

An Agent-Based Model of Context-Dependent Norm Emergence in Language Contact

Agent-based modeling is a computationally powerful tool of analysis which has recently been used in the social sciences to study aspects of community dynamics, including language change and choice. Model systems are computer simulations consisting of interactions among entities or programs called ‘agents,’ and rules governing the behavior of the agents (Gilbert 2008). The present study investigates linguistic norm emergence in a heterogeneous community with context-stratified interactions. Results indicate that the maintenance of context-stratified norms and dynamic equilibrium among variants is contingent on both locality constraints on interaction, critical period effects on learning, and the nature of the agents’ updating mechanism.

Background and Previous Work

An early application of agent-based models to language which has influenced much subsequent research is the model of Nettle (1998). Nettle’s model consisted of a grid of agents choosing between one of two variants by weighting the ‘social impact’ of each variant according to a function of its frequency of use in their neighborhood. In this approach, as for most subsequent studies, variant choice is discrete and not probabilistic: agents choose one variant over another if it has a higher impact value, and while agents may have a bias for one variant, the only other factors determining this choice are the number and proximity of other agents using the variant. Nettle found that manipulating the exponential parameters governing the weighting of these factors had significant effects on the model. In particular, a relatively strong weighting of social distance was necessary to produce community-wide spread of the ‘new’ variant after it was introduced into the model. As regards other factors determining variant spread, he notes that “with no status differentials between people and no functional biases, it is extremely hard to see how rare variants ever overcome the threshold problem and spread.” Nettle then proceeds to explore the question of status differentials by introducing hyperinfluential individuals into the simulation, and addresses the issue of ‘functional biases’ in terms of arbitrary bias towards one variant or the other. His results indicated that hyperinfluential individuals are particularly important in promoting the spread of a new variant.

A more sophisticated extension of Nettle’s norm emergence paradigm is represented by Fagyal et al. (2010). Unlike Nettle’s model, which used locality to indicate network distance, Fagyal et al. generated network links using a randomized and scale-free algorithm. Also, rather than the social impact function used by Nettle, Fagyal et al. modeled variant choice by an ‘indegree-biased voting model’ depending on network distance and connectivity (influence) of neighboring agents. Their model considered a multiplicity of variants, and found that not only were highly-connected, ‘hyperinfluential’ individuals essential to majority variant establishment, but that ‘loners’ with very few network links functioned as reservoirs of minority variants which could later be reintroduced into the broader population, and thus blocked long-term dominance of a single variant. This study highlighted the importance to community language dynamics of the ‘small-world network’, a network having the properties of a small diameter and high clustering.

Two more recent agent-based modeling studies informed the specific approach of the present study. Most important was the model of Mühlenbernd and Quinley (2012) simulating the development of register-differentiated doublets in Middle English associated with French

influence. This study illustrated the applicability of a game-based approach to community language dynamics, rather than the voting-style models used by Nettle and Fagyal et al., and added the important dimensions of status and context differentiation to the analysis. Modeling variant selection as the result of a context signaling game, this study concluded that context-dependent meanings could be constructed from interaction across contextually asymmetric groups, but that breakdown of context-dependence resulted from the removal of status differentials between the groups. The present model draws on the concept of context-dependent updating used by Mühlenbernd and Quinley, but simplifies it by removing the dimension of pure status and rather treating context assignment as a probabilistic function of group membership.

Finally, the work of Stanford and Kenny (2013) illustrates use of agent-based models to simulate specific empirical situations, rather than behavior in abstract theoretical spaces. Here, the authors modeled the spread of the Northern Cities Shift to the St. Louis Corridor in a coordinate space reproducing the geography of the relevant area, with a number of agents scaled to match proportions of actual population density. Agents produced vowels varying along a single dimension, and travelled back and forth in varying proportions between ‘Chicago’ and ‘St. Louis.’ By modulating learning parameters and structural properties of the simulated vowel space, the model could study patterns of diffusion of the northern vowel shift to ‘St. Louis’. A significant departure of Stanford and Kenny from the earlier approaches mentioned above is the use of an exemplar mechanism to update the agents’ behavior, and I consider the implications of such a mechanism for my own model. However, it should be noted that the present study falls more into the category of ‘abstract models’ than ‘facsimile models’ such as Stanford and Kenny.¹

