

Thresholds of Behavioral Flexibility and Environmental Turbulence for Group Success

Andrea Jones-Rooy

Ph.D. Candidate
Department of Political Science
Center for the Study of Complex Systems
University of Michigan, Ann Arbor
5700 Haven Hall, 505 S. State St.
Ann Arbor, MI 48109-1045
ajonrooy@umich.edu

Abstract

Agent adaptability – the ability of agents to change behavioral strategies when it is beneficial to do so – is presumed to be an important part of the robustness of complex adaptive systems (CAS). But, determining when changing behaviors is advantageous for agents has proven quite challenging in CAS research, as sometimes behavioral change is necessary, but other times it can impose costs that exceed benefits. I present the results from experiments using an agent-based model (ABM) designed to discover thresholds after which behavioral flexibility leads to improved societal-level outcomes in groups of agents in dynamic environments. The first major result is that there are thresholds in both levels of flexibility in agent behavior and in levels of turbulence in the environment below and above which there are marked differences in utility gains for agents. In particular, relatively high flexibility leads to lower overall utility scores, as well as, surprisingly, decreased diversity and increased inequality between agents. The second result is that at very high levels of environmental turbulence, the effects of the environment alone on agent utility overshadow any benefits to agents from flexible behavior strategies. This suggests, counter-intuitively, that the best strategy for agents in very dynamic environments is simply to keep behavior constant. The third major result is that there is an interaction between agent behavior and the environment: high flexibility of other agents can effectively make an environment more “dynamic”, which just fuels more flexibility, and leads to a scramble between different strategies with no utility gain. A final theoretical contribution of the paper is that the model is able to show drawbacks to flexibility without relying on costs to changing behavior, as is done in much of the literature on strategy change.

Introduction

In biological, ecological, and social systems, behavioral flexibility is generally held to be a useful characteristic of

constituent actors in a population, with flexibility referring to the capacity of actors to change their behavior or some characteristic of themselves in light of changes in their relevant environment. Work in biology has shown that species that are able to perform more diverse sets of behaviors outperform species with only one strategy for, e.g., finding mates or food (Rossmannith et al. 2006, Bundgaard 2000, Lande and Shannon 1996). Researchers have also found that the heterogeneity of species’ habitats correlates positively with more diversity in populations (Piha et al. 2007, Ehrlich and Murphy 1987, Weiss et al. 1988, Kindvall 1996).

With respect to social systems, March (1991) shows that some balance between exploration (trying new procedures or looking for new solutions) and exploitation (imitating existing strategies) is important for the success of firms. This is because exploration is risky: it doesn’t always yield performance or utility gains – in fact explorers may fail miserably – but the possibility of discovering improved strategies makes some degree of exploration desirable. Exploitation is less risky, but given that the payoffs are guaranteed, this means that as long as there is no exploration, there will be a ceiling on how well the firm can do beyond its existing performance. Complementary work in institutional analysis (North 2006, Hayek 1945) also emphasizes that while institutional stability is desirable, there is no guarantee that that particular institution will produce growth over time (North 2006, p. 363); thus, institutions ought also exhibit some degree of flexibility in order to be successful in the dynamic social environments in which they operate. Thus, it would seem that firms, governments, universities, and other organizations would all do well to undertake measures to be flexible in the face of change.

To the extent that stability is also valued in biological and social systems, it is generally held with regard because of the costliness of change; that is, it costs time and resources for a firm to come up with a new strategy for responding to new market pressures. Similarly, it costs individuals time, resources, and mental exhaustion if they need to invent a new strategy each time they confront a new scenario (Bednar et al., manuscript). With respect to species' flexibility, genetic mutation is similarly held in check, and for good reason: not all mutations are beneficial, and if the population is surviving well, they are not likely to benefit from rampant mutations with each generation.

