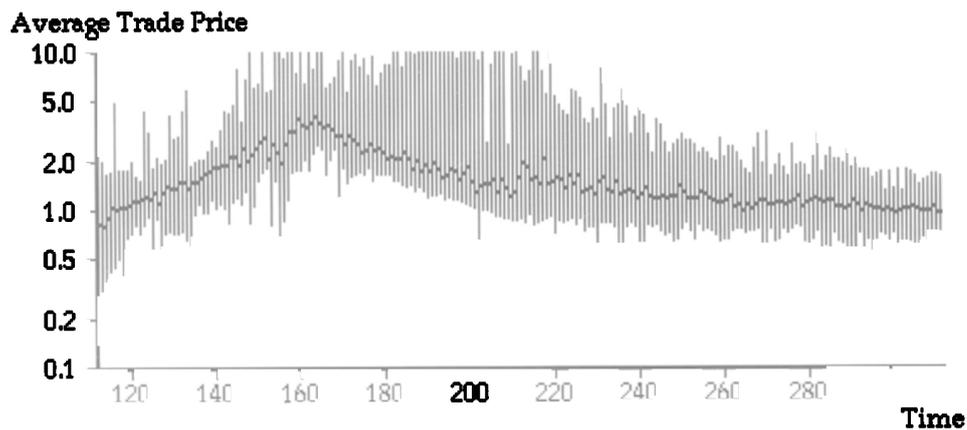
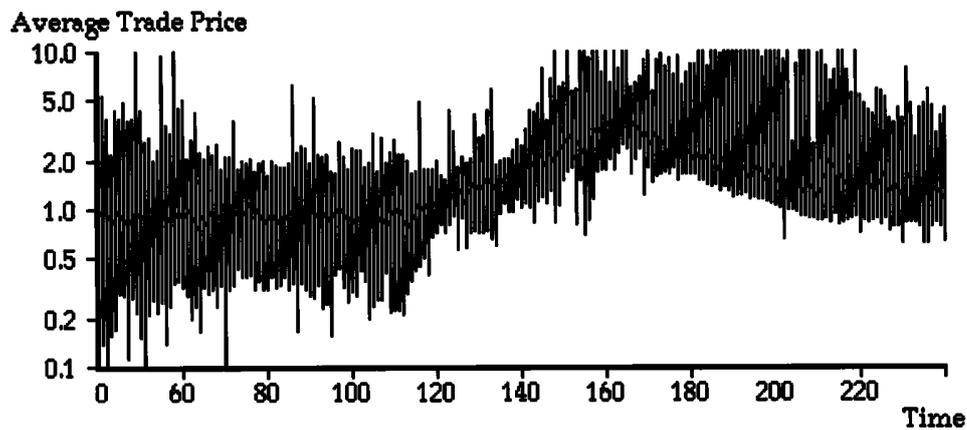
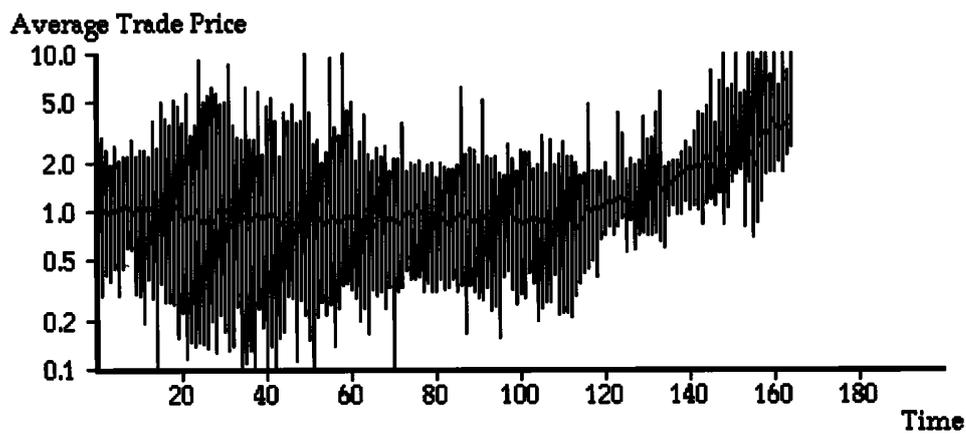
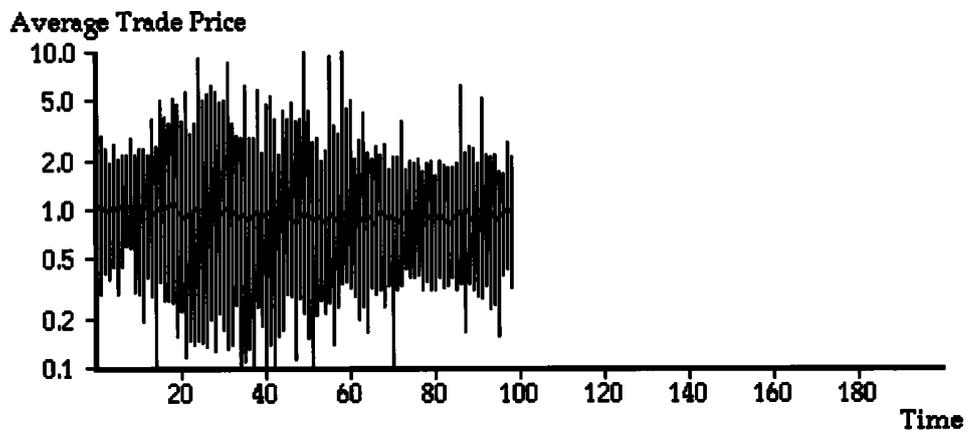
migrated from the polluted area. When we turned diffusion on, the pollution dissipated and agents moved back into the abandoned zones.

