Emerging Local Dynamics: A Model on Cultural Dissemination

Clara Orbe[1]

Mentor: Michelle Girvan[2]

[1] Brown University,

Department of Applied Mathematics, Providence, RI

[2] Santa Fe Institute,

Santa Fe, NM

August 19, 2006

Abstract

Social scientists have explored many ways to answer why differences persist as people tend to become more alike in their beliefs, attitudes, and behaviors. Robert Axelrod explored one such approach via an agent-based model on cultural dissemination where local influence does not lead to global homogeneity. The model represents culture as a string of bits where each bit is a cultural trait for which agents assume one of a possible set of features. Via this abstraction the model defines cultural exchange as random bit swapping between agents who converge locally to similar cultures in proportion to their existing cultural similarity. One mechanism that explains why we are not all the same, the model emerges stable homogeneous regions from an initially random population. I propose that collapsing bit strings into single outputs and positioning them on a one-dimensional value spectrum reveals counter-intuitive local agent movement that Axelrod's model does not capture. I ask, can looking at how agents travel locally on a simple spectrum surface interesting dynamics and is this behavior consistent with real-world cultural diffusion? I also construct a variation on Axelrod's algorithm and observe any global and local changes in cultural evolution.

I thank John Miller, Sam Bowles, Rajiv Sethi, Ginger Richardson and many others for their support and advice during my stay at the Santa Fe Institute. I extend a large thanks to the institute as a whole for supporting me in many capacities.

1 Introduction

1.1 Motivation and Background

Current events evidence that slight differences can manifest large affects. In the United States 2000 presidential race Ralph Nader of the independent Green Party gained 2% of the popular vote, a slight margin (i.e. votes 'stolen' from Democratic candidate Al Gore) that tipped the end vote distribution in favor of Republican candidate George W. Bush.[1] In majority voting systems, where success is all or none, marginal differences can impact on large scales. Given this motivation, I explore agent-based algorithms that generate these globally marginal differences.

As they impact in highly unpredictable ways, third parties often escape formal models. One route to exploring how third parties might effect change, spatial voting models position major parties on electoral landscapes shaped by voter preference distributions and electoral rules [3]. These models treat parties as adaptive, constantly adjusting their platforms to a shifting landscape [5]. The literature on third parties, though, debunks third parties as adaptive, suggesting that third parties never adjust their platforms to voter preferences since doing so gains them neither the majority vote *nor* independent voters [1]. I did not encounter any models that explicitly address **how** third parties effect change largely since complexities that are difficult to model inevitably weave into elections (i.e. voters' expectations of winners). In turn, I explore general models on social influence, in particular cultural diffusion models where simple algorithms manifest interesting local dynamics.

Models in comparative politics that explore state formation, domestic cleavages and succession conflicts try to address how cultures form [2]. Cultural formation is a multi-scale problem for which realistic modeling seems almost futile. Recent agent-based models, though, approach the problem 'bottom-up' by assigning agents observed rules of cultural exchange and emerging real-world global behaviors [4].

Applying the diffusionist[2] definition of culture, Robert Axelrod constructs a model on cultural dissemination that represents culture as a bit string where each bit represents a trait (i.e. religion). Here, culture refers to any set of individual attributes subject to social influence. His model acts an algorithm on a population of randomly distributed strings where neighbors swap bits in local exchanges of culture. Local rules of convergence observe that individuals who are culturally similar are more likely to interact and become more culturally similar [2]. See Figure 1.

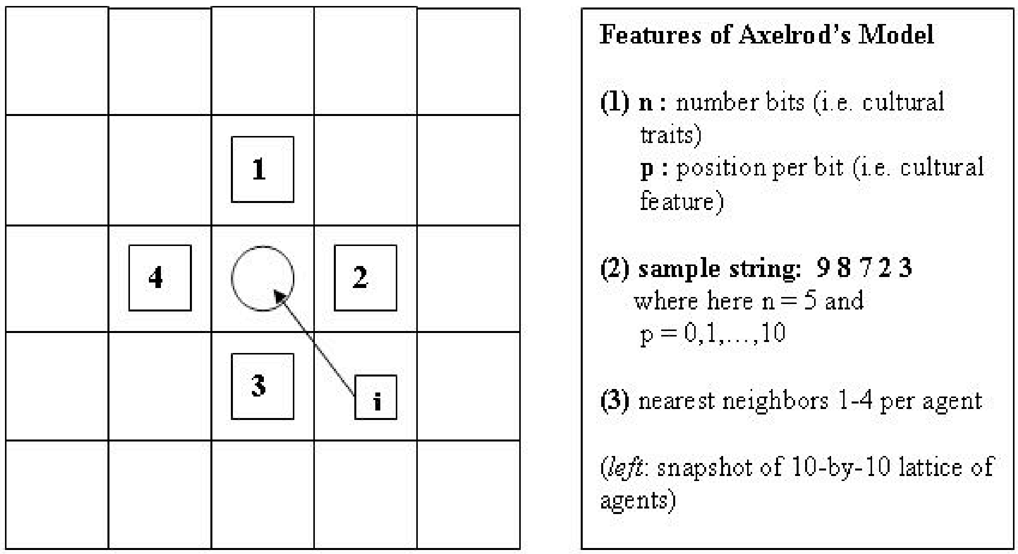

Figure 1: Population schematic and particulars of Robert Axelrod's model on cultural dissemination. See Figure 2 for algorithm rules.

Axelrod concludes that repeated steps of local convergence evolve an initially random population into a steady state of stable homogeneous sub-populations (i.e. regions where every agent is the same). I suggest that Axelrod's algorithm is more than just a mechanism for maintaining cultural difference–that is, when viewed locally, cultural evolution surfaces interesting counter-intuitive agent behavior that Axelrod's model, as presented, neither addresses nor suggests. *Can a local approach to cultural exchange surface interesting behavior with real-world corollaries?*

2 The Model

Axelrod establishes his model as a mechanism for globally maintaining cultural difference but, returning to my original motivation, I suggest that exploring local dynamics affords a better understanding of how slight differences form. While Axelrod's model looks over certain essentials of cultural dissemination in its neglect of social networks, it is still very useful as an algorithm based on a simple realistic notion of converging cultural exchange.