The brief overview above suggests the following implications regarding the relationship between actor flexibility (strategic/behavioral or characteristic/genetic) and their environment: in constant environments, actors generally would do better to exhibit less flexibility. Flexibility is costly, and if the current strategy insures survival, there is no need to endure that cost. If the environment is heterogeneous (over space or time), however, some degree of flexibility should prove to be beneficial for actors in that environment. Finally, an environment is more than its physical characteristics (e.g., a prairie landscape, or a market), which are effectively exogenously given. An additional important part of what distinguishes one particular environment from another is the other actors in that environment (e.g., the number of predators or competitive firms). If we are interested in understanding the fate of a particular group or individual in an environment, these endogenous actor factors are certainly as important as the exogenous physical ones.

In this paper I present an agent-based model that evaluates this relationship between actor flexibility and environmental dynamics. I consider both topological sources of environmental turbulence, as well as environmental dynamics that come from the changing behaviors or other actors in the system. In my model two types of agents attempt to spread themselves according to contrary agendas across a lattice. In the control case, the environment (the lattice) is constant and agents are confined to staying whatever type they are originally assigned. In subsequent trials, I allow the environment to change, then I allow the actors to change from one type to the other if it is performing better than their own, and, finally, I allow both the environment and actor types to change.

My core result is that under some conditions flexibility can lead to sub-optimality, even in a dynamic environment. This is the case when we consider sub-optimality in the strict sense of lower utility at the societal level than we would see had the agents not exhibited flexible behavior. In addition, flexibility can *also* be suboptimal in a broader sense: I find that more flexibility in agents' behavior leads to less diversity and more inequality. As the following section details, my model only considers two types of

actors. Thus, I treat diversity as the ratio of one type to another – a 50/50 system is the most diverse, while a 99/1 system is the least diverse, which is reasonable if we consider diversity as the probability of running into someone unlike oneself if we interact with others in our society randomly. Inequality here refers to the proportion of wealth (utility, or “contentment”, in this model) held by members of each type – an equal system is one in which each type earns about equal utility, while an unequal system is one in which one group earns much higher utility than the other. Another interesting and related result from the model is that, contrary to what we might expect, the group with the higher utility does not necessarily have more members.

A second result is that the model shows that environment plays a very prominent role in social outcomes. This in itself is not radically novel, but the effect does seem stronger than the literature might lead us to expect. The power of the environment exerts itself in two directions: on one hand, there are simply some situations in which no amount of flexibility can help agents overcome environmental disadvantage. On the other hand, even in situations where the physical environment would allow one group of agents to do well, the presence of a second groups can significantly derail the first's success – even if in the early rounds the first group was the more successful of the two groups. In this particular model, we are also able to observe evidence of some disadvantages from early success. If such derailing takes place after enough members of the first group have reached a particular satisficing threshold, they won't change strategies even if there are additional gains to doing so.

Finally, this model also contributes to the exploration versus exploitation literature in that it is able to generate sub-optimality in flexibility without relying on costs to exploring or changing. Risks are still embedded implicitly in the model, as the agents do not know when they are switching behaviors what the best behavior actually is, but the actual act of changing bears no cost in the model. Much of the literature on cognition, bounded rationality, and decision-making under uncertainty turns on the understanding that it is difficult for individuals to change their behavior (and even more so for groups). This model produces suboptimal outcomes from changing behavior without this stipulation.

Below, I explain the model in detail and identify several hypotheses about the behavior we might expect to observe. Then, I present evidence for the results I've described, after which follows a discussion of the results and comments on directions for further research. In the Appendix the reader will find tables and graphs to support the results. The java code for the model is currently available by emailing the author and will in the future be available for download directly from her website.

The Model

The model is very simple. This offers the advantage of being very generalizable, as it is readily applicable to a variety of social and ecological systems. The model has also been specifically designed so that the addition of context-specific features is straightforward.

In this model 100 agents of two different “Types” live on a 5x5 non-wrapping lattice and attempt to allocate that space among themselves in a way that maximizes their individual utility. One type of agent is more content when it is surrounded by as many other agents as possible, while the other type is less content the more agents are around it.