The model developed here considers the dynamics of context-differentiated norm emergence under different agent learning mechanisms and population conditions. The basic question I wish to answer is how stable register differentiation emerges in a heterogeneous society, and what aspects of learning and population structure influence the outcome of contact between linguistically distinct groups. The facts suggest that register and style differentiation are universal aspects of language use, and that prestige languages often emerge in multilingual situations. The classical diglossia of Ferguson (1959) is a prime example of this phenomenon, as is the style-shifting observed in sociolinguistics (Labov 2001). I develop a model which creates initial style-stratification through group asymmetries in context distribution, and allows agents to develop context-specific variant frequencies through interaction and learning. This approach draws on the insight of Bell (1984) that linguistic styles are constructed by ‘audience design’ towards the language of a typical speaker in a given context.

The model was designed and run in the agent-based modeling environment NetLogo. I make use of the fundamental framework of the Quinley and Mühlenbernd model in a simpler interaction process, where context assignment is directly related to group membership without the intermediate stage of status. Furthermore, my model contrasts the signaling-game approach to context-dependent updating used in that model with exemplar-based alternatives such as were

¹ For a discussion of this distinction and its significance, see Gilbert (2008).

² In the experimental dataset considered below, which consisted of 360 runs of 5000 steps each, the actual observed maximum was 157 agents.

³ This specific threshold value is undoubtedly an artifact of the updating weight of 0.1 and the fact that the experimental initial $P(g_1)$ values are clustered around 0.25, 0.5, and 0.75 due to the setup conditions.

used by Stanford and Kenny, and considers the implications of requiring a ‘critical period’ for agent learning.

Mechanism of the Model

The model is set up with 100 agents in a toroidal two-dimensional coordinate space having dimensions of 32 by 32 units in the model environment, where each unit is defined as an agent’s locality. Agents are divided into two groups g_1 and g_2 , in proportions specified during setup as $P(g_1)$. Members of each group are normally distributed over the three interaction contexts C_A , C_B , and C_C at each run of the model, such that members of g_1 are most likely to be assigned to C_A and least likely to be assigned to C_C at a given run, and members of g_2 are most likely to be assigned to C_C and least likely to be assigned to C_A . Table 1 gives the probabilities of being assigned to each context for members of each group.

Group	$P(C_A)$	$P(C_B)$	$P(C_C)$
g_1	68.2%	27.2%	4.6%
g_2	4.6%	27.2%	68.2%

Table 1. Distribution over contexts by group

Agents coordinate with each other using a set of two variants v_1 and v_2 , which are restricted to g_1 and g_2 respectively in the initial conditions, and afterwards spread through the population based on the result of interactions. Each step of the model after setup consists of five stages. Agents first (1) find partners, then (2) coordinate with them, and (3) update their likelihood of using variant 1 or 2 based on the result of the interaction. The final two stages are optional, and consist of (4) moving and (5) reproducing and dying.

1. Partner Selection

Agents select coordination partners in the model by picking out another agent at random having the same interaction context C_A , C_B , or C_C . There are two main variants of this mechanism, the ‘unrestricted’ selection method under which agents may choose a partner from any location in the population space, and the ‘local’ selection method under which agents may only choose a partner within their locality, defined as a distance of one unit in the coordinate plane. After partners are chosen, a diagnostic is run to ensure that partners are mutual. Agents that pass this test, typically between 20% and 60% of the total population, move on to the coordination stage.

2. Coordination

The coordination stage consists of an interaction in which agents send a signal of v_1 or v_2 to their partner, and receive a similar signal in return. The choice of variant is calculated from information stored in a ‘grammar’ sensitive to the context of interaction. This ‘grammar’, encoded in the model as a list of exemplars or probabilities for each of the three contexts, is

generated by the update mechanism in every interaction subsequent to the first. For the initial run, all g_1 agents signal v_1 in all contexts, and all g_2 agents signal v_2 .

3. Updating

The updating process is the core of the model, and governs the method by which agents modify their grammars after coordination. There are three distinct methods.