Having developed a model of bilateral trade in this chapter, we are now in a position to explore the effect of pollution on prices. To economists, environmental pollution is the classic negative externality. Externalities are important since their existence is an indication that an economy is not achieving efficient resource allocation.

To explore the effect of pollution on prices we let *one* resource, sugar, be a “dirty” good. That is, when agents harvest sugar from the landscape they leave behind production pollution. When they metabolize sugar they produce consumption pollution. Spice harvesting and consumption, by contrast, do not cause such pollution. Our experiment, then, is this. First, we will allow agents to trade. Then, after 100 periods, the agents begin generating sugar pollution and we track the effect on prices. At $t = 150$, pollution is turned off and diffusion processes are activated. Results of the experiment are logged in animation IV-3.

When the agents flee the polluted sugar mountains and move to the spice rich (sugar poor) regions, most of the sugar available to meet metabolic needs is what the fleeing agents have carried with them. Agents who need sugar must trade for it, and the relative sugar scarcity that

Animation IV-3. Price and Price Range Time Series for Pollution Accompanying Sugar Extraction and Consumption; Rule System $(\{G_1, D_1\}, \{M, T, P\})$



results causes the sugar price to rise.⁴⁹ The effect is dramatic, effectively an exponential price rise.

Then, when pollution generation is turned off—imagine this being the result of some technological windfall—and pollution levels are transported across the landscape by diffusion, the agents return to the sugar rich zones and the sugar price falls to its previous value of around 1.0. Such price adjustment dynamics are ignored in static microeconomics, where the implicit presumption is that, as a policy matter, it is safe to assume instant adjustment to a new equilibrium. For $t > 150$, we do indeed see adjustment back toward the original equilibrium. But from the perspective of agent society the process is far from instantaneous. In this case the artificial economy requires roughly twice as long to recover its statistical price equilibrium as it did to deviate from it. When transients such as this are long-lived, it makes little sense to focus all attention on equilibria. Artificial societies provide a means of studying price *dynamics*.