Two important features to note about the setup are, first, the lattice is non-wrapping in order to more closely approximate reality. In this model agents are concerned with how many neighbors they have, and in most towns and neighborhoods, there are borders, outskirts, and dead ends, to which individuals who want to be more isolated can move. We will see in the results that the existence of corners for agents to move indeed leads to some interesting outcomes that would not be possible in a toroidal lattice.

The second feature is the size of the lattice and the number of agents in the environment. At the setting I will use for all of my results in this paper, the population density of the model is 100 agents over 25 units of space, which means 4 agents per unit, or cell. The model is robust to most variations in density levels except for very high and low ones. This is because, as we will see below, a very dense environment will favor the group of agents that prefer to cluster, while a very sparsely populated environment will mean that the group of agents who prefer to spread will do best.

The features of the agents are as follows. Agents in this world take on one of two possible Types – clustering or spreading – with some probability (P). An agent’s Type defines the criteria by which the agent earns utility. In this way, Type is like a strategy, in that agents want to use the space in the lattice in a way that earns them the highest utility, and the Type defines the rules by which they do so. The first Type is Type C, which earns utility, or “Contentment” (K), by clustering agents together. As the number of other agents occupying a single agent’s neighborhood (defined below) increases, the K for a Type C agent increases. The second Type is Type S, which, opposite to C, earns utility by spreading out over the lattice. As the number of agents occupying the neighborhood in which an agent lives decreases, the K for a Type S agent decreases. Specifically, Contentment scores for each type are given by the following,

$$K_C = (\text{Number of Agents in Neighborhood})/100$$
$$K_S = (100 - \text{Number of Agents in Neighborhood})/100$$

where “Neighborhood” refers to the space of the lattice that is visible to each agent. This will be explained in

detail below, as it is the aspect of the environment that will change when I introduce habitat heterogeneity into the model.

Thus, we see that Type C agents get higher K the more neighbors they have, while Type S get higher K the fewer neighbors they have. For example, if a Type C has 65 other agents in its relevant neighborhood, that agent earns a K score of 0.5. If a Type S has 65 other agents in its relevant neighborhood, however, it earns a K score of 0.35. This contentment is evaluated at each timestep of the model, with the maximum score able to be earned by any agent being 1. The K scores are averaged across all agents of each type at each timestep, and then I evaluate them cumulatively over 50 timesteps, which constitutes one “run”. Thus, the maximum “Social Contentment” that can be earned by a group is 50 (a score of K = 1 each timestep for 50 timesteps)


Action in the model proceeds as follows. At each timestep every agent has the opportunity to move to a neighboring cell if it will increase its own utility. There can be more than one agent per cell. In principle, all 100 agents could end up in one cell. Thus, there are two Types with two different agendas that move simultaneously around the lattice. Moves are simultaneous to reflect the uncertainty in real life decisions. When we make a decision about where to go or what to do, we consider what we think other actors will do; however, we do not ever know for certain what the other actors will do. Even if other people have promised or otherwise made clear their intentions, we never know until after the fact what the outcome is with certainty. By having all agents make decisions simultaneously, this means they are fully taking into account the situation at present and making their decision based on that – which, with few exceptions, is the best most actors can usually do. (Again, promises or signaling might guide our expectations, but not to 100% certainty. Here, the present moment distribution proxies for promises/signaling – if a cell is very full at present, it’s more likely to have agents during the next time step than a cell that is currently empty.)

This setup of two Types also suggests that the proportions of one Type over another will matter. For example, if Type C agents are a minority, their utility scores will be limited, because there will only be so many agents with which they can cluster. If Type S agents are a minority, depending on where in the lattice the cluster of Type C agents is located, and the size of the neighborhood, Type S agents may be stuck in a neighborhood with many neighbors, and thus will also face an upward limit on the K score they can earn.

This concludes the description of the basic model. The next steps are to incorporate a dynamic environment and behavioral flexibility. First, we make the environment dynamic. Recall that the agents are on a 5x5 lattice. To simulate a heterogeneous environment, I allow the “vision” of agents to vary over three possible values: 0, 1, and 2.

