AGENTS 2002: "The Interplay of Differences"

Scott E. Page*

November 4, 2002

1 Opening Remarks

My title the interplay of difference refers to the substantial and various impacts of including or enabling diversity in physical, biological, human, institutional, and yes, since this is a conference on agent based modeling, artificial systems. Without diversity not only would there be almost no physics, no biology, no economics, no politics, or no sociology there would not be much use for agent based models. It is easy to construct and analyze a model of identical agents all in the same spatial and social locations interpreting the world identically and all taking the same action. Doing so is not very interesting or enlightening. The fun begins when we move agents outside of that oval office, and place them in social, economic, and family networks, endow them with distinct interpretations of the world and different intellectual toolboxes, assign them to roles and responsibilities within institutional and organizational structures, and enable them to define and evolve labels or types that simplify their interaction with others.

In this brief, wide ranging address, I share some (perhaps all) of my own thoughts on diversity and mix in and reinterpret some research by others. My remarks are meant as an invitation to others to explore, research, and play with models of diverse agents more than they are intended as a coherent and complete summary from on high. Agent based modeling research is itself bottom up.

This address is organized around four themes whose titles I borrow from literature: "The Classics", "Children's Literature", "Poetry", and "Science Fiction". Within each of these sections I'll emphasize a particular theme going off on riffs of varying coherency on topics ranging from architecture, to E. coli, to decks of playing cards.

*Center for the Study of Complex Systems, Departments of Political Science and Economics, The University of Michigan, Ann Arbor, MI 48106. I would like to thank Troy Tassier for reading earlier drafts of this work and Troy and Cosma Shalizi for their input on this project. This research is supported by a grant from the James S. McDonnell Foundation.

2 Diversity in the Classics

To begin my examination of the role of diversity in agent based models, I'm going to revisit some of the classic models in complex systems, or at least the ones that I teach. These include Conway's game of life, Wolfram's rules 30 and 110 for two dimensional CA, Schelling's segregation model, Watt's small world framework, my and Lu Hong's diversity model, Arthur's Bar Problem, and a novel little game called the prisoners' dilemma that like many others I first loved, then grew to find tedious but which I now recognize is incredibly important.

2.1 The Game of Life

Most members of the agent based modeling community have played with Conway's game of life and experienced how the simple rules enable elaborate patterns, from cycles (including my favorite the figure eight shown below), to gliders, to glider guns.[1]

The Game of Life Figure 8

			X	X	X
			X	X	X
			X	X	X
X	X	X			
X	X	X			
X	X	X			

This configuration which resembles an eight at a 45 degree angle creates a cycle of length 8 according the the Game of Life rules for a two dimensional cellular automata.

The Game of Life contains limited diversity on just a few dimensions. There is only one type of agent. The agents can choose from among only two actions. All agents follow the same behavioral rule. They do not tremble or mutate. And the spatial diversity is over a two dimensional grid with a checkerboard topology transformed into a torus. Yet, in spite of (and I will argue because of) this limited diversity in the model's assumptions, the Game of Life exhibits incredibly diversity at the level of emergent structures. In fact, the Game of Life has been shown to be a universal computer. So in a sense, you can do anything from these simple model.

[1]In the Game of Life, each cell on a checkerboard has eight neighbors. A cell can be alive or dead. A live cell remains alive if two or three of its eight neighbors are alive. A dead cell only comes to life if exactly three of its neighbors are alive. Time proceeds in discrete steps with synchronous updating.

The fact that a little diversity can go a long way will recur in our discussion of the classics.

2.2 A New Kind of Science: CA Rules 30 and 110

As captivating and evocative as the Game of Life may be as evidence of the potential for a little diversity to go a long way it is at best a definitive example, at worst an anomaly, and probably lies somewhere in between in the land of supportive anecdotes. A more compelling argument runs through the many pages of Steven Wolfram's *A New Kind of Science.* Here we see in gory detail the ability of so little to do so much (and, based on the book sales, for so many!). Much of Wolfram's analysis focuses on one dimensional cellular automata using a nearest neighbor topology. It has an even less diverse spatial structure than the game of life and certainly one far less interesting than what we encounter in the real world irrespective of scale. You could look from quarks, to genes, to molecules, to neurons, to people, to computers, to firms, to nations, to planets and you'd be hard pressed to find many interaction topologies as simple as the one dimensional CAs that Wolfram, for the most part, studies.

On top of this, Wolfram allows only two types of actions for most of the book. And, as in the Game of Life, all of the agents follow the same behavioral rule and interpret the world identically. They do not adapt, and they do not tremble. In fact, Wolfram doesn't appear to think of the cells in his CAs as agents but as automata, which is what they are. Thus, many of Wolfram's models like the Game of Life should be seen as belonging to a class of models that assume fixed spatial diversity and moderate action diversity. In the space of the possible, this class occupies a small but important location.