I explore two algorithms for cultural formation: (1) the first algorithm replicates the rules of Axelrod's model while (2) the second treats cultural diffusion as a more localized spreading. Note that Axelrod's model exercises diffusion on a population of 100 agents; I run evolution on 10000 agents simply as a more realistic scale. See Figure 2. I compare algorithms to see how slight algorithmic changes manifest global differences, if any at all.

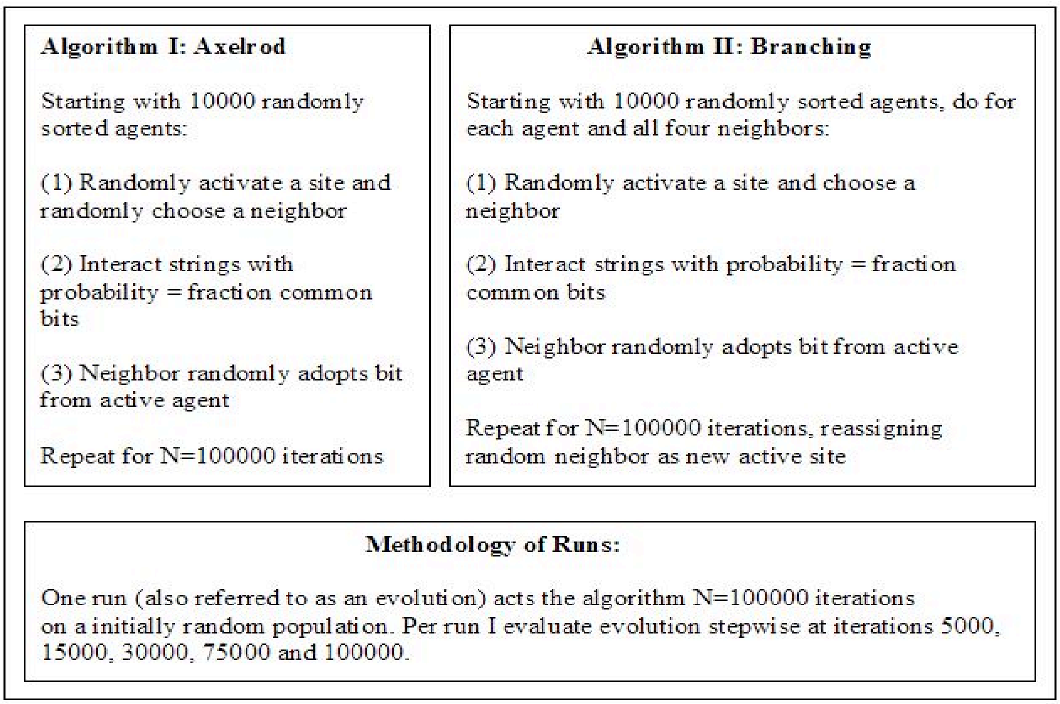

Figure 2:

2.1 Composite Value Spectrum

Cultural similarity, measured as the fraction of common bits between any two strings, guides cultural exchange. In turn, evolution depends on *hamming distance*, the total number of bits on which any two strings differ. While hamming distance indicates the likelihood of cultural exchange between neighbors I ask, 'How do strings relate outside of this nearest neighbor relation?' I propose an alternative way to relate strings that borrows from traditional voting models: a one-dimensional spectrum. I compute single composite values per string and, based on these outputs, position strings on an ordered spectrum. See Figure 3.

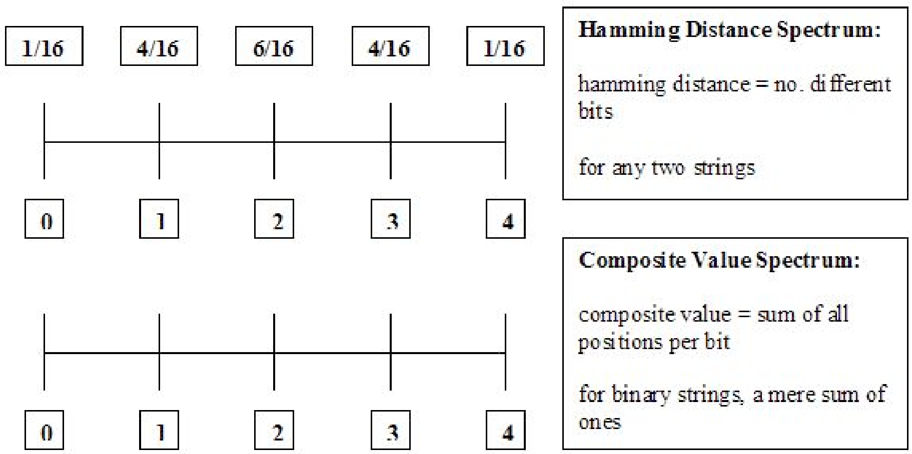

Figure 3: Hamming distance and composite value spectra. Values 0-4 correspond to distances and values respectively while the top proportions distribute strings normally in a 4-bit binary string space along the spectra (same for both).

Axelrod's distance metric (hamming distance) affords only the probability of how nearest neighbors will interact but the spectrum captures the indirect relations between non-neighboring strings in an evolving population. For any three agents A, B and C, where B connects A and C, a value spectrum tells how A and C move relative to one another, a far more powerful description of evolution than a mere probability of interacting *only given* A and C are nearest neighbors. I suggest that exploring these non-neighboring relations, explicitly manifested on a spectrum, might afford a more meaningful (and counter-intuitive) interpretation of evolution[3].

For convenience, I impose no meaning to the spectrum. One, though, can easily interpret a string as a set of issues for which one can assume the stances pro ($p = 1$) or con ($p = 0$). In this case, composite value represents an agent's overall support for the string of issues. In their own models, Durrett and Levin have proposed 0 and 1 as Democratic and Republican positions respectively. In this sense, composite value positions an agent at an ideological medium between Democrat and Republican extremes. Other models using bit string abstractions propose alternative ways of comparing strings (i.e. spatial voting models that interpret value as a voter's distance from a party's position on a given issue) [3].

3 Analysis

I posit several questions about cultural evolution on both local and global scales.

