Introducing the idea of cultural coherence to Axelrod's model of cultural dissemination

Project report

Course: Self-Organisation, Cognition & Social Systems
Date of submission: 25/01/2010

Students: Johannes Ackva, Ali Honari & Tomas Zwinkels
Instructor: Professor Dr. Charlotte Hemelrijk
Institution: University of Groningen (RuG)

1. Introduction

In the following paper, we will present an extension to Axelrod's (1997) model of cultural dissemination that showed how processes of local cultural convergence need not lead to global convergence, indeed can lead to the exact opposite: stable global cultural polarization. A *coherence-bias* - the idea that certain *configurations of culture* are inherently more likely than others - will be argued for, its implementation be described and changes of model outcomes due to the modification will be reported and discussed. Briefly stated, there are two main reasons for this modification that will be explicated in more detail below: First, the idea underlying our coherence-bias captures important aspects of different cultural and cognitive theories. Furthermore, the modification also makes the original Axelrod-results - local convergence and global polarization - more robust to the introduction of randomness: While introducing any amount of random interaction or mutation to the original Axelrod model always resulted in global cultural convergence (San Miguel et al. 2005) we show that by introducing a coherence bias assumptions on the absence of random interactions need not be made to produce Axelrod's ‘basic result’.

The paper is structured as follows: First, Axelrod's model and existing extensions will be shortly introduced, as this work is our starting point. Second, the rationale underlying the coherence-bias will be presented, followed, third, by a description of its implementation. Fourth, the hypotheses, experimental conditions (parameter space explored) and measurements will be described, before - fifth - the results will be presented and - sixth - discussed.

2. Prior research

Axelrod's (1997) model of cultural dissemination, a cellular automaton model with patches having n *features* which can be any of m *traits* as defining their culture, serves as the starting point of our analysis. In the model, at each time step (tick), a random patch is activated and one of its neighbors (in the standard setting, 4 neighbors) is randomly chosen as its potential interaction partner.
The following algorithm describes the interaction process:

1. The degree of cultural similarity is determined. Cultural similarity is defined as the degree of identical traits for identical features.
2. If cultural similarity is neither maximal (completely identical) nor minimal (completely dissimilar), in one of the non-identical features (randomly determined) the activated patch

copies the trait of its neighbor.

This simple algorithm encapsulates *homophily* (interacting more with those similar) and *social influence* (becoming more similar to those you interact with) and thus leads, on the local level, to convergence whenever interaction is possible - and to stable borders whenever it is not.

On the global level, this mechanism can - depending on the exact combination of parameters (see section 6 for some examples) - lead to mono-culture (completely identical traits for all agents), but, more interestingly, also to global cultural polarization in the face of local convergence: A number of regions consisting of identical patches being completely dissimilar from other regions and thus resulting in stable equilibria of cultural divergent regions.

While producing this result out of the simple and rather plausible assumptions of the model was an interesting finding, it got dramatically weakened by San Miguel et al (2005). They observed that, when introducing any randomness into the model - either through randomly chosen interactions (i.e. not checking for similarity to determine whether interaction is possible) or through random mutation (change of traits of a patch independent of interaction), mono-culture was the only outcome, independent of other parameters. This finding resulted from the fact that borders between regions are, when randomness exists, never stable - eventually leading to all borders to disappear as the result of random events that made priory incompatible cultures open to interaction and consequently convergence.

Given that the absence of any random interaction or mutation as assumed by Axelrod is obviously a very implausible assumption, these results shed doubt on the usefulness of the Axelrod-model in explaining stable cultural polarization. Being able to explain global cultural polarization in the face of randomness, thus, would be an important improvement of the model - especially if it could be reached by a theoretically motivated modification instead of an ad hoc fix of the model. Discussing the reasoning for such a modification will be the subject of the following section.