Judging from the book's many reviews as well as a handful of conversations in the hallway, what Wolfram shows can be interpreted anywhere from revolutionary to provocative to repetitive. My purpose here is not to pass judgment on all of those keystrokes, but to highlight what readers have obviously gleaned: you can do some incredibly cool stuff with two state one dimensional cellular automata. You can create randomness - what could be more diverse than that – and you can compute anything provided that you start in the right place.

As you scan through all of the pretty pictures in Wolfram's book (lacking permission none are reprinted here), you might think that these CA rules can produce anything. There is a sense in which that is true temporally: if rule 30 truly does generate a random sequence, then eventually it will produce any finite sequence of 0's and 1's along a single cell location over time.[2] There is also a sense spatially in which it is not true. A long time ago Arthur Burks and others described what are known as Garden of Eden states. These are states of the automata that cannot be attained temporally. By this I mean, if I look at all of the cells in a particular time step - a

[2] In other words, if we look at a vertical line of cells, we might see anything.

horizontal strip – there will be some configurations that I never see for a given rule.

To reiterate this distinction because it is an important one: rule 30 can give you any vertical slice you'd like but only a subset of the possible horizontal slices. Those that it cannot belong to the Garden of Eden. These Garden of Eden states must be the initial states of the system if they are to be visited at all, so those visits will either be brief or, in the case of equilibrium Garden of Eden States, eternal. Therefore, at least in the case of two dimensional CAs, to say that a little diversity goes a long way is eerily accurate, because while it goes quite far, it does not go all of the way of being able to generate all states.

The two classic models discussed so far include limited diversity of actions with a dash of fixed spatial diversity tossed in so that we can look at pretty pictures. We have ignored diversity in behavioral rules and representations, which could either be hard wired or allowed to emerge, as well as diversity in spatial representation and social connectivity. We turn to these types of diversity next.

2.3 Schelling's Segregation Model

In later introductions to *David Copperfield* Charles Dickens wrote that in his heart of hearts he had a favorite child and that child's name was David Copperfield. If there is a competitor to the Game of Life as the best of the simple, as the model in moments of ambition we most would like to call our own, it must be Schelling's model of tipping. I'm not violating any confidences in saying that scholars ranging from Cederman to Axtell to Gaylord adore this model for its elegance and intuitions. In the model, there are two types of agents, which I will call type 1 and type 2 agents. The agents are placed randomly on a two dimensional grid. These agents can move if they are unhappy with their current location. In some instantiations of the model these agents follow a single rule: move if more than X% of your neighbors are of the opposite type as you.[3] Schelling finds that the system tips. A few agents move, then more until eventually it segregates into regions of type 1 and type 2 agents even when the tolerance threshold is set at 40% so that an agent only relocates if more than 40% of her neighbors differ from her.[4]

The clarity and power of Schelling's message resonated across the social sciences: we cannot infer agent level racism despite the segregation we see in American cities as detailed in Denton and Massey's *American Apartheid* or in high school or college cafeterias. Moreover, we must always be careful not to commit the error of inferring focal or obvious microlevel forces from macro level patterns, the so called ecological fallacy. That the transition from micro to macro can be paradoxical with

[3] In Schelling's paper he allows the two populations to have different thresholds for moving. And, I can happily report from experiments of my own that this diversity of thresholds creates significantly more tipping.

[4] I should add here that the "tipping" is not as pronounced as one might expect from Schelling's account of the model. But in fairness to him, his technology didn't allow him to calculate the extent of tipping as easily as we can now.

so few moving parts is astounding. This is the real beauty of Schelling's model.

I want to mention one other lesson we can learn from the Schelling model, namely that identifiable differences can be used as food for both thought and action by other agents. When we let agents loose to evolve strategies and actions they will often use anything at their disposal to gain an advantage. Types are often something at their disposal. This phenomena can be seen in other agent based models as well. See especially the paper in *Nature* by Cohen, Axelrod, and Riolo whose agents use identifiable tags to create epochs of cooperation in the prisoner's dilemma.

2.4 Watts' Small Worlds Framework

In the small world framework of Duncan Watts and others, agents belong to closely knit groups (the small worlds) and they also have an occasional random connection – that guy your sister's roommate dated who moved to Seattle that you ran into at the airport one time. One way to appreciate the beauty of the diversity in this model is to think of it as the offspring of two other models and to compare it to its rather boring parents. One of the parents has all random connections. The other has all local connections. If we consider a collection of agents with random connections, we can prove all sorts of things about degrees of separation and spanning trees, but none of these are very surprising or interesting. Alternatively, if we consider a collection of agents with local connections, whether those be in cliques or on a lattice, we again find little beauty in the answers to questions like: are all agents connected? Yes. What is the separation between two agents? It depends on how far apart they are on the checkerboard. Etc ...