3.1 Part I: Global Analysis

Before imposing composite value as a window into local dynamics, I observe how the algorithms evolve initial populations of randomly distributed agents. I ask, *Globally, what does cultural evolution look like?*

(Q1) Do both algorithms evolve the population into sub-populations? That is are both algorithms consistent with Axelrod's claim that "local convergence leads to global divergence?"

(Q2) Does the persistence of a string depend on its position on the value spectrum? Do the algorithms 'prefer' strings of certain values over others?[4]

3.2 Part II: Local Analysis

Here I take a two-fold approach to evolution: (1) How does the population evolve with respect to different agents in the population? and (2) can a simple value spectrum on which agents move with evolution reveal how neighboring and non-neighboring agents interact in counter-intuitive ways? More specifically:

(Q3) How does the population move with respect to different strings? [5]

(Q4) As agents converge locally with each interaction (using the hamming distance metric), how do they move relative to one another along the spectrum (using the composite value metric)?

4 Global Results

4.1 Q1

Axelrod required 81000 iterations of his algorithm to complete evolution at a stable state of homogenous regions. Treating evolution as a run of 100000 iterations, what I deemed computationally feasible given time constraints, I found that neither algorithm evolved the populations into steady-states. From Figure 4 one can arbitrarily fit step-wise time-series runs of the algorithms but I reserve finding a theoretical steady-state for future work. It might be useful to conjure this expected time for evolution for the system but this is complicated by (1) the increased population (order of three) *and* (2) the decreased string space (from 100000 to 16). For now, simply note that evolution is well under way by 100000 iterations and that Axelrod's algorithm is faster, most likely because of the Branching algorithm tends to loops.

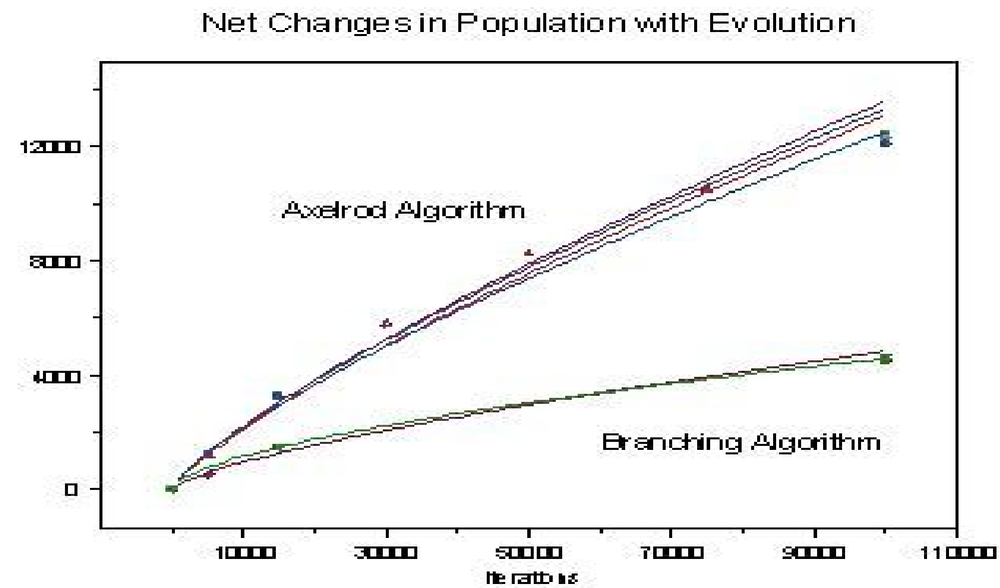

Figure 4: Step-wise evolutions of initially random populations via two algorithms. I run one evolution for $N = 100000$ iterations, noting changes step-wise at $N = 15000, 30000, 50000$ and 75000 where I define a net change as a bit that swaps position from its original state by the end of N iterations. Power curves connect data from several trials but purely serve to guide the eye (cannot extrapolate an eventual steady-state).

While neither algorithm generates a steady-state can one observe the emergence of sub-populations? Subpopulation formation entails (1) an increased cultural similarity between agents within a homogeneous region and (2) the appearance of straddlers between homogenous regions, bordered simultaneously by identical and almost opposite neighboring strings. Here I measure cultural similarity as an agent's neighbor hamming distance, averaged over four nearest-neighbor pairings. That is, for every agent i where hamming(ij) denotes the hamming distance between neighbors i and j for all possible neighbors j=1:4

average agent hamming(i) = average hamming(ij)

A distribution of average agent hamming values explicitly renders how agents are distributed according to their cultural similarity with nearest neighbors. From observations (1) and (2) above one expects an initially normal distribution to tend left as sub-populations form and individuals become more similar to their neighbors. See Figure 5. While Branching evolution clearly lags in increased cultural similarity and Axelrod's algorithm manifests a strong left-leaning distribution, neither algorithm supports the emergence of sub-populations (note the marginal rise in value 0). What this suggests is that, for 100000 iterations, one can at best extrapolate to a steady-state far in the future which implies that Axelrod's model, set on 100 agents, is probably an unrealistic representation of cultural subpopulation formation.