3. Rationale for the coherence-bias

When discussing conceptions of culture, Axelrod (1997: 206) distinguishes two possibilities: (1) Viewing culture as a distinct set of traits, the cultural diffusion of each trait essentially being independent of the others and (2) treating culture as an "integrated package" in which the meaning of all traits is dependent on the others. He then goes on to claim that his model is novel in explicitly taking into account that "the effect of one cultural feature depends on the presence or absence of other cultural features" (Axelrod 1997: 207) - thus, being closer to the second view. While it is true that his model takes into account the amount of cultural similarity between two agents by influencing the likelihood and substance (becoming more similar on a still dissimilar feature) of their interaction, this is a *dyadic effect*, not an effect of the culture of any of the agents itself. Thus, contrary to Axelrod's claim, the effect of one cultural feature does *not* inherently depend on the presence or absence of others, but only so in dyadic relations where similarity matters. As he puts it (Axelrod 1997: 220):

> "The social influence model shows how homogeneous cultural regions can arise without any intrinsic relationship between the separate dimensions that become correlated"

While it is certainly true that co-variation of cultural traits need not imply that there is an intrinsic relationship between them, this does not imply the opposite - namely, that the co-occurrence of cultural traits is purely random. As, for example, cultural theorists claim (Thompson, Ellis & Wildavsky 1990: 2) the *continuous* co-occurrence of certain beliefs and values under evolutionary conditions - i.e. cultural mutations and selection pressures on cultural configurations[1] - is only plausible when these cultural configurations consist of traits that mutually reinforce each other and, taken together, form a coherent set of preferences and beliefs.

Similarly, in an attempt to develop a general formal model of preference and belief formation Chai (2001) coined his model "coherence-model", making the minimization of internal contradictions (and contradictions with prior actions) the central principle to base the selection and change of beliefs and preferences on - i.e. rooting it in observations on coherence and the minimization of cognitive dissonance stemming the cognitive sciences.

Empirical observations on socio-cognitive idea interaction dynamics - in the field of social psychology (Gray, 2002; 517-531) and more formal attempts to scrutinize human semantic memory and ‘belief processing’ in connectionist semantic memory models (McCleland, 2003) - generally show three important socio-cognitive ‘effects’ for which one could claim it to be important to include them in models of social belief dynamics:

---

[1] Given that always various sets of ideas compete for adherents, those that make less sense out of the actually observed world will eventually be weakened.

1.) Interaction generally leads to some form of homogenisation of ideas (social influence)
2.) Weakly embedded / less cognitively connected ideas are relatively more prone to change.
3.) Strongly embedded / highly cognitively connected ideas are relatively hard to change.

Thus there are good reasons to assume that in many cases co-occurring cultural traits have intrinsic relations to each other (one of the two only making sense if the other is present / absent) and when those traits co-occur, they mutually reinforce each other - making their disappearance less likely. This is the core theoretical idea behind the coherence bias. Giving this a formal algorithmic implementation in our model is the subject of the next section.

4. Implementing the coherence-bias

To introduce these effects into the classic Axelrod model - without having to deal with the computational (and time practical) burden of running connectionist algorithms that describe and model the cognitive relation between cultural feature values – we introduced a coherence bias. Where in the original Axelrod model after initiating the interacting patches have a similar change of changing *any* of the 'non shared' cultural feature values the coherence bias reduces the likelihood of features values to be 'flipped' to the value of the interaction partner whenever these feature values are 'internally cohered'[2], here addition the coherence bias is larger when more traits are cohered.

The vector that contains the probabilities that any specific feature value is flipped in our model is given

by:
When $C_i \neq 1$
and $p = (1 / o)$ when $C_i = 1$

Here

C id a vector containing a count of the amount of similar trait-values per trait-value
b id the parameter defining the coherence bias ($0 =< b$)
n is the amount of features
m is the amount of traits
o is the amount of features that is 'open for change' i.e. that contains a feature value that is different

---

[2] We considering features to be 'cohered' when the have the same feature value (for example index 1 and 3 in the vector 12134).

then the feature value on that position from the interaction partner.