On the Evolution of Foresight

Against our simple agents it may be said that they are myopic temporally. A simple way to remedy this is to have them make decisions not on the basis of their current holdings but instead *as if* they were looking ahead ϕ periods. Formally, let the agents now move to maximize

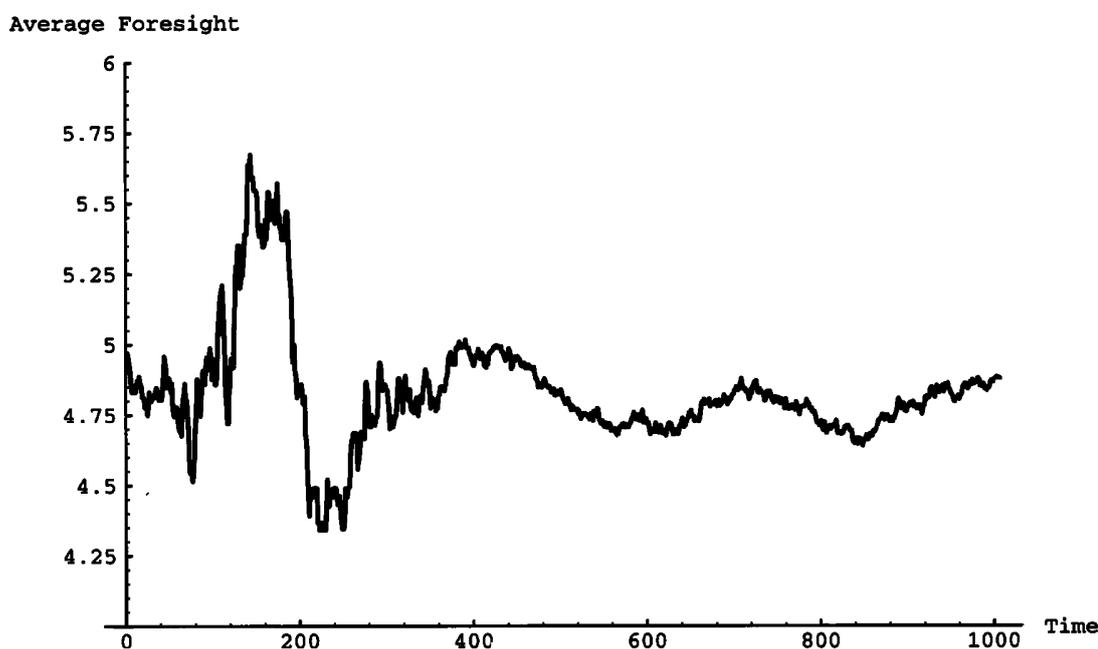
$$W(w_1, w_2; \phi) = (w_1 - \phi m_1)^{m_1/m_T} (w_2 - \phi m_2)^{m_2/m_T}, \quad (6)$$

where the parenthesized terms on the right hand side are set equal to zero if they evaluate to a negative number.

To study how this simple kind of foresight can modify agent behavior, we initially let ϕ be uniformly distributed in the agent population in the range $[0, 10]$, and then turn sex (that is, **S**) “on.” Once more, we can “watch” evolution unfold (see figure IV-18) by tracking the average foresight in the population.

Clearly, some foresight is better than none in this society since the long-run average foresight becomes approximately stable at a nonzero level. However, large amounts of foresight, which lead agents to take actions as if they had no accumulation, are less “fit” than modest amounts.

49. Recall that prices are ratios of spice-to-sugar: A sugar price of 5 means that a buyer of sugar would sacrifice 5 units of spice to acquire 1 unit of sugar.

Figure IV-18. Evolution of Mean Foresight under Rules $(\{G_1\}, \{M, S\})$ 

Emergent Economic Networks

Since agents in our model interact directly with each other rather than through the price system, there is a dynamic interaction structure that can be studied independent of explicit economic variables. That is, there exist well-defined networks that depict agent interactions and the evolution of such interactions.⁵⁰ Here we first describe the networks that have been implicit in the trade processes discussed above, networks of trade partners. Then we introduce a new relationship between agents, a credit rule, and study the network of lenders and borrowers that emerges.