When vision = 0 agents can only see within their own cell. When vision = 1 agents can see their own cell plus one cell further in every direction (for 9 cells total). When vision = 2 agents can see their own cell plus two cells further in every direction (for 25 cells total for any agent in the middle of the grid, and fewer for those toward the edges). Importantly, when contentment is calculated, it is only with regard to the “neighborhood” that the agent can see. If vision = 0, then the K expressions given above will only include the number of agents sharing a cell with the agent whose contentment is being calculated.

The three lattices below illustrate this “vision” feature. The shaded grey areas are the relevant neighborhoods to Agent X, positioned in the center, at different levels of vision.



Clearly, Type C agents will do best when Vision = 2 and Type S agents will do best when Vision = 0. (Indeed, 30 runs of each scenario with all C Types and the all S Types confirms: the average cumulative score for Type C is 50 (perfect) when Vision = 2, and approximately 47 for Type S when Vision = 0. This is because when vision = 0 agents are restricted in their movement, so it is sensitive to initial conditions.) To further illustrate, consider the lattices below, where **Vision = 0** on the left, and **Vision = 2** on the right:

0	0		0	0
0	0		0	0
0			0	0
	0		0	X

0	0		0	0
0	0		0	0
0			0	0
	0		0	X

If Agent X is Type S, then it receives a K score of 1 when Vision = 0 because it is the only agent in its relevant

neighborhood (where the relevant neighborhood refers to as far as the agent can see), and $(100-0)/100 = 1$. If the environment changes to Vision = 2, however, Agent X’s score decreases to 0.97, because now there are three other agents in its relevant neighborhood.

When I run trials of the model where vision changes, I allow it to change with some probability (P_V), where $P_V = 1$ means the vision changes at each timestep and $P_V = 0$ means the environment is static. An additional major parameter that we vary is the probability with which actors change Types. As we just saw above, if Vision consistently = 2, then it pays to be a Type C agent. Allowing agents to change from Type S to Type C if some particular (preset) satisfaction threshold is not met means they can capitalize on this environmental condition and earn much higher K scores than if they had to remain Type S. As with P_V , the probability of changing Type (P_T) can vary from $P_T = 1$, where agents *always* change types when their threshold isn’t met (maximum flexibility), to $P_T = 0$, where agents never vary their Type, no matter how poor their performance (minimum flexibility).

Finally, we can also vary the threshold below which agents are triggered to change strategy, where lower thresholds correspond with quicker satisficing. Recall that a perfect K score over 50 runs is 50. If the threshold above which agents stop changing their type is 1, this means that very early on we should expect to see agents stopping changing strategies; even if later on the environment changes such that it would be advantageous to change strategies, these agents will not change. (While this may seem like a disadvantage, we actually see in the model that, contrary to what we might expect, extremely high thresholds are surprisingly associated with lower total K scores.)

To summarize, the events that take place during a single timestep are the following:

1. Agent X is located somewhere on the lattice
2. Agent X considers all relevant neighbors (given by Vision)
3. Agent X considers the number of agents in neighboring cells; if Agent X is Type C, will move to a new cell if the number of agents in the neighboring cell exceeds the number in Agent X’s current cell (if there are more than one cells with the most, Agent X will move to the cell with the highest number). If there are no neighboring cells with higher numbers of agents, Agent X does not move. If Agent X is Type S, the exact same process holds, with the exception that now Agent X evaluates cells with respect to which has the lowest number of agents.
4. Once Agent X moves (if the agent decides to move; if not, once the decision to “not move” has been made), Agent X considers its Contentment (K). If K is below the Agent’s satisficing Threshold (T), then Agent X changes Type with

some probability (P). (Note: in the base case, $P = 0$.)

5. The current timestep concludes and a new one begins with Agent X considering all current relevant neighbors. If Agent X changed Types in the previous timestep, Agent X now plays the new round as the new Type.