For example:

The feature vector from:
agent A is [00123]’ and the feature vector from his interaction
partner is [21453]’
and b = 1 and n = 5 and m = 10

Then C (the vector containing a count of similar features) = [22111]
And o = 4 (notice that the last feature value from agent A is similar to the value of his interaction partner and thus four items are ‘open for change’).

The probability value for the first item will thus be:

The final probability vector will thus look like:
((1/16 1/16 1/4 1/4 1/4)’

Note that the trait flipping in this algorithm relates to cognitive property number one. In addition notice that the probability that non-cohered items (items 3,4 and 5) change is higher then the probability that the cohered items (item 1 and 2) change, in addition one can see that the probabilities do not need to sum up to 1 - providing a situation where cultural change might not occur – thus together including property two and three.

5. Hypotheses & research design

From the introduction of the *coherence bias* we expect several changes to the model outcomes: First, rather trivially, we expect that the *coherence bias* will make the emergence of more coherent cultural regions[3] more likely. This we will assess by measurements of one of our dependent variables - *global*

---

[3] It should be noted here that our measurement of number of regions slightly varies from that used by Axelrod (1997). In our measurement non-adjacent regions of the same culture are counted as 1 instead of 2 cultures, i.e. our measurement is biased towards counting less regions than the original Axelrod implementation did. In that sense, it does not weaken our below reported findings as, even with a slightly biased measurement towards less regions, we observe more regions. The reason for the difference from our implementation from Axelrod's is that Axelrod's model was not programmed in NetLogo

coherence - the averaged coherence over all patches[4]. Furthermore, we expect more regions to form - based on two mechanisms: As the attraction to coherence in same cases counteracts *social influence*, we expect more cultural regions to form. Consider the following example:

culture 1: 4 1 3 3 3
culture 2: 4 1 5 5 5

Here, in the original Axelrod-specification we would expect these two cultures to converge, already being identical in 40% of their features. The introduction of a *coherence bias* makes that less likely as there is an attraction to a rather pure 3-culture / 5-culture for culture 1 / culture 2, respectively. Related to this mechanism and working in the same direction is the expectation that attraction towards coherence will result in faster reached stable equilibria. As the time needed to reach equilibrium itself influences the number of regions (the longer it takes, the more time there is for converge, Axelrod 1997: 217-219), this strengthens our expectation that the coherence bias will increase the number of regions in the end. In sum, our hypotheses are the following:

H1: As soon as the coherence bias $b > n/m$ the number of regions will increase.
H1.1: The stronger the coherence bias, the more regions there will be.

H2: As soon as the coherence bias is $b > n/m$ global coherence will increase.
H2.1: The stronger the coherence bias, the higher global coherence will be.

To test these hypotheses, we employ the research design summarized in the following table (described in more detail in the results section) - holding size (10x10 patches) and termination (after 2 million ticks) constant:

---

and the Axelrod-model NetLogo version provided to us had this mistake which we only realized in the analysis.

[4] **Global coherence:** The average coherence over all patches. Averaged from the coherence of each single patch, the measurement varies in the same range - between 0 and 1.