By combining the random connections with the local connections, the small world model generates degrees of separation between agents that are remarkably close to what arise in the random model, and this is accomplished with relatively few random connections. The small worlds model has many other nice properties of which most of us our familiar. All I wanted to emphasize was how these properties rely on just a little bit of diversity.

2.5 Brian Arthur's Bar Problem

The Bar Problem, like many of the models that we now explore, can be found in Schelling's *Micromotives and Macrobehavior*. In Brian Arthur's implementation of the Bar Problem there are a collection of agents each of whom must decide whether or not to go to the bar. All of the agents have identical payoff structures. They want exactly sixty agents to be at the bar. The agents evolve rules to decide whether or not to go to the bar based upon the time series of past attendances. These rules are simple. Arthur restricts the set of possible functional forms that they can take. Yet, the agents collectively can evolve strategies which lead to approximately the correct number of agents going to the bar each time. If all agents use the same rule, then either all or none of the agents go each week. Instead, they evolve a diversity of rules that works. When several agents' rules tell them not to go, other agents' rules kick in and tell them to go. It's a form of stability through diversity, a point we will revisit later in this talk.

If we compare the Bar Problem to the Santa Fe Artificial Stock Market model, which by comparison is baroque, unwieldy, and difficult to analyze, we begin to see the reason for limited diversity. Too much diversity on too many dimension and we get a mangle that is hard to sort out. But a little diversity appropriately placed can generate just enough complexity to prick our interest.

2.6 Hong and Page's Diversity Model

In my model with Lu Hong on diverse agents solving problems, we endow agents with diverse encodings of problems - what we call *perspectives* - and diverse analytical tools to solve those problems - what we call *heuristics*. We can think of heuristics as playing cards and each agent having up to thirteen cards. Simple combinatoric arguments show that the number of collections that an agent might have is huge, as many as the number of possible bridge hands. So here finally, we have a model that does not have a little diversity. It has diversity in abundance.

We find that diverse collections of agents perform incredibly well. In fact, under a set of fairly unrestrictive assumptions, we find (and prove) that a relatively small group of randomly chosen diverse but intelligent agents will outperform a group of the best agents on a hard problem. So, for example, in a model with five possible heuristics, we endowed each agent with five. The ten best individual agents performed worse than ten randomly chosen agents. The reason for this is that the best agents will attempt to solve the problem similarly so having two of their heads is not much better than having one. Also, all of the agents are smart, so the random agents have good heuristics as well.

The first lesson to take from my model with Lu is that small collections of simple diverse agents can be collectively extremely bright. Even though the model can include substantial diversity and still work, all of the diversity is not needed to make the point that diversity trumps ability. It's an extension of an idea that's been around since Adam Smith and expanded upon by Hayek. Recall the old lines about how markets aggregate individual knowledge. The railroad worker knows about railroads; the farmer knows about wheat; and therefore the economy knows about everything. But while the the Smith-Hayek argument about markets as super agents is really a comment on the power parallel processing along dimensions, our result is on the power of simultaneous diverse processing that overlaps dimensions.

The second lesson to take from the model is that if you want to include lots of diversity, which we do by allowing every agent to have his or her own heuristics and perspective, as well as all of the agents interacting in a soup, then do not have the agents influence one another. In the Hong Page model, the agents work on a solution to a problem. So all of the agents interact but with the problem solution not with one another.

2.7 The Prisoner's Dilemma

The prisoners' dilemma is by far the most famous, popular, analyzed, and modeled game. There have been more agent based models of the prisoners' dilemma than I would want to count. One thing that has been shown in the agent based PD models is that agents can learn to evolve cooperative strategies like Tit for Tat. However, in my experience they can only learn to do this if one of two conditions is satisfied. Either the space of strategies must be small: *limited strategy diversity* or the agents must be connected by a topology and not in a soup: *limited number of interacting players.* If you begin with a soup of many players who can evolve all sorts of strategies, you will wait a long time for cooperation to break out. In fact, within the soup cooperation won't have a place at which to gain a foothold so the strategy will have to be both sophisticated and general. I am not saying that you cannot get cooperation, you can. It will just take a while.

My point here is that the early models of agents by John Miller, Robert Marks, Bob Axelrod, Nowak and May, and others that generated cooperation in the PD did so by only admitting limited diversity – agents did not have big strategy spaces and they were also often small in number or placed on a lattice.

If a little diversity goes a long way, what can we say about a lot of diversity or even a modicum of diversity? Do they take us even further? Surprisingly, the answer is often no. Agent based modelers from Axelrod with his KISS principle (Keep It Simple Stupid) to Wolfram with his empirical evidence that three state and five neighbor CA rules are not as interesting on average as two state three neighbor CA rules extol the virtues of stark models. Even the Sugarscape model of Epstein and Axtell, which by comparison appears Guadian with its sugar and spice, epochs of cooperation (everything nice?), trade, and disease, is relatively simple at it's core.