One “run” of the model is 50 timesteps. For nearly all runs, 50 was more than enough to reach convergence to a long run equilibrium in terms of the proportion of Types in each population, as well as to establish a clear Type winner (the Type with the highest K) for that particular run. For the few cases where even after 50 timesteps there was not clear convergence, this lack of convergence was only with respect to outcomes; i.e., it was not necessarily clear by 50 timesteps which Type would come out with the highest K. However, for all runs the long-run distribution of Types was apparent within the first 10-20 timesteps (see graphs in the Appendix).

Results

There are three central results. First, flexibility can lead to sub-optimality, (a) directly in terms of lower K scores, and (b) indirectly in terms of reduced diversity and greater inequality between agent Types.

In terms of the directly effects of flexibility on utility, it turns out that long periods of switching between types means that agents do not stay at a particular Type long enough to ever earn many points. That is, it takes a few periods of agents attempting to cluster before they get close enough to really earn high scores (generally the 0.8-1 range). Similarly, a pack of Type S agents also need a few timesteps to move around before they’ve spread themselves out. If the agents are too “picky” and they switch back and forth all the time, the entire period during which they are switching they aren’t earning as many points as they might if they would just pick one Type and stay there. What’s more, since these agents have high thresholds, it takes them that many more timesteps of switching in order to reach those high thresholds, after which they finally converge on a consistent distribution of Types.

With respect to the indirect effects, as can be seen in the graphs in the Appendix, flexibility can mean that the agents switch Types frequently early on, and then get locked in to a certain ratio of types once the satisficing threshold is triggered (and this is robust to any threshold). The agents who do well early on, and thus trigger their “stay at this Type” threshold earlier than other agents end up losing out tremendously when/if the system ends up being one that favors agents of one Type over another. In addition, the lower the threshold, the more agents there are who can capitalize on this new “knowledge” of which strategy is best, which exacerbates inequalities – we’ll see far more of the successful type than the unsuccessful,

which just makes the successful types even more successful. (Again, this is particularly startling if the system began slightly biased in favor of the Type that ends up losing!). The lack of diversity appears here, too, for in these cases we see an average of 91% of agents taking on the dominant Type, leaving only 9% stuck in the minority with no way out given the parameters of the model.

The second major result is that environment matters, and in two ways. The first is with respect to the physical landscape, or topology, itself. As hinted in the description of the model, if $Vision = 2$, then there is little hope for any Type S agent to receive a high K score, no matter where on the lattice it moves. Even if we allow these unfortunate Type S agents to change types, unless the threshold $= 0$ and the $P_T = 1$, there will always be a not-insignificant period during which agents who started out as Type S earn low K scores.

The additional way that environment matters is that how well one Type does depends on what the other Type is doing. As can be see in the graphs in the Appendix, many times if Type C and S are doing about the same for the first few runs, all it takes is for one Type to earn a slight advantage before the system suddenly turns to favor that Type. This is because if Type S agents end up moving to cells on some turn that earn them enough points to pass a threshold (and this result is robust to threshold levels), this means there will be more Type S agents in the next turn, as some Type C agents will switch. This, then, means that on the subsequent turn, Type S agents will do even better, because there will be that many more of them spreading out rather than clustering. By the same token, Type C agents will do that much worse, because there will be fewer agents with which to cluster. Note that half the time this result held for the reverse: if C got even a slight advantage on one turn, the population would veer toward favoring C’s.

This points to the final result that there are disadvantages to flexibility that do not rely on costs to changing behaviors. Originally when I was implementing the model I planned to add a “cost to change” effect in order to generate limitations on benefits from flexibility. Surprisingly, it turned out I did not need to add this effect in order to see negative aspects of flexibility in the agents. To be sure, I added a constant cost component to changing strategies, and in fact I found it slightly improved outcomes, as agents flipped back and forth between unsatisfactory strategies less. This result was not particularly strong, but suggests a direction for further research that would include costs as a function of the magnitude of the change, or costs that vary over time.