<table>
<tr><th>Analysis</th><th colspan="6">Compared experimental conditions</th><th>Rationale</th></tr>
<tr><td colspan="8">A. The general impact of the coherence bias on basic results</td></tr>
<tr><td>Comparison 1</td><td colspan="3">Original Axelrod Model</td><td colspan="3">Axelrod replication</td><td>Replication of Axelrod (1997): Test equivalence</td></tr>
<tr><td>Comparison 2</td><td colspan="3">Axelrod replication</td><td colspan="3">Axelrod with random interaction</td><td>Replication of San Miguel et al. (2005): Test equivalence</td></tr>
<tr><td>Comparison 3</td><td colspan="3">Axelrod replication</td><td colspan="3">Axelrod with coherence</td><td>Test H1 and H2</td></tr>
<tr><td>Comparison 4</td><td colspan="3">Axelrod with random interaction</td><td colspan="3">Axelrod with random interaction and coherence</td><td>Test H1 and H2</td></tr>
<tr><td colspan="8">B. Varying the strength of the coherence bias</td></tr>
<tr><td>Varying the coherence bias in the original Axelrod replication (5F, 10T condition)</td><td>0</td><td>0.1</td><td>0.5</td><td>1</td><td>2</td><td>5</td><td>Test H1.1 and H2.1</td></tr>
<tr><td>Varying the coherence bias in Axelrod with random interaction (5F, 10T condition)</td><td>0</td><td>0.1</td><td>0.5</td><td>1</td><td>2</td><td>5</td><td>Test H1.1 and H2.1</td></tr>
<tr><td colspan="8">C. Varying the range of interaction</td></tr>
<tr><td>Varying the range of interaction (to 1) (5F, 10T condition)</td><td colspan="3">Axelrod replication</td><td colspan="3">Axelrod with coherence</td><td>Test H1 and H2</td></tr>
<tr><td>Varying the range of interaction (to 3) (5F, 10T condition)</td><td colspan="3">Axelrod replication</td><td colspan="3">Axelrod with coherence</td><td>Test H1 and H2</td></tr>
<tr><td colspan="8">D. Varying the degree of random mutation</td></tr>
<tr><td>No random mutation (5F, 10T condition)</td><td colspan="3">Axelrod replication</td><td colspan="3">Axelrod with coherence</td><td>Test H2</td></tr>
<tr><td>Random mutation of 0.001 (5F, 10T condition)</td><td colspan="3">Axelrod replication</td><td colspan="3">Axelrod with coherence</td><td>Test H2</td></tr>
</table>

Table 1: Research design

6. Results

The following table gives an overview over the basic results, while the text below describes them in more detail:

Analysis	**Compared experimental conditions**		**Rationale**
A. The general impact of the coherence bias on basic results			
Comparison 1	Original Axelrod Model	Axelrod replication	Replication of Axelrod (1997) successful: Results very similar
Comparison 2	Axelrod replication	Axelrod with random interaction	Results indeed replicate findings from San Miguel et al, 2005 : Random interaction and muation always leads to mono-culture
Comparison 3	Axelrod replication	Axelrod with coherence	**H1:** More regions in all coherence-conditions No mono-culture anymore. Except for the 5 Trait / 15 Features condition (more regions than expected), the relative pattern between the different parameter combinations is identical for the coherence-conditions **H2:** Higher global coherence in the coherence-conditions
Comparison 4	Axelrod with random interaction	Axelrod with random interaction and coherence	**H1:** The randomness does not lead to mono-culture, but instead around 5 cultures persist in each condition There is no clear pattern among the different parameter conditions in the coherence-

							conditions (more model runs would be needed to establish the lack of (statistically) significant differences) **H2:** Higher global coherence in the coherence-conditions
B. Varying the strength of the coherence bias							
Varying the coherence bias in the original Axelrod replication (5F, 10T condition)	0	0.1	0.5	1	2	5	**H1.1** Global coherence does not increase as soon as coherence-bias >= 0.5 **H2.1:** Number of regions does not increase as soon as coherence-bias >= 0.5
Varying the coherence bias in Axelrod with random interaction (5F, 10T condition)	0	0.1	0.5	1	2	5	**H1.1** Global coherence does not increase as soon as coherence-bias >= 0.5 **H2.1:** Number of regions does increase with increasing coherence-bias
C. Varying the range of interaction							
Varying the range of interaction (to 1) (5F, 10T condition)	Axelrod replication			Axelrod with coherence			**H2: Confirmed**

Varying the range of interaction (to 3) (5F, 10T condition)	Axelrod replication	Axelrod with coherence	**H2:** Confirmed
D. Varying the degree of random mutation			
No random mutation (5F, 10T condition)	Axelrod replication	Axelrod with coherence	**H2:** Confirmed
Random mutation of 0.001 (5F, 10T condition)	Axelrod replication	Axelrod with coherence	**H2:** Confirmed

Table 2: Summary of basic results

The replication of Axelrod's original experimental design with regards to variation in number of features and traits (A in the table above) yields almost identical results - making us confident that our slight deviation from the original model - different region-counting and point of termination - do not significantly alter our implementation from the original model. It is especially noteworthy to point out that all conditions that produced mono-culture in Axelrod's original model led to the same result in our implementation.