Commodity Flows through Networks of Trade Partners

In this chapter we have specified rules for local trade between heterogeneous agents and have studied the markets that emerged. All trade was between neighboring agents. There thus exists a network of trade partners.⁵¹ To depict such a network, let each agent be a node of a graph and

50. Other models of trade networks include Kauffman [1988] and Tesfatsion [1995].

51. Since all trade partners are neighbors, but not conversely, the trade partner network is a subgraph of the neighborhood network, defined in Chapter II.

draw edges between agents who are trade partners. Such trade networks are endogenous in that they depend in a complicated way on agent behavior (that is, the movement rule, trade rule, and so on). They change over time, of course, as agents move around the landscape. Animation IV-4 gives such an evolution.

It is useful to think of the edges in such networks as channels over which commodities flow.⁵² Notice that although any particular agent trades with at most 4 neighbors, agents who are quite distant spatially may be part of the same graph, that is, connected economically. In essence, such graphs portray large-scale flows of goods across the landscape.⁵³

Credit Networks and the Emergence of Hierarchy

So far the agent societies studied in this book have been “flat”—there is no sense in which some agents are subordinate to others. This stems from the fact that agent interactions are usually short lived, lasting one (or at most a few) periods, or are symmetrical (such as when agents are neighbors of one another, or are mutual friends).

We can produce *hierarchical* relationships among agents by permitting them to borrow from and lend to one another for purposes of having children. The following local rule of credit produces such relationships:⁵⁴

Agent credit rule L_{dr} :