In the end, we're left with a bit of a paradox. Our most compelling agent based models of complex adaptive systems do not have many moving parts, what many would consider a defining characteristic of something that is complex. Diversity is necessary for complexity in that complexity models require diversity to get them started. This initial diversity can be in types, in behavior, in actions, or in space. And, in the most compelling models we see that only a little diversity has been added along one or two of these dimensions.

3 Children's Literature

My second theme concerns the diversity or variation that arises from mutation and the process by which those mutation rates are chosen or evolved. I rely on a single work of children's literature to frame my discussion: *Goldilocks And The Three Bears.* To remind you of the story, Goldilocks happens upon the bears' cabin and finds three chairs, three bowls of porridge, and three beds. When Goldilocks tastes the porridge she discovers that one bowl is too hot, one is too cold, but one, the baby bear's bowl, is "just right." The metaphor of finding the just right temperature applies in agent based modeling in multiple contexts. When we hear someone talk about the exploration/exploitation trade off, the edge of chaos, or the evolution of evolvability, they are really talking about finding the baby bear's porridge, though being good academics they use a whole lot more jargon.

In most biological, ecological, and social science models, mutation is the primary cause of diversity. Crossover in biological contexts and partial imitation in human ones often bootstraps this diversity to create even more. But it is important to note that crossover alone may not guarantee full exploration of the space depending upon the initial population. All of us have learned the same lessons when it comes to mutation rates: Too much mutation leads to a system that boils and bubbles and never settles. Too little mutation leads to a system that freezes. Somewhere in the "in between" the system comes alive, generating patterns and complexity.[5] In my brief comments, I am going to focus on the discovery of and properties of that in between region.

3.1 Exploration and Exploitation

I'll begin by showing the trade off between exploration and exploitation. Scholars of agent based modeling and complex adaptive system have devoted substantial time and energy developing models that explicate the trade off between exploration and exploitation. Too much exploration and the system eternally boils. Too much exploitation and the system stabilizes prematurely. If our model was built to compute something or to find an efficient outcome, in the former case the system never stabilizes and therefore fails to exploit successful structures. In the latter case, the system settles on a bad solution, possibly the best in the initial population. In contrast, if our model was intended to represent some physical, biological, or social phenomenon, too much exploration and we just have random agents bumping into one another, but if we have too much exploitation, the system is not interesting. Therefore, if we imagine an exploration/exploitation slide bar, we want to position it in what I like to call "the interesting in between" the region separating stasis and randomness.

To see the exploitation/exploration trade off in action, consider the following problem. Each of two agents must choose whether to graze on some subset of thirty fields. An agent gets a payoff of one from being the only grazer on a particular field, a payoff of negative two from being one of two grazers on a field, and a payoff of zero if the agent does not graze on the field. In this multiple location commons problem, there are 2^{30} socially efficient equilibria one for each partition of the fields among the two agents. Below are data on average payoff in period one hundred from a series of runs in which the mutation rate varies from 0.0001 to 0.20 from a model using a genetic algorithm with a population of size thirty and a simple tournament selection operator. Performance equals the percentage of the maximal possible payoff achieved.

[5] Oscar Wilde once quipped that "Moderation is a fatal thing. Nothing succeeds like excess." We will see how the opposite is true in agent based models. Excess is a fatal thing. Nothing succeeds like moderation.

Mutation Rate	Performance
0.0000	23.33
0.0005	66.67
0.0010	83.33
0.0015	100.00
...	...
0.0200	100.00
0.0205	100.00
0.0210	100.00
0.0215	96.67
0.0220	96.67

This simple example shows how with too much exploitation, we cannot find the optimum and with too much exploration we cannot exploit successful structures. This example also shows an enormous range for acceptable mutation rates. So while there may be an optimal mutation rate, namely the one that on average would most quickly enable the agents to locate and maintain the optimum, we need not expend much energy to find it. In this case, the slide bar has a huge margin of error. This often occurs in agent based models. Mutation rates can vary substantially without much effect on outcomes.

3.2 The Edge of Chaos

The slide bar metaphor also applies in a loose way to the concept of the "edge of chaos" put forth by Langton and others on the basis of some simple models. The edge of chaos as initially formulated was intended to capture the idea that a small increase in action interdependency can shift a system from being "complex" to being chaotic. Life and even the Game of Life, were thought to exist on this precipice at the edge of chaos. As Crutchfield first explained and John Miller and I have elaborated in a simpler context, the edge of chaos may not exist, at least in the case of one dimensional CAs. Instead as Miller and I describe, CA rules may be classified as structured (all 0's and all 1's are stable), unstructured (a CA of all 0's gets mapped to all 1's and vice versa), and partially structured (both all 0's and all 1's go to the same state). Chaotic rules belong to the unstructured class and complex rules belong to the structured class. In other words, they are not near each other in CA rule space. Therefore, there is no edge there. The edge appeared to exist because the CA rules were projected onto a one dimensional measure, the number of 1's in the CA rule, making the complex and random CAs neighbors on that projection even though they are not close neighbors in the space of CA rules.