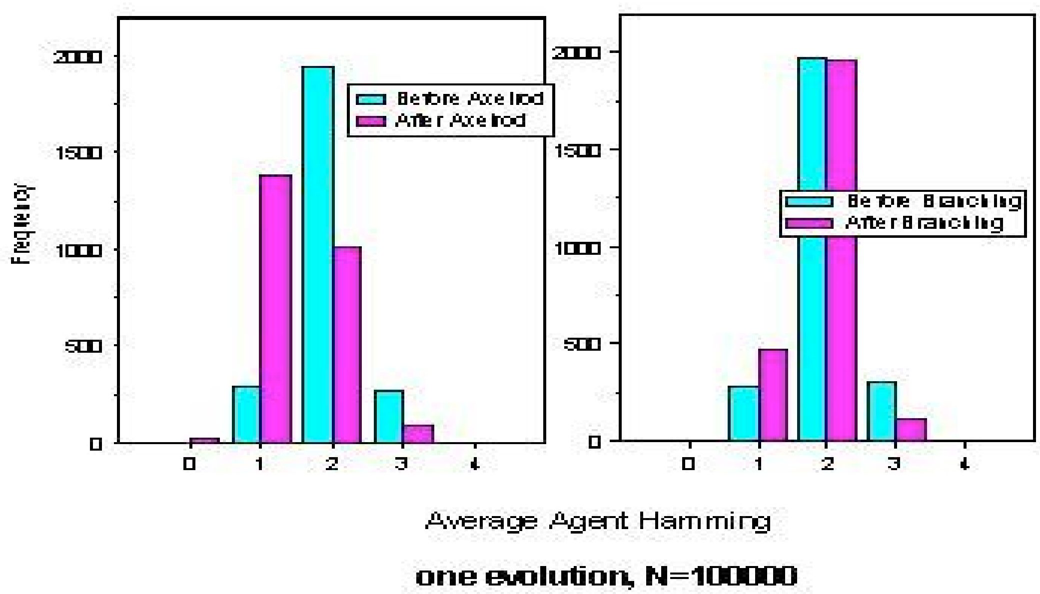

Figure 5: Changing distributions of agents' degrees of similarity with neighbors before and after evolution. Here cultural similarity is an averaged hamming distance; little points to the formation of sub-populations by either algorithm.

4.2 Q2: Does Evolution Prefer Certain Values over Others?

As constructed, the composite value spectrum normally distributes strings in a 4-bit binary space. Since both algorithms dictate successions of random events one expects evolution to preserve this normal distribution, evidenced nicely by the plots in Figure 6. Slight changes in value frequency result purely from variations in the systems and they are neither consistent nor predictable. Therefore, evolution generates few globally interesting results and Part II locally explores these slight differences as hubs of counter-intuitive agent movement.

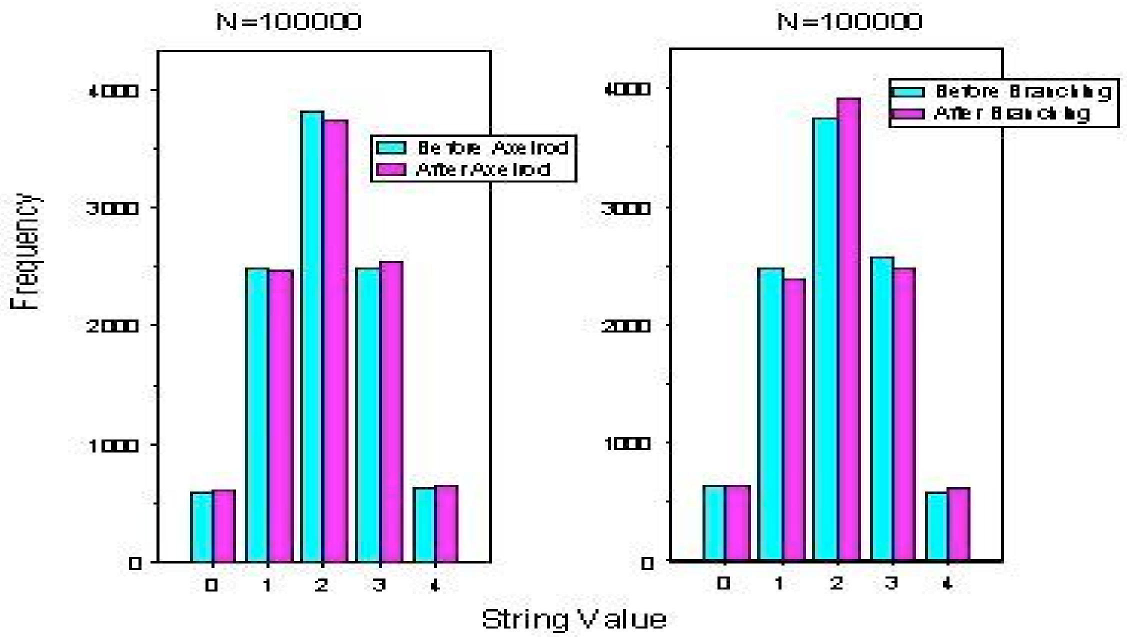

Figure 6: The distribution of string values before and after evolution.

4.3 Q3: Evolution as Relative Staying and Moving In/Away

The schematic below treats evolution from the perspective of a single agent. To each agent I associate a hamming distance of 0, 1, 2, 3, and 4 into which, with each iteration, other agents in the population stay or move in/out. Axelrod's algorithm explicitly activates an agent and carries out a cultural exchange, moving on to another activation and repeating the same steps. What affects, though, do these changes have on other neighbors in the system who are not explicitly activated or implicated by the algorithm? That is, consider evolution outside of the nearest-neighbor context: Does evolution (viewed locally) afford a new insight into how agents interact?

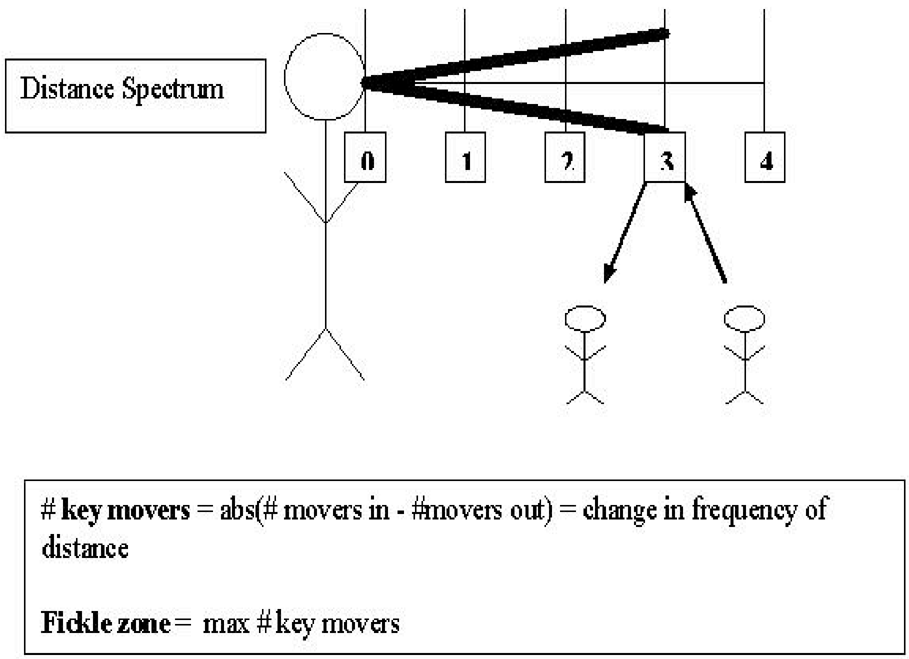

Figure 7: A schematic applying Axelrod's distance metric that positions strings at hamming distances relative to agents in the population. As evolution proceeds, agents move in and out of these **fickle** zones. In this figure I arbitrarily designate distance 3 as the most fickle, visualizing it as the region with the highest number of key movers.