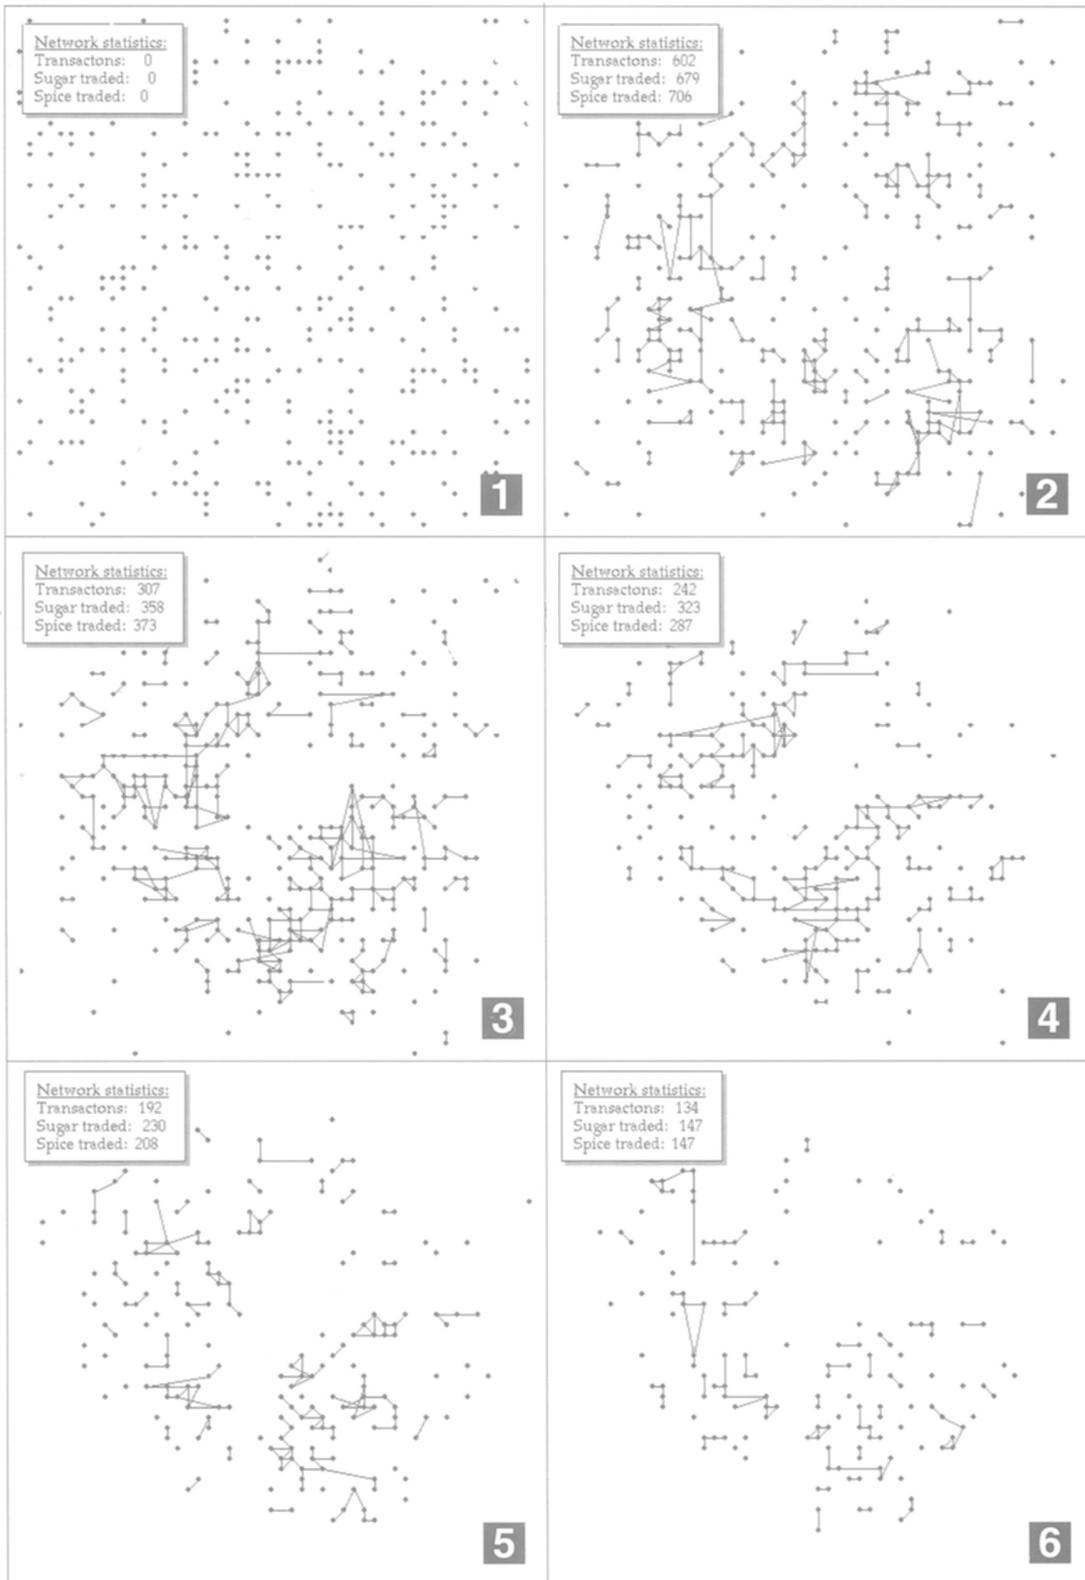
- An agent is a *potential lender* if it is too old to have children, in which case *the maximum amount it may lend is one-half of its current wealth*;
- An agent is a *potential lender* if it is of childbearing age and has wealth in excess of the amount necessary to have children, in which case *the maximum amount it may lend is the excess wealth*;
- An agent is a *potential borrower* if it is of childbearing age *and* has insufficient wealth to have a child *and* has income (resources gathered, minus metabolism, minus other loan obligations) in the present period making it credit-worthy for a loan written at terms specified by the lender;

52. As in Chapter III, lines across the entire lattice connect trade partners who are neighbors on the torus.

53. For why this network does not have a pure von Neumann structure, see footnote 29 in Chapter II.

54. Insofar as a primary consequence of this rule is that older agents lend to younger ones, a kind of finely grained overlapping generations model results.

Animation IV-4. Emergent Trade Network under Rule System ($\{G_1\}$, $\{M, T\}$)



- If a potential borrower and a potential lender are neighbors then a loan is originated with a duration of d years at the rate of r percent, and the face value of the loan amount is transferred from the lender to the borrower;
- At the time of the loan due date, if the borrower has sufficient wealth to repay the loan then a transfer from the borrower to the lender is made; else the borrower is required to pay back half of its wealth and a new loan is originated for the remaining sum;
- If the borrower on an active loan dies before the due date then the lender simply takes a loss;
- If the lender on an active loan dies before the due date then the borrower is not required to pay back the loan, unless inheritance rule I is active, in which case the lender's children now become the borrower's creditors.