All of the theory aside, there is a looser sense in which there is an edge of chaos. If we plot the efficacy of a system or of a search algorithm as a function of the mutation rate, such as I did in the table above, we sometimes see gradual increases up to the optimal mutation level and then a sharper drop off in performance. This extreme performance decline from a little push of our exploration/exploitation slide bar toward more exploration suggests “an edge”. The “of chaos” part we'll just have to chalk up to marketing. Systems that haphazardly move between states need not be chaotic in the formal sense.

Notice though that there did not appear to be much of an edge in example I gave above, but watch what happens if we increase the number of agent types to ten but keep the environment essentially the same. This will make the environment more complex and therefore create more of an edge in payoffs from an increase in mutation. With ten types assume that an agent gets a payoff of nine if the agent is the sole grazer but loses ten for each other agent that grazes. This imposes a much stiffer penalty for stepping on what has evolved to be another agent's turf. This change in payoffs alone will cause the drop off to appear more severe than in the previous example, but as we see in the table below, even taking the payoff difference into account the reduction is larger.

The Exploration /Exploitation Trade off Revisited

Mutation Rate	**Performance**
0.0175	100.00
0.0180	100.00
0.0185	90.37
0.0190	100.00
0.0195	75.93

Considering both the two agent and the ten agent versions of our simple multiple field grazing model, it seems fair that the mutation slide bar can create an edge but need not.

Kauffman offers an alternative formulation of the “edge of chaos” and life existing at its edge based upon his NK model. He finds that as you increase K, the number of connections between agents, you get an increase and then a decrease in performance. Later in this essay, I will discuss the similarities between increasing mutation rates and increasing the number of types of agent interactions. For now, I merely call your attention to the fact that Kauffman's edge is similar in spirit but different in construction from Langton's and that both suffer from some logical flaws.

3.3 The Evolution of Evolvability

The discussion so far suggests that the best mutation rates balance exploration and exploitation to avoid falling off of the edge into the boiling porridge region. The adage "moderation in all things" appears to apply even to diversity. But we know more than that. The payoff as a function of rate of mutation is single peaked in the mutation rate. We also know that single peaked functions can be solved easily by either gradient ascent methods or evolutionary search. If we combine these two ideas, we bump into an irresistible concept: "the evolution of evolvability."

I first ran across this phrase in a paper by Bedau and Packard. In a fascinating (but for our purposes perhaps overly complex) model, they allow the mutation rates to evolve. The ideas of evolving search parameters was not original to them I'm sure, but the recognition that by evolving the mutation rate into the baby bear region, the system is evolving not only good system level performance but also the ability to evolve itself is a profound insight. And in light of the fact that performance appears single peaked in mutation rate with a substantial central plateau even if the system's ability to evolve the mutation rate was crude, evolution would succeed. It would lead to evolvability. Even better, if the system exhibits an "edge of chaos" and if it also begins with a mutation rate that is too high, then any change that leads to lower mutation rates would be of high marginal value and we would expect it to be even easier to evolve evolvability.

Can we infer from these few models that scholars have uncovered a partial explanation for the existence of successful evolutionary systems? I think so. At a minimum, we have not found evidence to the contrary, that getting the mutation rate in the evolvability region would be hard. Biological constraints on how mutation rates might change notwithstanding, getting the mutation rate correct appears a whole bunch easier than say evolving the Krebs cycle, the human eye, or a pocket watch.

To quickly summarize our second theme, we see that diversity drives system performance. Without it the system has no ideas or attributes to explore. But with too much diversity the system boils. Agents cannot maintain and exploit successful building blocks. In some cases, it appears that ramping up the diversity in the population leads to gradual increases in performance followed by a precipitous falls. This collapse can be though of an edge, as a cliff on a rugged landscape. Or, if you really want to push the metaphor, as "the edge of chaos" where chaos is loosely interpreted as a system that never settles down or exploits successes. Finally, in that payoffs appear single peaked in the mutation rate, the evolution of mutation rates that enable evolution to occur, the so called evolution of evolvability, appears a plausible conjecture.