Conclusion

This model provides some preliminary insights into points after which behavioral flexibility can be disadvantageous, even in a dynamic environment. It is important to note that

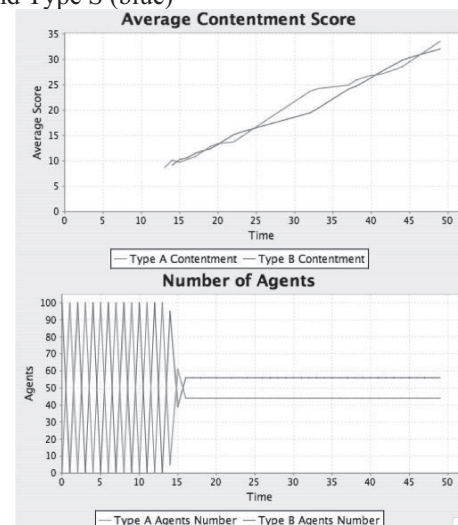
these results are for inherently competitive systems. Non zero-sum situations may have different dynamics. That said, the results presented here from a competitive system are of course not to suggest that flexibility is bad. Indeed, if Type S agents were stuck in an environment with Vision = 2 and did not have the ability to change, then we would see consistently low scores with no way out (see Appendix 2). Future work to build on this is application of these insights to real world cases where we see groups stuck in suboptimal situations due to precisely the dynamics described here. Additional work after that is to determine the configurations to tweak the group out of these spirals of suboptimality, which can include everything from low payoffs to lost diversity to increased inequality.

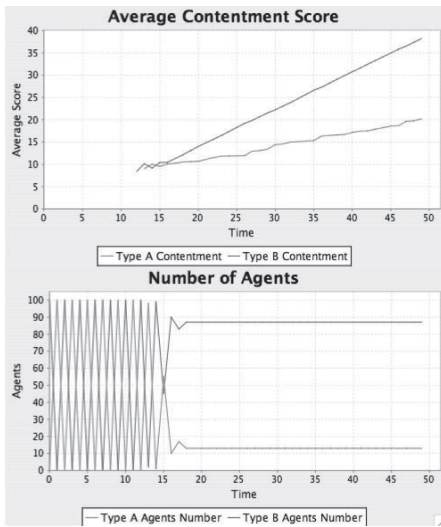
References

- Bednar, Jenna, Aaron Bramson, Andrea Jones-Rooy, and Scott Page. mss. The Emergence of Cultural Signatures and Persistence of Internal Diversity: A Model of Conformity and Consistency (*under review*).
- Bundgaard, Boerma. 2000. Does inbreeding affect the extinction risk of small populations? Predictions from *Drosophila*. *Journal of Evolutionary Biology* 13(3), pp. 502-514.
- Ehrlich, Paul R. and Dennis D. Murphy. 1987. Conservation Lessons from Long-Term Studies of Checkerspot Butterflies. *Conservation Biology* 1(2), pp. 122-131.
- Hayek, Friedrich. 1945. The Use of Knowledge in Society. *American Economic Review* 35(4), pp. 519-530.
- Kindvall, Oskar. 1996. Habitat Heterogeneity and Survival in a Bush Cricket Metapopulation. *Ecology* 77(1), pp. 207-214.
- Lande, Russell and Susan Shannon. 1996. The Role of Genetic Variation in Adaptation and Population Persistence in a Changing Environment. *Evolution* 50(1), pp. 434-437.
- March, James G. 1991. Exploration and Exploitation in Organizational Learning. *Organization Science* 2(1), pp. 71-87.
- North, Douglass C. 1994. Economic Performance Through Time. *The American Economic Review* 84(3), pp. 359-368.
- Page, Scott E. 2007. *The Difference: How the Power of Diversity Creates Better Groups, Firms, Schools, and Societies*. Princeton, NJ: Princeton University Press.
- Piha, Henna, Miska Luoto, Markus Piha, and Juka Merila. 2007. Anuran abundance and persistence in agricultural landscapes during a climatic extreme. *Global Change Biology* 13(1), pp. 300-311.
- Rossmannith, Va, Volker Grimm, Niels Blaum, Florian Jeltsch. 2006. Behavioral Flexibility in the Mating System Buffers Population Extinction: Lessons from the Lesser Spotted Woodpecker *Picoides Minor*. *Journal of Animal Ecology* 75(2), pp. 540-548.
- Tilly, Charles. 1992. *Coercion, Capital, and European States AD 990 – 1992*. Oxford: Blackwell.
- Weiss, M. J., D. E. Cole, K. Ray, M. P. Whyte, M. A. Lafferty, R. A. Mulivor, and H. Harris. 1988. A Missense Mutation in the Human Liver/Bone/Kidney Alkaline Phosphatase Gene Causing a Lethal Form of Hypophosphatasia. *Proceedings of the National Academy of Science* 85(20), pp. 7666-7669.