This rule may not seem at first glance to be particularly parsimonious. However, it is the simplest one we could think of that bore some resemblance to real-world credit arrangements.

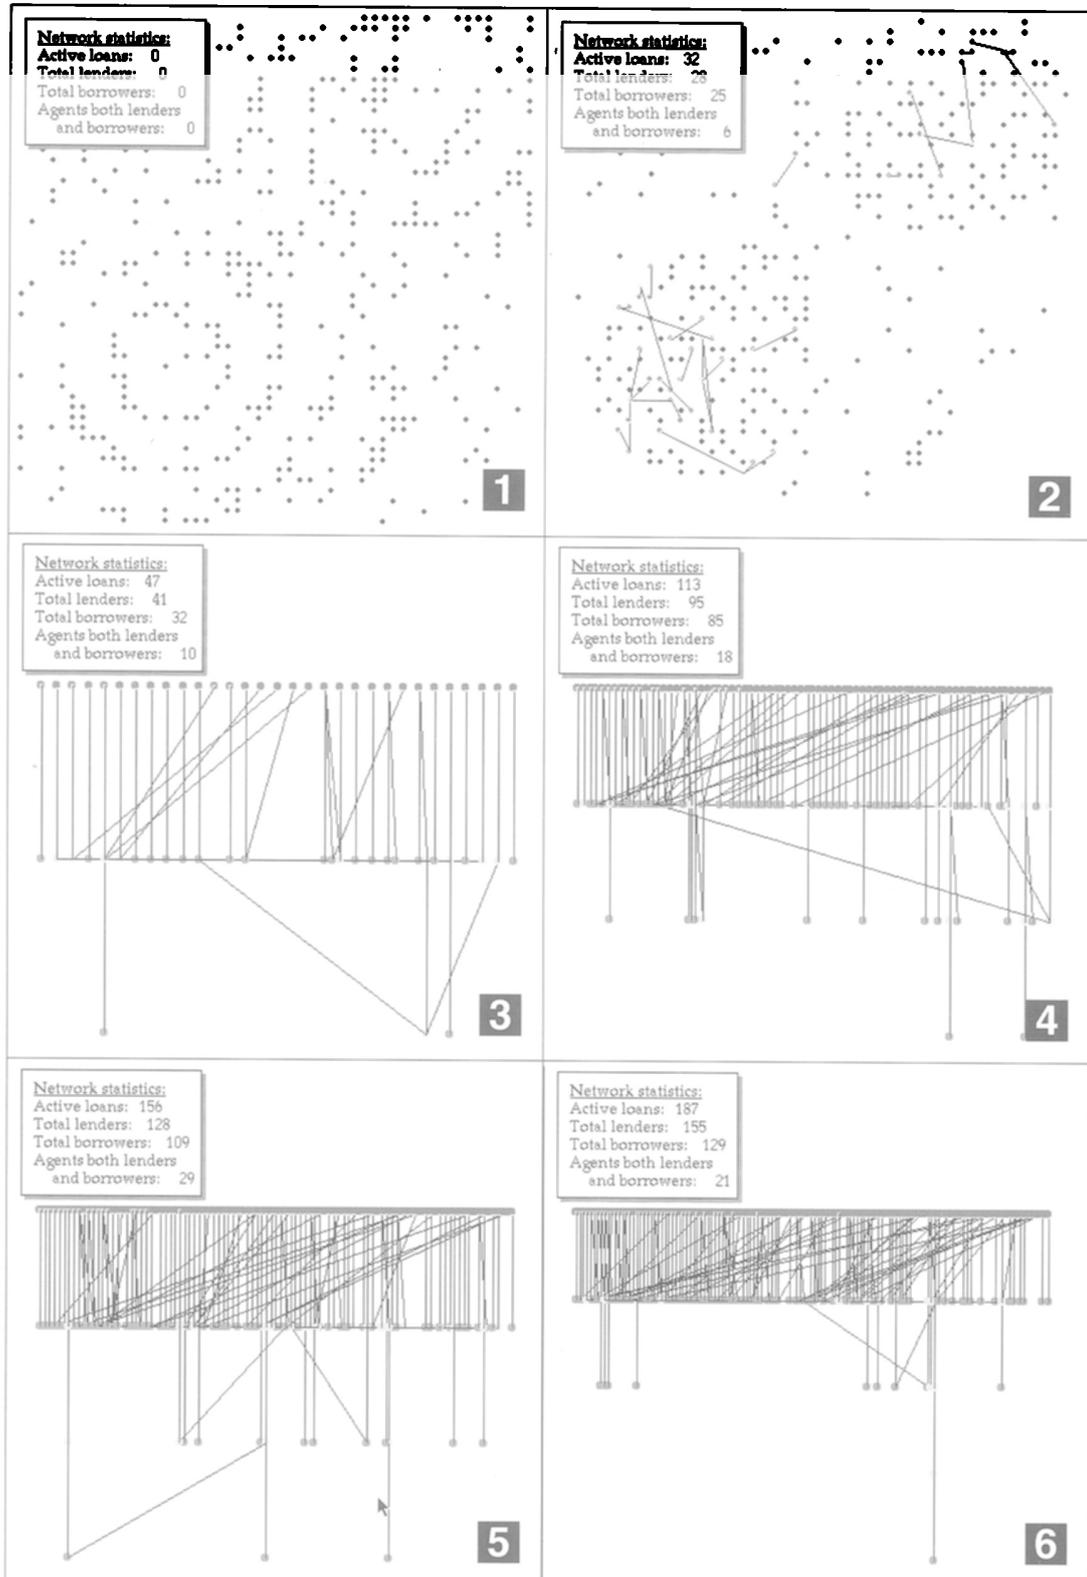
We return to the one commodity landscape to illustrate the operation of this rule. When agents move, engage in sexual activity, and borrow from and lend to one another, there result credit relationships like those shown in animation IV-5.⁵⁵ This animation begins by displaying agents spatially, coloring lenders green, borrowers red, and yellow those agents who are both borrowers and lenders. Subsequently, the hierarchical evolution is displayed. Agents at the top of the hierarchical plot are pure lenders, those at the bottom of any branch are pure borrowers, and agents in between are simultaneously borrowers and lenders. For this run as many as five levels of lenders-borrowers emerge.