The results are reported in detail in the following table:

(we / Axelrod)	**Amount of Traits**		
Amount of features	**5**	**10**	**15**
5	1 & 1	3,4 & 4	18,7 & 20
10	1 & 1	1 & 1	1,6 & 1,4
15	1 & 1	1 & 1	1 & 1,2

Table 3: Comparison of Axelrod (1997: 212) with our data on the original Axelrod model

The same holds for our replication of San Miguel et al. (2005) results on the introduction of random interaction. Setting the parameter of random interaction to 1%, each of the conditions yields mono-culture in our as well as their implementation.

To test H1 and H2 we now first compare the original Axelrod model (coherence bias = 0) to an otherwise identical model with a moderate coherence-bias (coherence bias = 1) over the range of 9 different combinations of features and traits (as reported in Axelrod 1997: 212), the results are reported in the following tables:

Coherence bias	**Amount of Traits**		
Amount of features	**5**	**10**	**15**
5	5,8	17	35,1
10	5,1	6,8	9
15	5,6	4,5	5,5

(Axelrod - no coherence bias)	**Amount of Traits**		
Amount of features	**5**	**10**	**15**
5	1	3.4	18.7
10	1	1	1.6
15	1	1	1

Tables 4 & 5: Comparing the original Axelrod model with moderate coherence-bias model

With the exception of the 5 traits / 15 features condition, the results confirm our expectations - in general the relative sizes of the different parameter conditions are similar to the original model and, consistently over all conditions, the introduction of the coherence-bias leads to more stable regions - with mono-culture not occurring at all anymore. Almost needless to say, the global coherence values are also much higher in the moderate coherence bias conditions:

Global Coherence org, Axelrod	**Amount of Traits**			
Amount of features	**5**	**10**	**15**	
5	0,34	0,31	0,12	
10	0,26	0,19	0,14	**average**
15	0,24	0,15	0,12	**0,21**

Global Coherence with bias	**Amount of Traits**			
Amount of features	**5**	**10**	**15**	
5	0,96	0,76	0,53	
10	0,93	0,76	0,67	**average**
15	0,93	0,79	0,58	**0,77**

Tables 6 & 7: Global coherence for the conditions reported in Tables 4&5

Concluding the general comparisons, we found that - when looking at a model with 1% random interaction - H1 and H2 are again confirmed: While random interaction of 1% in the original case always led to mono-culture, this does not happen at all in the conditions with a moderate coherence-bias (= 1) - with an average of 5 stable regions in the end, without clear patterns among the nine different parameter combinations.

Turning to the impact of changing the size of the coherence bias (B in the table above), we find that in the case without random interaction, the size of the coherence bias does not increase the number of regions anymore as long as it is >= 0.5. This finding suggests that the moderate coherence bias we used above (= 1) was already sufficient to show the full working of the coherence attraction, with a stronger bias not making a noteworthy difference anymore. It also leads to the partial dismissal of H2.1 for the case of no random interaction: While, at low values an increase of the coherence-bias leads to an increase of regions, this effect levels off at 0.5 - with an additional increase having no additional effects and a similar finding for the level of global coherence.

Figure 1: The impact of the size of the coherence bias on the number of regions

Coherence bias	**0,00**	**0,10**	**0,50**	**1,00**	**2,00**	**5,00**
Regions	5,80	3,60	17,20	16,80	17,60	18,00
Global Coherence	0,27	0,25	0,69	0,72	0,75	0,75

Table 8: The impact of the size of the coherence bias on global coherence

Interestingly, these observations are different for the case of random interaction (1%): Here, increases of the coherence-bias show an increasing effect on number of regions over the whole range of values. This might be explained by random interaction acting as a counter-force to the stabilization of the system and, with more ticks needed to lead to stable regions, more time for the coherence-bias to work.