First, the ‘payoff-sensitive’ update method treats coordination as a game with payoffs, which the agents refer to in order to modify future behavior. If coordination succeeded (the variant signaled was the same as the variant received), agents become 10% more likely to use the same variant at the next round of coordination within that context. If coordination failed, however, and a different variant was received than was sent, agents become 10% less likely to use the same variant at the next round of play. The probability of using a different variant at the next round of coordination is simply the inverse. More formally, at any period of interaction number t in the relevant context,

$$p_{t+1} = \begin{cases} p_t + 0.1 & (v_t = v'_t) \\ p_t - 0.1 & (v_t \neq v'_t) \end{cases}$$

and

$$q_{t+1} = 1 - p_{t+1}$$

where p_{t+1} is the probability that the agent will choose the same variant at the period of interaction following period t as was chosen at period t with probability p_t , q_{t+1} is the probability that the agent will choose the variant not chosen at period t at the next period of interaction, v_t is the variant chosen by the agent at period t and v'_t is the variant chosen by the agent’s partner at period t .

The ‘exemplar’ update method uses an entirely different procedure, according to which an agent stores the variant received by its partner in the most recent interaction in a given context, and chooses that same variant as its signal for the next round of coordination in that context.

Thus, if an agent in context C_A receives the signal v_1 from its partner, then the next time the agent interacts in C_A it will send the signal v_1 . Under this method, agents store very little information themselves, and more complex behavior is a function of the population as a whole.

The ‘memory-averaging’ update method combines aspects of the exemplar method and the payoff-sensitive method by being both probabilistic and memory-driven. In this case, agents store a memory of the previous 50 interactions in each context (memories are initially populated with v_1 for agents in g_1 and v_2 for agents in g_2) and choose between v_1 or v_2 in interactions according to the frequency of each variant in the memory. Varying the length of the memory was not found to affect the behavior of the model under this condition.

4. Motion

Agents may be allowed to move one unit in a random direction within the model space after the interaction and updating stages. If motion is selected, it may be obligatory for all agents or obligatory contingent on failure to interact in the current run. In the latter case, only those agents which did not have a mutual partner in the current play move one patch in a random direction. It

was found that causing all agents to move during each run was extremely similar in its effects to having no locality restriction on partner selection, so in the following discussion only the more restricted motion process will be discussed. It should also be noted that in the case that locality restrictions on interaction are not present, the geography of the model space is completely irrelevant to the working of the model, and motion is entirely superfluous. Thus, it is only for local partner selection that motion will be considered as a factor.

5. Generational Dynamics

The final optional stage when running the model is to introduce demographic dynamics by recording agents' age and causing them to generate new agents and die. This option has the primary function of allowing for age-specific learning patterns, such that agents do not update if their age is greater than 100 steps. In order to maintain population stability, death may only occur after an agent has generated a single offspring. This means that the maximum possible population is 200, twice the initial numbers, and death probabilities are set so that typically just over 130 agents coexist simultaneously.² Group proportions are also guaranteed to remain stable over multiple generations.

Conditions and Outcomes of the Model

Four kinds of long-term behaviors were found to be generated by the model : (1) emergence of stable context-stratified frequencies of variant use, (2) emergence of stable frequencies with no context stratification, (3) population shift to a single variant in each context, and (4) complete shift to a single variant. Different combinations of parameters favor different outcomes. In particular, (2) is the only outcome observed under the memory-averaging update method when there are no restrictions on learning, and (3) is uniquely associated with the exemplar update method when there are no restrictions on partner selection or learning.

Outcomes (1) and (4) may emerge under a wider variety of conditions. (1) is associated with payoff-sensitive updating under restricted conditions of partner selection or learning, and with memory-averaging updating under critical period learning restrictions. (4) is the only stable long-term outcome of payoff-sensitive updating with no interaction or learning restrictions. Either (1) or (4) may also emerge from exemplar updating under the same conditions, with the difference that in the exemplar case the distribution of variant frequencies over contexts in (1) is not expected to correspond to the distribution of group frequencies over contexts, and that (4) emerges only as a special case of (3), under identical conditions.