Social Computation, Emergent Computation

The theory of general equilibrium is essentially a body of results on the *existence* of equilibrium. In the neoclassical story, the Walrasian auctioneer is a *mechanism* for achieving such an equilibrium. Once the market-clearing price is determined, the population of price-taking agents

55. The Sugarscape software system implements the n commodity generalization of credit rule L. For the sake of simplicity, the single commodity (sugar-only) form is used here.

Animation IV-5. Emergent Credit Network under Rule System $(\{G_1\}, \{M, S, L_{10,10}\})$



produces a socially optimal allocation of goods through exchange. The auctioneer is essentially an *algorithm* for the computation of prices. In this picture of the economic world no single agent has enough information (about endowments, preferences) to compute an efficient allocation on its own. Yet this allocation results through the cumulative actions of individuals. However, a particularly curious characteristic of this picture of decentralized decisionmaking via prices is that the auctioneer algorithm requires *centralized* information. That is, all agents must report their demands to the auctioneer who ultimately furnishes an authoritative price to the population that all agents *must* use in their trade decisions.

In reality, of course, there is no auctioneer, no *central price computation authority*. Rather, prices emerge from the interactions of agents. These interactions occur in parallel and asynchronously. It is *as if* agents are processing nodes in some large-scale parallel, asynchronous computer. Trade is the algorithm the nodes execute; nodes communicate prices to one another and change their internal states through the exchange of goods. The computation topology (architecture) is endogenous and ever-changing. Under what circumstances do such computations terminate in a market-clearing price? And how would any particular agent know that they had terminated—is it even possible to discern whether an equilibrium price has been achieved?⁵⁶ Does the fact that some nodes die while new nodes are regularly added to the social computer mean that notions of computation termination must be stochastic in nature?