Appendix 1: Graphs of Outcomes

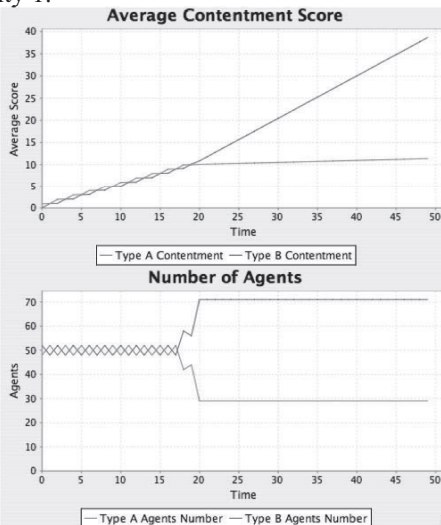
First, we see equality and inequality in two sample runs. Here $P_T=1$, $P_V=1$, and begin with equal number Type C (red) and Type S (blue)




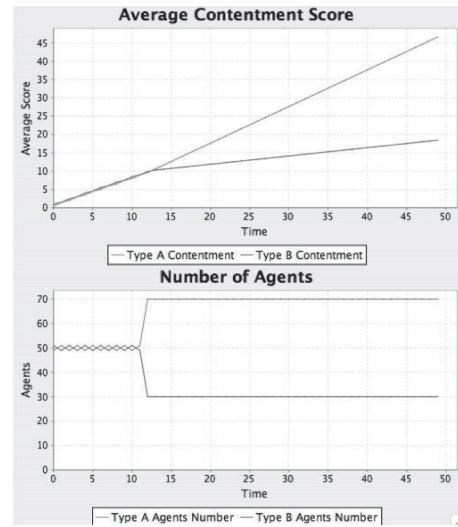


Equality is on the left, inequality on the right. Notice that in the right the two strategies begin at about the same well-being, but just slightly more Type S starts to tip it, until there are a lot more Type S, and Type C is permanently lower simply because they do not have the numbers required to increase their K.

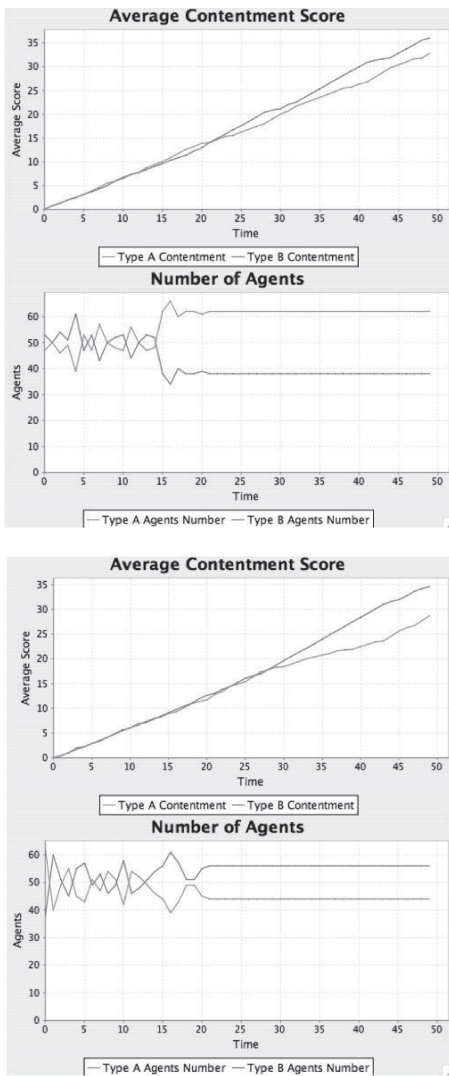
Below is another graph showing the influence of just one or two timesteps. This case begins with 50 of each Type of agent, and all are endowed with 100% probability to change Type, and the environment changes with probability 1.