The following discussion considers model behavior for each of these cases in more detail, using data from an initial experimental set of 360 runs of the model, 5 for each relevant condition, stopping in each case after 5000 stages. Table 2 below shows the set of parameter conditions tested. Conditions combining no motion and local updating were not considered for analysis, since under these conditions interaction is infrequent (typically less than 10 agents per step) and the majority of the population remains static. Also, note that motion conditions are irrelevant to the behavior of the model under the unrestricted partner selection method.

² In the experimental dataset considered below, which consisted of 360 runs of 5000 steps each, the actual observed maximum was 157 agents.

		Critical Period		No Critical Period	
		Restricted Motion	No Motion	Restricted Motion	No Motion
Payoff-Dependent Updating	Local Interaction	Context-Dependent Frequencies*	N/A	Context-Dependent Frequencies	N/A
	Unrestricted Interaction	Context-Dependent Frequencies		Complete Replacement	
Memory-Averaging Updating	Local Interaction	Context-Dependent Frequencies*	N/A	Group Proportion Matching	N/A
	Unrestricted Interaction	Context-Dependent Frequencies		Group Proportion Matching	
Exemplar Updating	Local Interaction	Context-Dependent Frequencies	N/A	Context-Dependent Frequencies*	N/A
	Unrestricted Interaction	Context-Dependent Frequencies*		Context-Dependent Replacement	

*These cases are defective in being excessively static or random, as discussed below

Table 2. Conditions tested by the BehaviorSpace experiment and associated outcomes.

1. Model Behavior in Unrestricted Conditions

Given entirely unrestricted conditions of interaction and no critical period, each update mechanism is associated with a characteristic outcome. Payoff-sensitive updating results in complete replacement of one variant by another, exemplar updating results in context-dependent replacement, and memory-averaging updating results in rapid levelling to initial frequencies across all contexts.

Figure 1 illustrates the replacement resulting from payoff-sensitive updating. Replacement is strictly in the direction of the variant associated with the majority group, even when the difference is very small, as it is in this example. The top graph is a run having initial proportions of 46% g_1 , and the bottom graph is for a run with initial proportions of 53% g_1 . Replacement is also fairly rapid, reaching completion in all contexts well before 1000 steps of the model (approximately 5 ‘generations’).

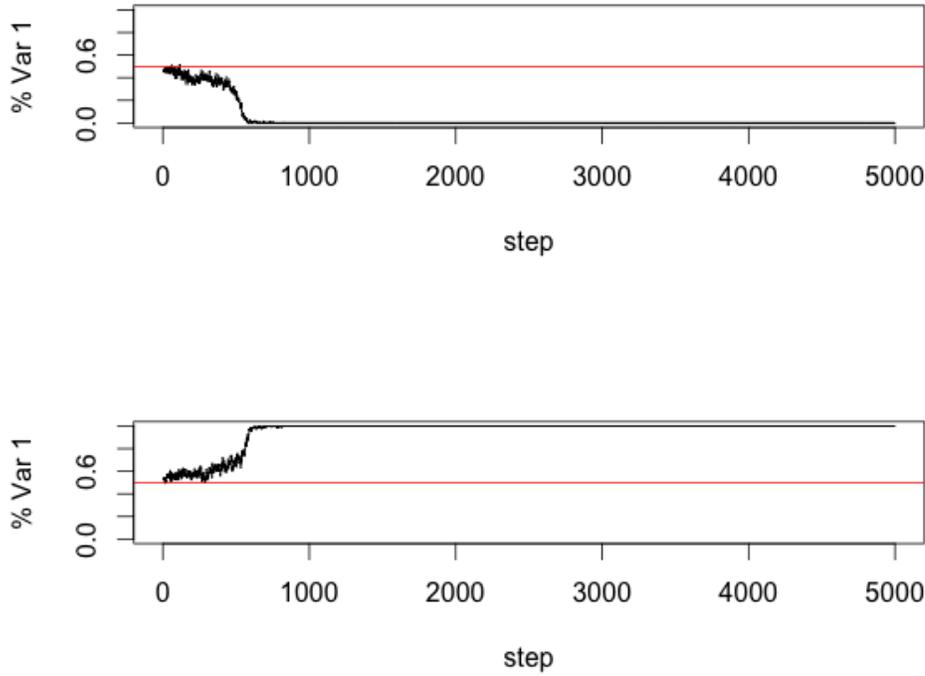


Figure 1. Variant replacement in payoff-sensitive updating with no restrictions.