4 Poetry

I want to begin by reminding you of two of the most over quoted lines of poetry in the English language. I will then relate them back to two themes from complex systems, namely path dependence and lever points, and then I'll argue with a tinge of sadness why these poems resonate so and what agent based modeling and complex systems has to say about that. The first line is Robert Frost's "two roads diverged in a yellow wood ...". It has obvious uses in advertising colleges and the like. The idea being that if you take the right road now, you'll be happier. The other line is from none other than Shakespeare and it puts a negative spin on the same theme "There is a tide in the affairs of men which taken at the flood leads on to fortune. Omitted all the voyages of his days are bound in shallows and in miseries." There's a happy thought. Both poets are talking about how decisions in our lives effect our futures, that history matters. Frost is talking about the path. To him the entire path matters. Shakespeare is talking about a lever point, a critical juncture.

These lines profoundly affect people because they highlight and romanticize those rare instances, those times when we make choices that influence our lives. Most days aren't like that. The courses of our lives are much more robust. Lots of butterflies flap their wings without changing the next word I type. It is this unromantic robustness that I want to focus on. Is it related to diversity in any way?

From an ecological perspective there would seem to be one obvious positive relationship between diversity and robustness. And by diversity here I mean diversity of species and not so much diversity within a species. If we reduce the number of species in a particular ecosystem, we might destabilize it, and the ecosystem may not be able to stabilize itself without the loss of many other species. The impact will be especially severe if we wiped out what is known as a keystone species. But this line of argument is ridiculous because if we introduced a new species we might also destabilize the system. Therefore, more diversity would also appear to make the system less robust.

Instead, we would like to ask, are more diverse ecosystems better able to respond to fluctuations in their environment or the extinctions of some species? This question is difficult to answer. To see why, I will describe two extremely similar models. In the first, diversity will imply substantially more robustness. In the second, the opposite will be true.

4.1 The Route Selection Model

Imagine a world in which there are twenty time periods per cycle and twenty possible actions. Suppose that each agent must take each action once each cycle to survive and that each action takes exactly one period. We can then write the behavior of an agent during a cycle as an element of the permutation group on twenty objects. To complete the model, suppose that the payoff from an action during a particular is decreasing in the number of others taking that action in that period. Think of agents who have to visit a set of stores. They would prefer to go to less crowded stores. Therefore, everyone wants to be doing the opposite of what everyone else is doing.

Given some number of agents, there can be an enormous number of efficient collections of routes. Some of these are rather simple. For example, with five actions, denoted by A,B,C, D, and E, if equal numbers of the agents choose the sequences: ABCDE, BCDEA, CDEBA, DEABC, and EABCD, then the collection of sequences will be efficient. There are also much more diverse efficient collections. Consider the collection in which there is exactly one agent choosing each member of the permutation group on five elements. Since an equal number of these sequences will have each action in each location, the collection is efficient.

Now let's compare the robustness of these two collections or sequences. Suppose that one action and one time period are eradicated. For simplicity, assume that it was action E. Now the five sequences in the population will be ABCD, BCDA, CDAB, DABC, and ABCD. The first and fifth sequences are identical. Forty percent of the agents could benefit by changing their sequences. Given the simplicity of the model, the agents will be able to stabilize into a new configuration but it will take some time.

In contrast, in the case where the agents each choose a unique strategy from the permutation group on five elements, losing one action and one time period will have much less of an effect. In fact, the collection of sequences will remain efficient because the permutation group on five elements will be transformed into five copies of the permutation group on four elements once one action has been dropped. In this example, we see a strong connection between diversity and robustness.

In the second example, each agent must choose some subset of N fields on which to graze his sheep as in an earlier example. But what I want to do is change the payoffs so that each agent wants to graze with each other agent on exactly k fields. If there are only two types of and a field is dropped, there are only two possibilities: the agents met there or they did not. If they did not meet there, then their actions remain efficient. If they did meet on that field, then all that the agents need to do is fine one new field on which to meet, a task that should be accomplished with a few timely mutations.

If we increase the diversity so that there are now ten agents that each must meet each other type on exactly k fields, we see more interesting phenomena. First suppose that the ten agents have settled on a pattern of grazing which is not diverse: *all ten agents choose the same k fields.* The dynamics that will arise if we wipe out one field are not identical to what we would see with only two agents. If they did not graze on the field that is wiped out, then they are still optimally arranged. If they did graze on the field, then they have to coordinate on a single new field or evolve some pattern in which each meets another once. These processes would take some time. Second suppose that the agents have diverse grazing patterns. Reducing the number of fields is certain to wipe out several matches between pairs of agents. agents then need to find new fields upon which to match one another, but coordinating on a new set of fields will not be easy unless there is an excess of fields. Assuming that there is not an excess of fields there are two possibilities one being that they will not be able to find a field. In this case, efficiency falls and the system is less robust. In the second, they do find another field on which to meet but in doing so they will then be meeting with other agents more than they would like. This may cause further relocations. The two possibilities are lower performance or a rather elaborate sequence of action changes. Therefore, we cannot help but think of this system as less robust than the other.