Below, we see outcomes when Vision is fixed at 2 (Left) and 0 (right). Clearly Type C outperforms S on the left, and vice-versa on the right. Here $P_T = 1$, but notice that some C's and S's still get "stuck" in the losing team because their thresholds were met before the split in K scores.



Below we see that decreasing flexibility can lead to greater equality. The below two graphs show typical runs when Vision change = 1, but the probability of changing Type is only 0.5. We begin with 100% type A in both.



Appendix 2: Results from Base Cases

I first present outcomes of the model in an unchanging environment. Recall that vision can take on values 0, 1, or 2. The table below shows the average outcomes over 30 runs for each “vision environment” when there are an equal number of Type C’s and Type S’s and no agents change their type.

Average Aggregate Contentment (K) Per Type in 3 Constant Environments

	Vision = 0	Vision = 1	Vision = 2
Average K Type C	2.2	27.8	50.0
Average K Type S	47.9	42.7	20.1

Table 1: Average Aggregate Contentment (K) Per Type in 3 Constant Environments. Here the population is 50% Type C and 50% Type S. The probability of vision change and the probability that an agent will change Type are both zero.

The next table presents average aggregate outcomes over 30 runs where I vary the turbulence in the environment. Specifically, I vary the probability that the agents’ vision will change at any timestep. Again, there are 50 Type C and 50 Type S agents.

Average Aggregate Contentment (K) Per Type in Changing Environments

	Prob. Change Vision = 1	Prob. Change Vision = 0.9	Prob. Change Vision = 0.5	Prob. Change Vision = 0.1
Average K Type C	33.0	31.6	34.2	31.3
Average K Type S	36.6	37.4	34.8	35.9

Table 2: Average Aggregate Contentment (K) per Type in Changing Environments. Here the population is 50% Type C and 50% Type S. The probability of an agent changing Type is zero, while the probability that an agent’s vision (environment) will change is varied between 1 (i.e., most turbulent: it changes every timestep), 0.9, 0.5, and 0.1).

Now I want to see the behavior of the two Types when they are in a homogenous society; that is, how well do C’s do in the absence of S’s, and vice/versa?

Average Aggregate Contentment (K) per Type in Homogeneous Society

	Vision = 0	Vision = 1	Vision = 2
Average K Type C	2.1	49.6	50.0
Average K Type S	47.9	44.5	37.5

Table 3: Average Aggregate Contentment (K) per Type in a Homogeneous Society. This table presents the average K over 30 runs when only one Type is present. The probability of the agents changing type and the probability that the environment changes are both zero.

In the below tables I allow agents to change strategies (threshold = 10) while keeping the environment constant. The table below presents the average aggregate contentment scores for both Types S and C after 30 runs of the model. In these cases I begin with 100% Type C.

	Vision = 0	Social Score
Ave. K Type C	11.3	178.5
Num. Type C	15.8	
Ave. K Type S	38.5	3241.7
Num. Type S	84.2	
	Vision = 2	Social Score
	47.3	3102.9
	65.6	
	20.1	691.44
	34.4	

Now I replicate the above experiment, but time beginning with 50% Type C and 50% Type S.

	Vision = 0	Social Score
Ave. K Type C	11.4	307.8
Num. Type C	27.0	
Ave. K Type S	38.8	2832.4
Num. Type S	73.0	
	Vision = 2	Social Score
	46.7	3530.5
	75.6	
	19.2	468.5
	24.4	