I chose to look at evolution locally with respect to the strings 1110 and 0000, 'moderate' and 'extreme' respectively (see Q4 for my motivation for choosing these perspective strings). Per evolution of a random population I observed how strings moved relative to 1110 or 0000 where I measured movement as changes in relative hamming distance distributions. That is, for a fickle zone 2 with respect to 1110, I specifically targeted agents of altered strings who moved in/stayed/moved out of a hamming distance 2 of 1110. I identified movers in as agents of some other relative distance 'moving into' a relative distance of 2 (other definitions follow in suit and are self-explanatory). As the schematic reads, I captured the notion of changing distances relative to the agent (1110 or 0000) via the term **fickle zone**. In the end, most local dynamics are felt in zones where agents move in for which no other agents move out–that is distances with the greatest change in 'key movers' [6]. I ran several evolutions per algorithm and found that one can neither predict which zones are most fickle nor the distribution of key movers per zone. Interesting dynamics factor in when I reintroduce composite value in Q4. For that reason, I reserve plots of changing hamming distance distributions for the next section.

4.4 Q4

Looking at evolution as a relative series of moving in/out does not seem important until I refer back to the composite value spectrum again. In Q3 I found that, though certain distances are more fickle than others (purely as a function of the normal distribution of strings in the string space), each evolutionary run is independent of others. For that reason, fickle zones are not predictable. At this point, I isolate the most fickle zones from the previous trials and try to explain how agents staying in or moving in/out of these distances travel on the value spectrum. Do these agents, particularly key movers for which no other agents balances their effects, move unexpectedly along the spectrum? This is an attempt, via a local view and an alternative distance metric, to better understand those very agents who generate slight differences that can effect largely in political scenarios [7].

Figures 8-10 present three different scenarios. Figures 8 and 9 plot the final values of agents associated with hamming distances 1 and 2 from 1110 respectively while Figure 10 hones in on various fickle zones with respect to 0000. I discriminate between Axelrod and Branching algorithms when necessary but, overall, both algorithms show similar dynamics, evidencing that slight algorithmic changes do not manifest different agent behavior on the composite value spectrum.

The figures below require some explanation as they incorporate results from EQ 3 that I intentionally reserved for this section. I ran each algorithm 30 times and chose representative runs where I define representative as consistent for a random system (i.e. predictable). Figure 8 observes evolution with respect to 1110; the bottom plots identify fickle zones 1 and 2 for two separate runs. How then, do I derive any meaningful interpretation about agent behavior? I do this by isolating maximum fickle zones (with associated arrows) and generating distributions of the final values of all agents associated with these fickle zone (i.e. where these agents end up on the composite value spectrum after 100000 iterations). Whereas hamming distance affords one only an 'in/stay/out' notion of movement, composite value visualizes evolution as agents travelling relative to one another along the spectrum. In Figure 8 the plots for fickle zones 1 and 2 respectively show that most agents end up at values 2 and 4 with relative frequencies 3:1 (stayers) while those who move out travel to values 1 and 3 [8]. This is not particularly interesting as one would expect this type of moderate movement along the spectrum. See the schematics in Figure 11 for this intuitive look at local movement along the composite value spectrum.

I repeat this interpretation for Figures 9 and 10. Figure 9 also observes intuitive behavior for Axelrod's algorithm while Figure 10 affords counterintuitive, if baffling, behavior by observing evolution with respect to 0000. That is, looking at evolution relative to 0000, an extreme along the spectrum, affords a window into interesting local dynamics one would not suspect in Axelrod's original model. Agents associated to highly fickle zones all behave similarly on the spectrum, implying the emergence of collective agent behavior (an interesting enough result in itself). Moreover, these agents move out of these fickle zones in curious ways. For example, in plot IV of Figure 10 agents exit a fickle zone 0 of 0000 and, after evolution, end up at precisely a composite value 4, the opposite extreme of the spectrum! Intuition would tell us that (1) evolution least affects extreme strings (purely as a function of their low fraction in the string space) and (2) these strings would change at most to composite values of 1 or 2. However, enacting the algorithm shows precisely a *radical* move from one end the spectrum to the other, implying that, while these agents might be few (and almost globally invisible) they exhibit radical behavior on the local scale. This extreme behavior, if surfaced on a political spectrum like that proposed by Durett and Levine could manifest few but radical differences. See the schematics in Figure 12 as examples of this counter-intuitive movement along the value spectrum.