These two examples suggest that the relationship between diversity and robustness is too subtle to just say that increasing diversity increases robustness or that increasing diversity decreases robustness. We've just seen two similar models and seen how in one case diversity increases robustness in the other case it decreases. To gain a better understanding of what is going on we'll need to turn to science fiction.

5 Science Fiction

My third theme is the most provocative, the least advanced, and the most speculative. It is science fiction for the simple reason that I haven't advanced my argument to the point where it could be called scientific fact. This section should be seen as an attempt to makes some sense of the first three themes. What I will argue is that the number of types, the number of interactions, the mutation rate, and the spatial arrangement of agents all contribute to the viability of a complex adaptive system similarly and that the complexity models that we most appreciate and understand all satisfy what I call the *limited interplay.* I want to distinguish the word interplay from the word in that you and I could interact but you might have not affect on me, as is the case in the Hong Page diversity model which supports all sorts of diversity. If we interplay then what you do affects my fitness or payoff.

I will begin by developing a connection between rates of mutation and the number of agents with whom an agent interacts. The connection is straightforward: if an agent interacts with ten other agents, then if any one of those agents mutates its action, then the agent under consideration may also have to change his action. Thus, the effective mutation rate could be ten times as high as when the agents interacts with only one other. I'll give an example where this is exactly the case. Second, I will show how spatial arrangements in effect reduce the number of types, thereby allowing models to stay complex for two reasons. The first reason is that the systems might boil. The second reason is that spatial segregation enables local clusters of agents of the same type to survive. This will lead to my discussion of the subtle and confusing relationship between diversity and robustness.

5.1 Mutation and Type Equivalence

The first functional form is what I will call the *single grazer commons problem* (SGCP). In the SGCP, there are In the SGCP there are N fields. Each agent can choose some subset of those fields to graze. In this way, the actions of the agents can be written as binary strings of length N. The payoff to being the only agent on a particular field equals one, and the cost of having two agents on the field equals some number larger than one. Given these payoffs, an optimal configuration of agents will have exactly one agent grazing on each field.

It is straightforward to show that for any N and for any M, a configuration is an equilibrium in action space if and only if it is the optimal configuration. If a configuration is not optimal then at least one of two conditions must be met: (i) two agents are grazing on the same field or (ii) there exists a field upon which no agents are grazing. If (i) holds than one of the agent types that is grazing on the same field as another can increase its utility by no longer grazing on that field. If (ii) holds, then any agent type can increase its utility by grazing on that field.

This first model nearly satisfies, what I will call the *reducible other agents property*. This means that the strategies of all of the other agents can be aggregated to form a strategy that could be that of a single agent. If we change the payoff structure so that the payoff equals 0 if you do not graze, 1 if you graze and no one else grazes, and minus the number of agents if you graze and some other agent or agents also graze, then the reducible other agents property would be strictly satisfied. If the reducible other agents property is satisfied then increasing the number of agents is equivalent to an linear increase in the mutation rate. Suppose that the system has stabilized. At this point, the interplay in SGCP for a single agent is approximately the same whether there is one agent with a mutation rate of 0.01 or ten agents with a mutation rate of 0.001. Out of equilibrium, there is a similar though less precise correspondence. If system performance is good then an agent will not have much interplay. There will not be many other actions by agent types that influence the agent. Therefore, in this model we see the that it is possible to create an equivalence between the number of types and the mutation rate. It is also clear that increasing the mutation rate increases interplay. If an agent has settled down on a particular strategy or action it is as if the other agents do not interplay with that agent. But if the agent keeps experimenting then the other agents do interplay with the agent.

5.2 Interplay Not Number of Types Determines Complexity

The second point that I want to make is that by virtue of the payoff structure SGCP has low interplay regardless of the number of types. If the interplay hypothesis is correct, the number of types should not matter much for time of convergence. The table below that shows time to an efficient equilibrium as a function of the number of types supports that intuition. In this and other results from this section I use a genetic algorithm with a population of size thirty to represent each agent and employ a standard tournament selection operator.

Time to Efficiency in the SGCP

Agent Types	Number of Periods
2	23.52
3	30.88
4	32.98
5	40.18
6	45.64
7	47.10
8	53.46
9	56.48
10	62.64

These data show that the time increases only slightly and that is due to the fact that the initially zeros and one are equally likely are in my initial populations of strings. When there are ten types the strings in an efficient configuration have many fewer ones than that.