Insofar as real economic agents engage in trade to improve their welfare, one might view the parallel, asynchronous exchange activities of agents not as the social computation of prices but as a distributed algorithm for the production of agent welfare. Artificial economies are laboratories where we can study the relative performance of distinct trade rules (algorithms) and alternative computational architectures (agent networks) in producing agent welfare.

Social computation concerns how societies of interacting agents solve problems that agents alone cannot solve, or even pose. Emergent computation concerns how networks of interacting computational nodes solve problems that nodes alone cannot solve.⁵⁷ Notice that these two fields have much in common.

56. Technically, these questions are very complicated. In the context of parallel, asynchronous computation, the relevant literature concerns the snapshot algorithm; see Bertsekas and Tsitsiklis [1989: 579–87].

57. For more on emergent computation, see Forrest [1991].

Summary and Conclusions

In many ways, the central question of economic theory is this: To what extent can economic markets efficiently allocate goods and services among agents? For example, does the ostensibly good performance of markets like stock exchanges tell us anything about the functioning of decentralized markets such as those for environmental goods and services? In our model we have just such a decentralized market—decentralized spatially—and we have found mixed results, to say the least, concerning the achievability of equilibrium prices and globally optimal allocations, under a wide variety of conditions.

Policy Implications

Foley [1994] has thoroughly criticized the way in which conclusions about economic policy are drawn from the model of Walrasian competitive equilibrium. In particular, the orthodox criticism of price regulation is that it is *irrelevant* if prices already fall within the limits set by the regulations or *distorting* if it actually constrains price movements. But in decentralized markets there is no single price.

If a significant amount of trading takes place at different price ratios, price floors and ceilings can serve to protect agents against relatively disadvantageous trades, and thus to mitigate the endogenous horizontal inequality produced by the market. [Foley 1994: 342]

Therefore, a clear role for economic regulation may exist when prices are heterogeneous.

Certain economists ascribe nearly magical powers to markets. Markets are idealized to operate frictionlessly, without central authority, costlessly allocating resources to their most efficient use. In this world of complete decentralization and Pareto efficiency, the only possible effect of government intervention is to “gum up” the perfect machinery. While this extreme view is perhaps little more than a caricature—and few would admit to holding it in toto—it is also, unfortunately, a position frequently promulgated in policy circles, especially when there is no econometric or other evidence upon which to base decisionmaking.⁵⁸

A different way to frame the issues raised in this chapter is as follows:

58. On the limited extent to which economic theory provides solid foundations for policy, see Hahn [1981] and Kirman [1989].

Do plausible departures from the axioms of general equilibrium theory produce markets that behave almost as well as ideal markets? While few would admit that extant markets function ideally, there is little cogent theory of performance degradation in real markets resulting from incomplete information, imperfect foresight, finite lives, evolving preferences, or external economies, for example.

The emphasis in the economics literature has been on the *existence* of static equilibrium, without any explicit microdynamics. Why cannot prices oscillate periodically on seasonal or diurnal time scales, or quasi-periodically when subject to shocks, or even chaotically?⁵⁹ Is it not reasonable to expect generational or other long-term structural shifts in the economy to produce prices that follow a trend as opposed to staying constant? Might not far from equilibrium behavior be a more reasonable description of a real economy?⁶⁰ From the computational evidence above, we think that there is good reason to be skeptical of the predominant focus on fixed-point equilibria. Economies of autonomous adaptive agents—and of humans—may be far from equilibrium systems. And, in turn, far from equilibrium economics might well turn out to be far richer than equilibrium economics.

59. Indeed, there is a growing *theoretical* literature that admits these possibilities; see, for example, Bala and Majumdar [1992]. However, these results do not seem to have made their way into policy discussions as of this writing.

60. As Farmer has observed, "To someone schooled in nonlinear dynamics, economic time series look very far from equilibrium, and the emphasis of economic theories on equilibria seems rather bizarre. In fact, the use of the word equilibrium in economics appears to be much closer to the notion of attractor as it is used in dynamics rather than any notion of equilibrium used in physics" [Anderson, Arrow, and Pines 1988: 101].