Exemplar updating under unrestricted contexts also results in replacement, but here the replacement is context-dependent, rather than convergence to a single variant across all contexts. Complete convergence does occur, but only as a subcase of context-dependent convergence. There is a significant random element to the direction of convergence within contexts, but statistical trends within contexts do emerge, as shown by the summary of our results given in Table 3. Over 30 runs, 17 had initial $P(g_1) < 0.5$ and 13 had initial $P(g_1) > 0.5$. In all runs, replacement was complete in C_A and C_C by 5000 steps, with a majority of replacements being in the direction of the majority group's variant for C_A . C_B showed a similar pattern to C_A in tending to be replaced by the majority group's variant, but additionally failed to undergo replacement in a handful of cases. This is probably because groups are evenly distributed over C_B , so the replacement process went to completion more slowly and had not yet reached completion by 5000 steps. For C_C , the results are harder to interpret, since this context showed a tendency to replace with v_1 in all cases (though more strongly when $P(g_1) > 0.5$) despite the predominance of g_2 in this context. This may be an indication that the choice of predominant variant in the exemplar-updating mechanism is too subject to random factors for the tendencies described here to be significant.

	$P(g_1) < 0.5$			$P(g_1) > 0.5$			
	v_1	v_2	neither	v_1	v_2	neither	
C_A	3	14	0	C_A	9	4	0
C_B	4	10	3	C_B	9	2	2
C_C	10	7	0	C_C	11	2	0

Table 3. Counts of replacements with given variants by context over 30 runs with exemplar updating in unrestricted conditions.

Using the memory-averaging update method, complete replacement never occurs in any context. Instead, under unrestricted conditions, frequencies of v_1 and v_2 rapidly converge towards initial group proportions within each context, resulting in a complete lack of context- or group-dependent differentiation. Figure 2 plots the distribution of overall v_1 frequencies over the last 100 steps against initial g_1 proportions for all 30 runs with these conditions, comparing the distribution with a line of complete equivalence. It is clear that there is a significant tendency for variant proportions to be less uneven than the initial group proportions, but in no case is the deviation from initial proportions greater than 13.6%.

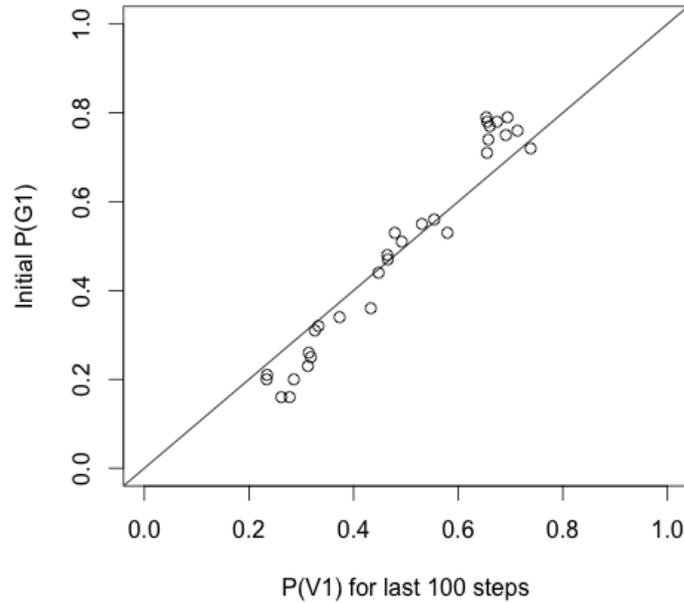


Figure 2. Final $P(v_1)$ averaged over 100 steps compared to initial $P(g_1)$, and the line $y = x$ for comparison.

The breakdown of context differentiation is even more complete, with the largest total difference between context-specific proportions over the last 100 stages of a run being 5%. Figure 3 charts a typical trajectory for this case, with initial $P(g_1) = 0.21$. Here the top graph plots overall frequencies of v_1 , the middle graph plots frequency of v_1 among members of g_1 , and the bottom graph plots the ratio of v_1 frequencies in C_A to v_1 frequencies in C_C on a log scale. The population converges to an overall $P(v_1) = 0.23$, $P(v_1|C_A) = 0.23$, $P(v_1|C_B) = 0.23$, and $P(v_1|C_C) = 0.24$. Establishment of these equilibrium conditions is relatively rapid, with frequencies stabilizing well before 1000 steps.