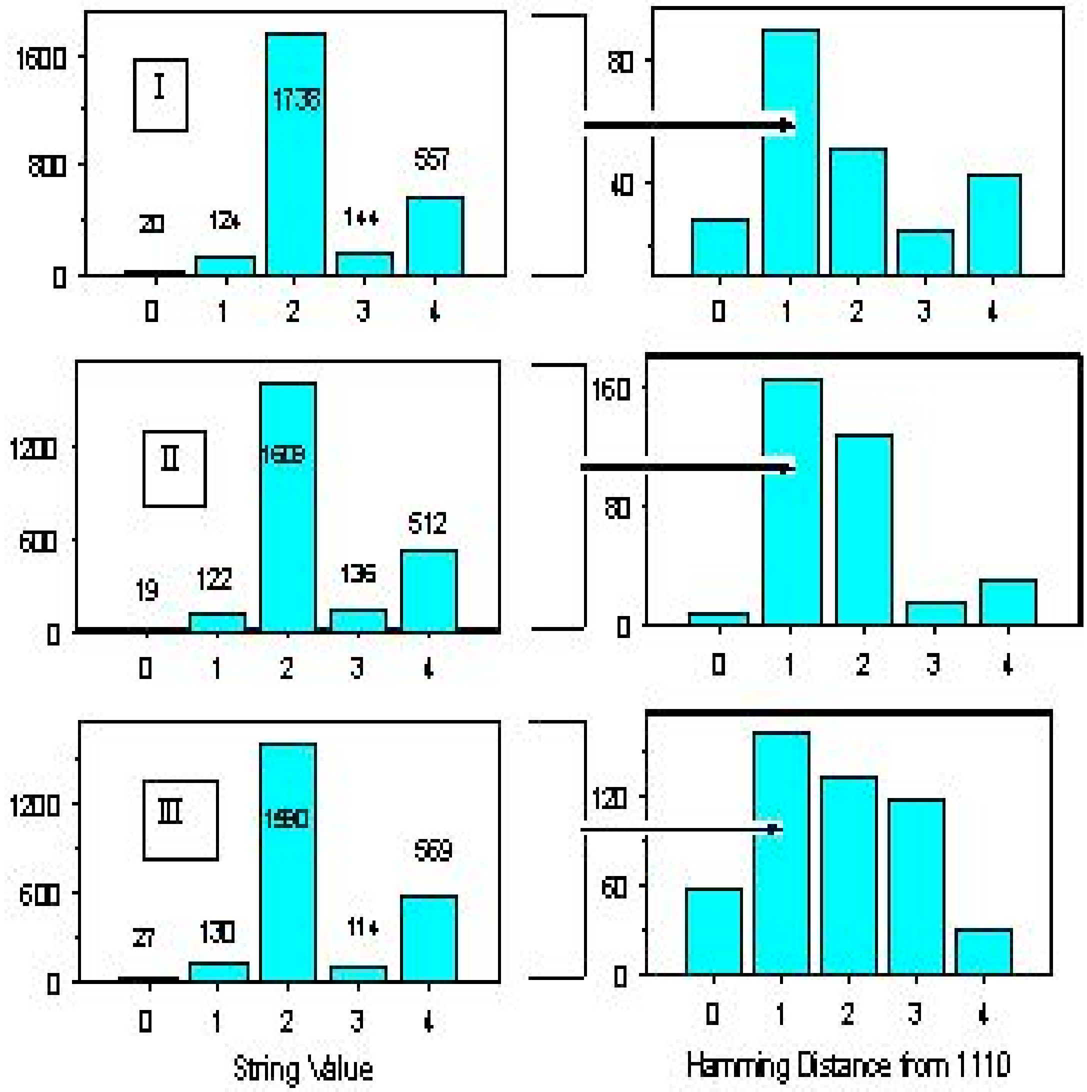

Figure 8: Observe final composite values of agents associated to the selected fickle zones. The top plots describe agents who move in/stay/move out of these fickle zones whereas the schematics in Figures 11 and 12 visualize only agents moving out of these hamming distances on the value spectrum. All plots taken from runs of Axelrod's algorithm (Branching produced the same results).

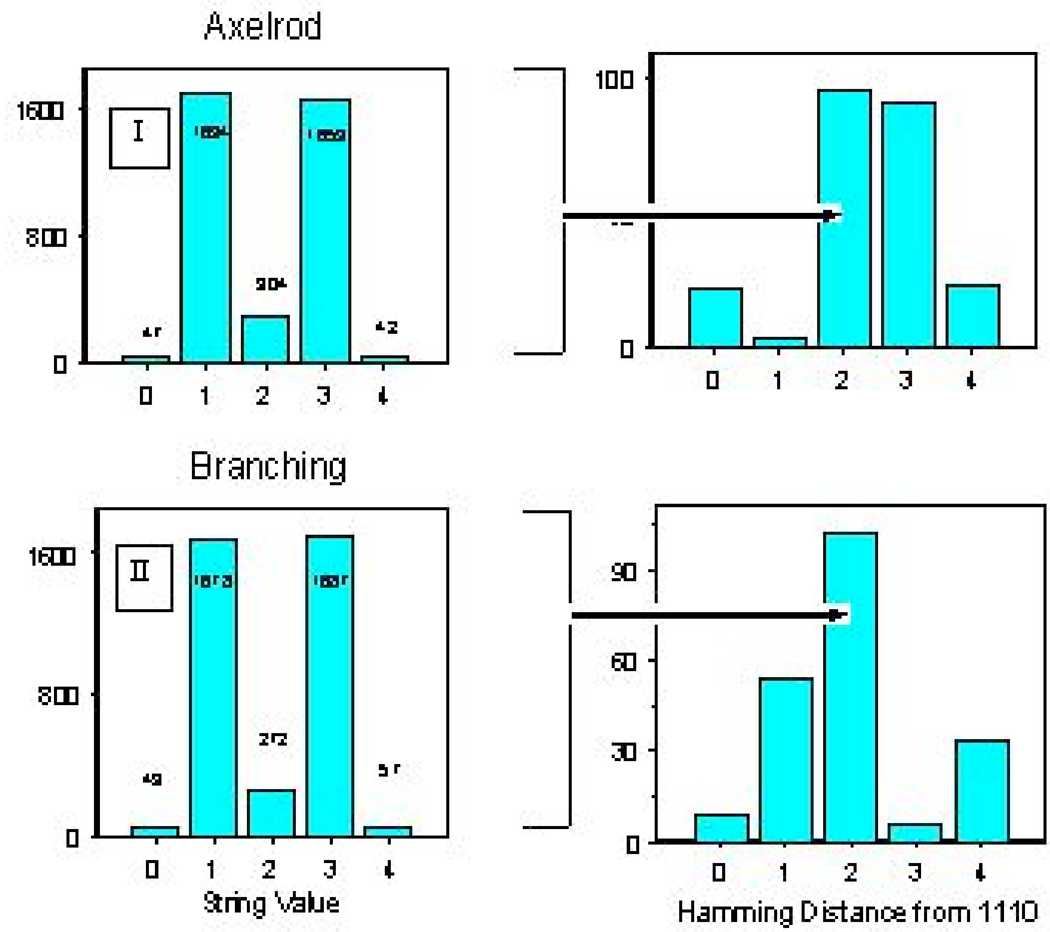

Figure 9: Observe final composite values of agents associated to the selected fickle zones. The top plots describe agents who move in/stay/move out of these fickle zones whereas the schematics in Figures 11 and 12 visualize only agents moving out of these hamming distances on the value spectrum. Both Axelrod and Branching algorithms consistently generated these distributions.