In that case, the leveling off of the effect of the coherence bias probably takes place at a higher absolute value which has not been covered by our parameter range. It is also interesting to note that, as in the prior case, global coherence does not change much after the coherence-bias has reached 0.5, despite the region patterns diverging in both cases - i.e. global coherence and number of regions seem to be, to some degree, unrelated in the latter case.

Figure 2: The impact of the size of the coherence bias on the number of regions (with random int.)

Coherence bias	**0,00**	**0,10**	**0,50**	**1,00**	**2,00**	**5,00**
Regions	1,00	1,00	2,80	3,00	5,00	7,60
Global Coherence	0,25	0,26	0,98	0,99	0,96	0,94

Table 9: The impact of the size of the coherence bias on global coherence (with random int.)

For C & D, interactions of coherence-bias and range of interaction / random mutation, respectively, all results are as expected and thus only briefly reported here: As expected, increased range of interaction (leading to fewer regions) and a moderate coherence-bias (leading to more regions) counter-act, leading to the combined condition (increased range of interaction, moderate coherence-bias) lying between the highest number of regions (increased range of interaction, no coherence) and the lowest one (normal range of interaction, coherence). Introducing mutation of 0.0001 (on average every thousand ticks on of the n feature values from a random patch is changed to another value) without a coherence bias leads to mono-culture. We have seen earlier that the presence of a coherence bias in the case of random interaction ‘saved’ the model from ending up in mono-culture. Again; while random mutation 0.0001 in the original case always led to mono-culture, this does not happen at all in the conditions with a moderate coherence-bias (= 1) - with an average of 9.8 stable regions in the end.

Regions	**Range**	
Coherence bias	**1**	**3**
0	3	1
1	19	9.8

Global Coherence	**Range**	
Coherence bias	**1**	**3**
0	0.28	0.30
1	0.79	0.88

Table 10 & 11: Results of varying range of interaction

Regions	**mutation**	
Coherence bias	**0**	**0.0001**
0	4.8	1.4
1	17.8	6.5

Global Coherence	**mutation**	
Coherence bias	**0**	**0.0001**
0	0.26	0.24
1	0.72	0.97

Table 12 & 13: Results of introducing mutation

7. Discussion & Conclusion

As apparent from the reported results the introduction of the coherence-bias changes the outcomes of the model significantly - generally, increasing the number of regions and, quite trivially, the level of global coherence. What is especially noteworthy is that mono-culture does not occur anymore as soon as any random interaction is introduced, which is a strength of the model modification as this result was an implausible one of the original model. While this was a positive by-product of our model modifications, this was not the initial purpose - rather, we aimed at incorporating a theoretical concept we believed to be of great importance - cultural coherence - to the original Axelrod model. As some of the open questions addressed in the results section indicate, much more could be analyzed regarding the model - to better understand surprising findings, clarify mechanisms and look at the impact of additional parameters (such as size, higher random interaction etc...) and, when simulating more cases, statistical analysis would be possible, too. Doing this, however, will require another paper.

8. References

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

Chai, S.K. 2001. *Choosing an Identity: A General Model of Belief and Preference Formation.* University of Michigan Press.

Gray. (2002) *Psychology: Fourth edition.* Worth Publishers: New York.

McClelland, J. L. and Rogers, T. T. (2003). The Parallel Distributed Processing Approach to Semantic Cognition. *Nature Reviews Neuroscience*, 4, 310-322.

Miguel MS, Eguiluz VM, Toral R, Klemm K (2005). Binary and multivariate Stochastic models of consensus formation. *Computing in Science and Engineering; 6-7; p* 67-73

Thompson, M., Ellis, R. & Wildavsky, A. 1990. *Cultural Theory*. Boulder: Westview Press