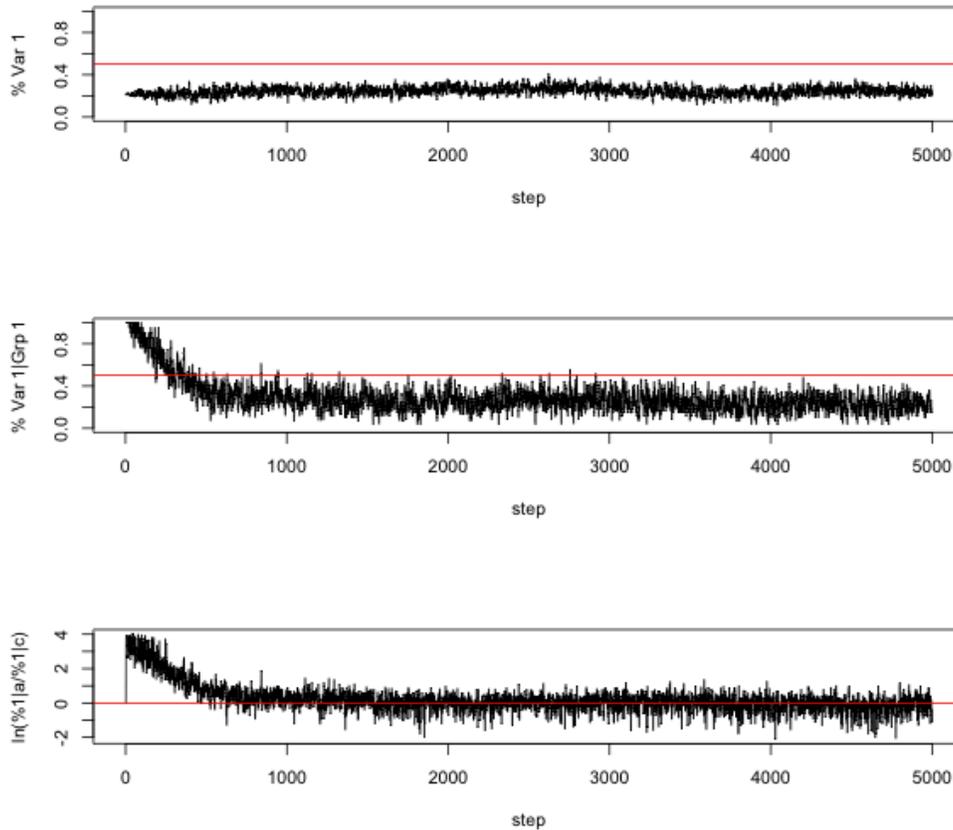


Figure 3. Convergence of variant frequencies to match group proportions with unrestricted memory-averaging updating. Measurements are compared to a baseline (in red) of 0.5 for the proportion plots and of 0 for the log-scaled ratio plot.

2. Model Behavior with Locality Restrictions

When locality restrictions on interaction are added to the model using the local partner-choosing algorithm and limited motion, it behaves very differently. For payoff-dependent and exemplar updating, stable context-dependent frequencies emerge as the long-term outcome, while the memory-averaging algorithm continues to converge towards a single group proportion at a slower rate. This ‘slowing down’ of the model is partly or entirely due to the fact that instead of approximately 50% of agents interacting at each step, only about 10% do so under locality restrictions. This is essentially an artifact of the model design. However, we have observed that in the unrestricted case equilibrium states are reached before 1000 steps, so we may expect that in the restricted case the time for equilibrium may be increased by a factor of 5 to 5000 steps, and may trust any pattern emerging by the end of this period as reflecting a stable outcome.