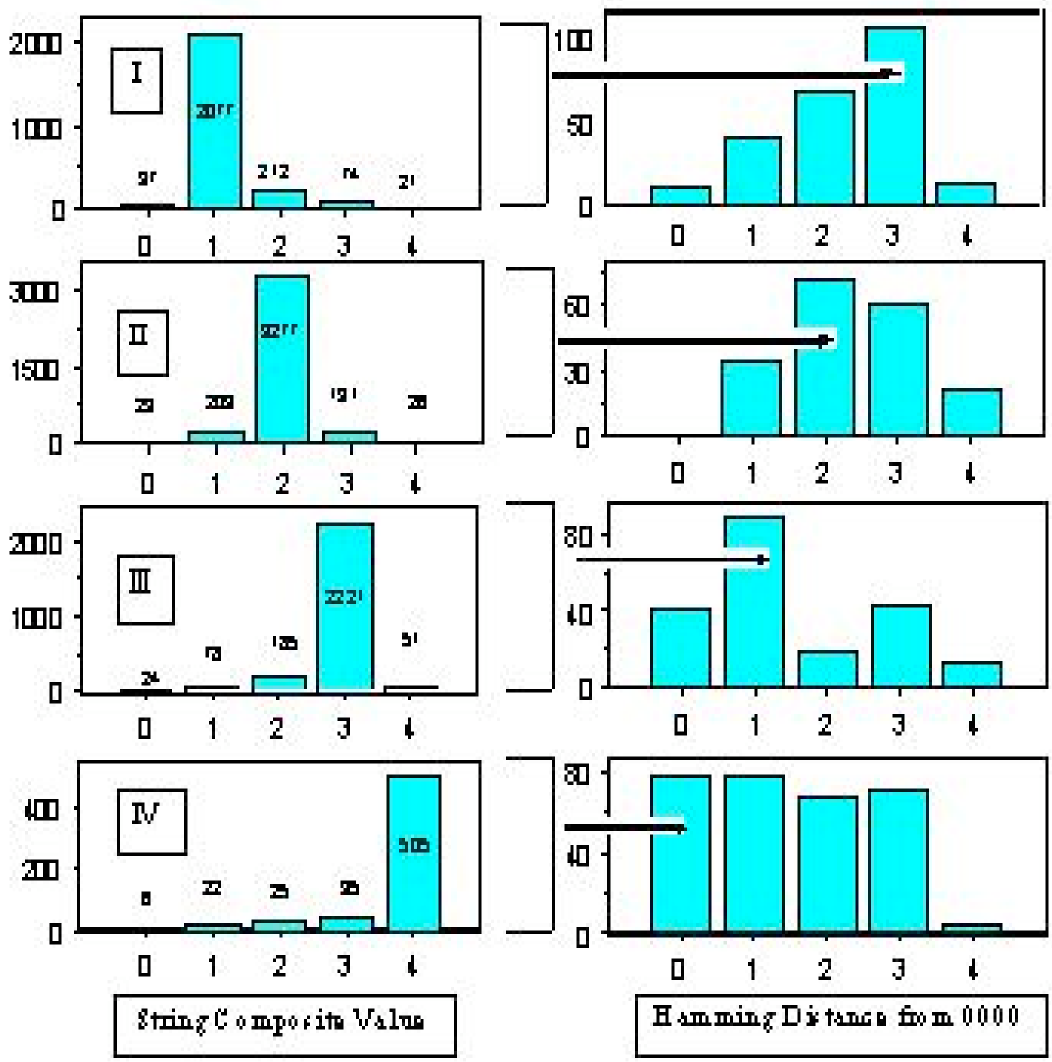

Figure 10: Observe final composite values of agents associated to the selected fickle zones. The top plots describe agents who move in/stay/move out of these fickle zones whereas the schematics in Figures 11 and 12 visualize only agents moving out of these hamming distances on the value spectrum. All plots taken from runs enacted by Axelrod's algorithm (simply a result of more rigorously exploring this algorithm).

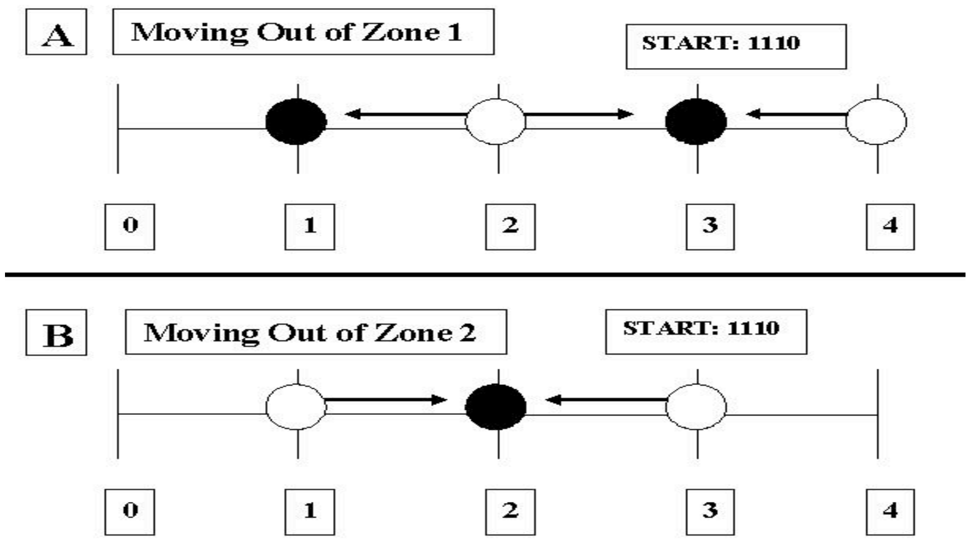

Figure 11: A. Agents moving out in scenarios I-III of Figure 8. B. Agents moving out in scenarios I and II of Figure 9.