The second model is the same one that I discussed in the previous section where each agent wants to meet each other agent on exactly k out of N. In this game, there are N choose k equilibria in which all of the agents choose the same k fields. There are also more diverse equilibria as well as inefficient equilibria. Suppose that $N = 4$ and $G = 2$. Consider the following configuration

Type 1: 1110
Type 2: 1101
Type 3: 1011
Type 4: 0111
Type 5: 0111

It is straightforward to show that the Type 1, 2, and 3 agents match each of the other types on exactly two fields. Agents 4 and 5 match each other on three fields. Since Types 4 and 5 have the same action, it suffices to show that Type 4 cannot improve by changing its action. A simple calculation shows that if the Type 4 agent chooses not to graze in any field, to graze in only one field or to graze in two fields, the agent is worse off. Similarly, if the Type 4 agent grazes in all four fields, the agent is also worse off. Therefore, the Type 4 agent must graze in exactly three fields. It turns out that any action that grazes in exactly three fields gives the same payoff. Therefore, once the population settles into the configuration such that each of the five agents grazes on exactly three fields, it will never leave that configuration.

It should be clear from this example that this is a model with a lot more interplay than the previous model and that moreover the interplay increases in the number of agent types.

Time to Efficiency

Agent Types	Number of Periods
2	2.08
3	3.60
4	6.16
5	10.02
6	18.46
7	57.90
8	272.32
9	867.96
10	1650.52

What we see in these data are that time to efficiency increases dramatically in the number of types but that is because the interplay increases as the number of types increases.

5.3 Space and Interplay

Many of our models of complex systems that generate cool patterns take place over networks and not in soups. I want to first describe some real world experiments that corroborate agent based theory and then I will tie the model back to the concept of interplay.

In a recent paper in *Nature*[6], Kerr, Riley, Feldman and Bohannan conduct experiments on a real life rock paper scissors game. Cells in E. coli, a form of bacteria, contain genes that encode a toxin. They also contain genes that encode a protein that makes the cell immune as well as a gene that encodes a protein that causes the toxin to be released. This allows us to classify the types of expressed cells as resistant to the toxin (R), sensitive to the toxin (S), or toxic (T). It turns out that these cell types interact according to a rule not unlike the children's game rock-paper-scissors. T cells will replace S cells because the S cells are sensitive to the toxin. S cells have a growth rate advantage over R cells, so over time S cells will replace R cells, and to complete the cycle R cells replace T cells because they also have a growth advantage.

If these three types of cells are placed on a two dimensional lattice and interact locally, then diversity is maintained. S grow faster than R and R grow faster than T but as the S become prominent locally, the T cells begin to grow very fast. When the locality assumption is relaxed the toxic cells (T's) wipe out all of the sensitive cells (S's) and then the resistant cells (R's), which grow faster than the T's take over the population.

[6] Kerr, Benjamin, Margaret A. Riley, Marcus W Feldman, and Brendan J.M. Bohannan, (2002) "Local dispersal promotes biodiversity in a real-life game of rock-paper-scissors" *Nature* Vol 428: July 2002. pp 171-174.

Kerr, et al ran experiments in three environments that they call flask, static plate, and mixed plate. In flask all three strains of E. coli were well mixed. In static plate, interactions are mostly local, and in mixed plate, there is some mixing and some locality. They found that diversity was maintained in the static plate environment but not in the other two. Local interactions and dispersal would seem to be necessary for diversity maintenance.

They make the very nice point that if producing the toxin costs money then we could see these non transitive relationships all of the time. In fact, let's think of the prisoner's dilemma with three types of players, All C, TFT, and All D. The All D players will defeat the All C players, the TFT players will defeat the All D players, and the All C players will grow faster than the All D players. Others have run agent based models with strategies like this and found that if the agents are in a soup that one strategy wins out but if the agents are on a relatively sparse network or on a grid diversity is preserved.

What I want you to contemplate is the obvious point that spatial segregation decreases interplay. It's hard to construct a model with ten types of agents all growing in a soup that remain viable but it is relatively easy to do so on a grid because of the fact that the types isolate themselves and do not interplay with all of the other types.

5.4 Interactions and Interplay

In Kauffman's NK and NKC models, the N term doesn't have much of an effect on performance. What matters are the K and C terms. In the NK model if we think of each site on Kauffman's strings as an agent, then K is a great proxy for interplay. The problem that I have with the NK model is that the interplay is fixed and random. There is limited ability for the agents to lower the amount of interplay. There are no micro foundations for why one site interplays another. In a system in which agents can spatially position themselves and choose actions that generate certain levels of interplay. This is why Kauffman's "evolution" to the edge of chaos argument is tenuous. He dials up K but each time he does so he wipes out all of the existing connections and creates new ones. Not only is this not what would happen in a system it misses a golden opportunity, namely to establish firmer foundations for the evolution of evolvability. The Packard and Bedau paper shows how this can be accomplished through an evolvable mutation rate. But as this science fiction section hints, it may really be interplay that matters and evolvability may evolve through type reduction, spatial segregation or through increasing or decreasing the number of interactions and the potential interplay.