Figure 4 illustrates the context-dependent frequency result arising from locality restrictions under payoff-sensitive updating, with initial $P(g_1)$ of 0.24. These graphs give the same statistics as in Figure 3: overall v_1 frequency, frequency of v_1 in g_1 , and the log-transformed ratio of v_1 frequencies in C_A to v_1 frequencies in C_C . This last is an important indicator of context stratification of frequencies, since the strictly positive value indicates that v_1

is more frequent in C_A than in C_C . Context stratification in the direction of initial distributions, such that $P(v_1|C_A) > P(v_1|C_B) > P(v_1|C_C)$, was observed for all 15 runs of the model under these conditions. It was also the case that $P(v_1|C_C)$ was vanishingly small (< 0.02) for g_2 majorities greater than 70%, as was $P(v_2|C_A)$ for g_1 majorities greater than 70%.³

The initial group-dependent associations of the variants are weakened, since we see $P(v_1|g_1)$ moving from 100%, the initial condition, to approximately 50%. However, the equivalent plot of $P(v_1|g_2)$, shown as the green line in the middle graph of Figure 4, indicates that there has been no adoption of v_1 by this group. In fact, the majority group uniformly exhibits over 90% use of its associated variant, while the minority group shifts towards use of both variants, as illustrated by Figure 5.