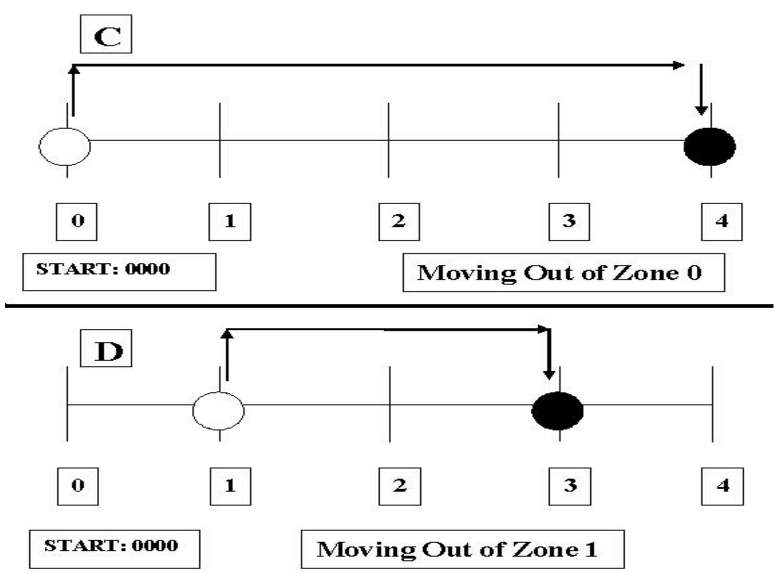

Figure 12: C. Agents moving out in scenario IV of Figure 10. D. Agents moving out in scenario III of Figure 10.

5 Conclusions

I summarize these results as:

- Agents in hamming zone 1 of 1110 move to values 1 and 3
- Agents in hamming zone 2 of 1110 converge to a moderate center
- Agents in hamming zone 0 of 0000 move to the opposite extreme on the spectrum
- Agents in hamming zone 1 of 0000 move to value 3, skipping a moderate center

Only the latter two are genuinely interesting but proof enough that a local view and the use of an artificial construction like a one-dimensional spectrum can afford counter-intuitive agent behavior. The question I have yet to explore is if this extremist radical behavior has real-world corollaries and, if not, determining if the root of the problem is Axelrod's algorithm or the spectrum as a useful construction.

5.1 Future Work

While I generated surprising results but faced inevitable time constraints I hope to address the following in the future:

1 Test the robustness of my results, as recommended by Eric Smith[9] in discussion. One method might be arbitrarily generating an alternative definition of opposite strings (i.e. 0100 and 1011) and observing how agents move along a composite value spectrum that positions these agents at its ends. Robust results would evidence similar counter-intuitive movements.

2 Run both algorithms longer to test when and if steady-states of stable homogenous regions appear (a test on Axelrod's model that is scaled far too down in order to test its applicability to large populations).

3 Incorporate some rational agent behavior where agents adopt bits as a function of how other agents in the population have engaged in cultural exchange [6]. Incorporating this might better address third parties since agents, in the last hours of an election, must weigh the risks/benefits of voting a particular way.

A Additional Notes

[1] So argues the claim of the 'wasted vote,' a notion that eclipses how essential third parties are in pressuring major parties to publicly address otherwise ignored issues. [**?**]. In this paper I do not judge third parties; rather, I only contend that some fraction of Nader's voters would have voted for Gore in the event of a two-candidate race.

[2] A group of anthropologists.

[3] As constructed, all strings of the same value behave the same on a political spectrum. Some might see this as a serious shortcoming as it eclipses the difference between, for example, strings 1110 and 0111. This can be resolved simply, though, by outputing a weighted composite value (a summation that factors in constant strengths unique to each bit).

[4] Intuitively one assumes evolution would preserve a normal distribution of composite values (see value spectrum for distribution); I introduce this question more to demonstrate the limits of any global analysis and to motivate a look into local dynamics.

[5] I measure evolution as changes in a population's hamming distance distribution from certain strings. This makes sense as hamming distance guides cultural exchange where I interpret a changing distribution of hamming distance as a population evolving toward or away from an agent.

[6] This is as much a test of the algorithms as a test if a simple one-dimensional spectrum can afford a better understanding of a system where local random dynamics seem to conceal any explanations of collective agent behavior.

[7] I concede that other zones are interesting, particularly since these changes are relatively few. For now I arbitrarily choose the maximum key movers as the site of interest (more in Q4). Future work on other zones will test the robustness of my findings in Q4.

[8] One can explain this relative frequency 3:1 by generating all possible strings that are within 1 hamming distance of 1110. Associate to these strings 0110, 1010, 1100, and 1111 composite values 2,2,2,4, and one understands the relative distribution of final values 2 and 4 as 3:1.

[9] Professor, Santa Fe Institute

References

[1] Adams, J., Clark, M., Ezrow, L., and G. Glasgow, 2006. Are Niche Parties Fundamentally Different from Mainstream Parties? The Causes and the Electoral Consequences of Western European Parties' Policy Shifts, 1976-1998. *American Journal of Political Science*,**50(3)**,513-529.

[2] Axelrod, R., 1997. The Dissemination of Culture: A Model with Local Convergence and Global Polarization. *The Journal of Conflict Resolution*,**41(2)**,203-226.

[3] Kollman, K., Miller, J.H., and S.E. Page, 1998. Political Parties and Electoral Landscapes. *British Journal of Political Science*,**28**,139-158.

[4] Rousseau, D., and A. M. van der Veen, 2005. The Emergence of a Shared Identity. *Journal of Conflict Resolution*,**49(5)**,686-712.

[5] Stokes, S.C., 1999. Political Parties and Democracy. *Annual Review of Political Science*,**2)**,243-267.

[6] Young, H.P., 2006. The Spread of Innovations by Social Learning. *Not printed*

[7] Wikipedia, the free encyclopedia. *third party (United States)*. July 11, 2006. http://en.wikipedia.org/wiki/Third_party_(United_States).