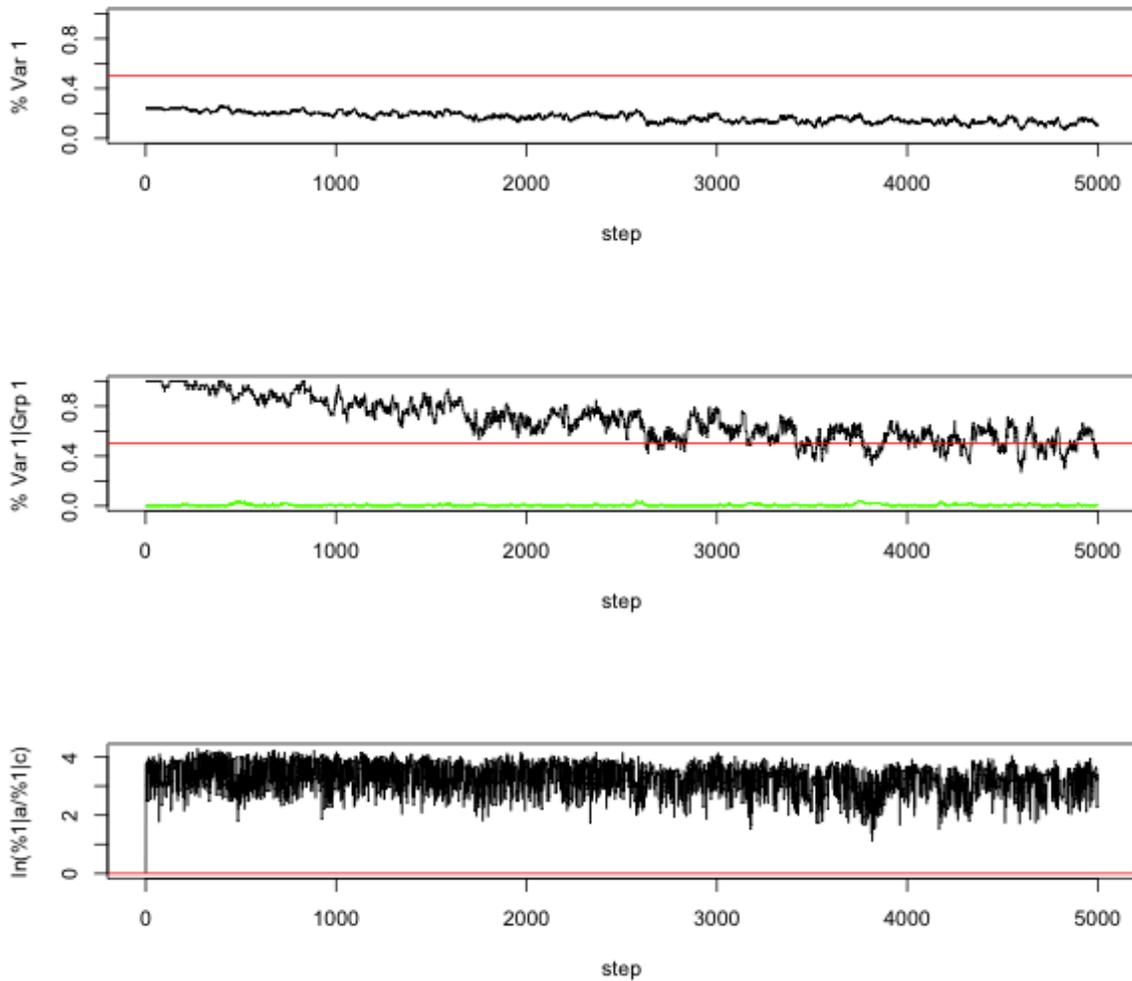


Figure 4. Stable context-stratified frequencies for locality-restricted motion under payoff-sensitive updating.

³ This specific threshold value is undoubtedly an artifact of the updating weight of 0.1 and the fact that the experimental initial $P(g_1)$ values are clustered around 0.25, 0.5, and 0.75 due to the setup conditions.

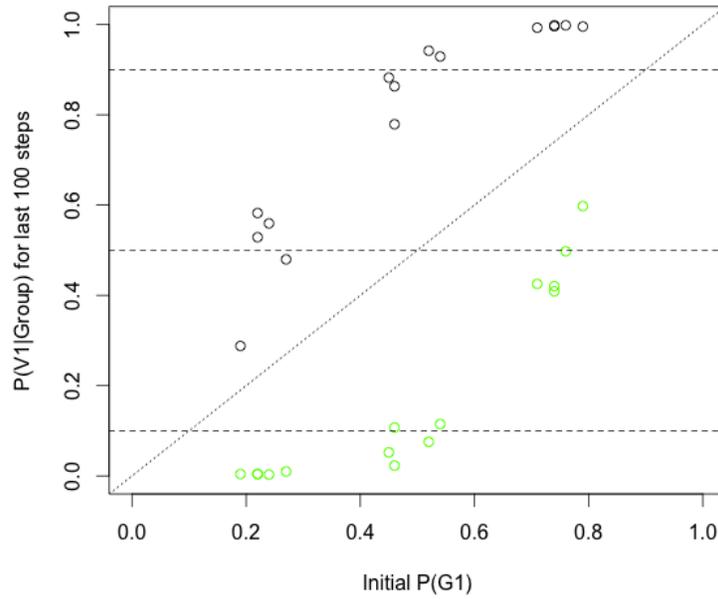


Figure 5. Final v_1 frequencies for g_1 (black dots) and g_2 (green dots) by initial proportions of g_1 .

This asymmetrical group behavior may explain the emergence of stable context stratification at the population level: the distribution of the stable majority group across contexts is determining the overall distribution of variants in each context, and the accommodating minority group is having a much smaller effect. Note also, however, that the degree of adoption of the majority variant by the minority group follows a linear relationship parallel to the line of equivalence with initial $P(g_1)$ (the diagonal line in the chart), such that a smaller minority shifts further towards the majority variant.

Under the exemplar-update mechanism, locality restrictions produce a much more random outcome, without the strict hierarchy $P(v_1|C_A) > P(v_1|C_B) > P(v_1|C_C)$, but still with some evidence of stable context-dependent frequencies. Table 4 below shows averaged final $P(v_1)$ for each context over 8 runs. The differences are slight, with frequencies for each context more or less randomly distributed around $P(g_1)$. On the other hand, these frequencies may remain stable over many runs, as illustrated by Figure 6 for three separate cases. From our small dataset, it appears that under these conditions stable the emergence of stable context stratification is disfavored by $P(g_1)$ near 0.5 (even group distribution). Considering that under these conditions, the probability of using v_1 at a given point for any single agent is either 1 or 0, even the weak stable variation exhibited here is remarkable.

$P(g_1)$	$P(v_1 C_A)$	$P(v_1 C_B)$	$P(v_1 C_C)$
0.23	0.189	0.227	0.227
0.28	0.254	0.316	0.340
0.35	0.426	0.422	0.448
0.40	0.549	0.475	0.391
0.42	0.368	0.476	0.571
0.47	0.450	0.522	0.531

0.69	0.764	0.758	0.744
0.72	0.766	0.779	0.781

Table 4. Averaged final v_1 frequencies for exemplar updating with locality restrictions.

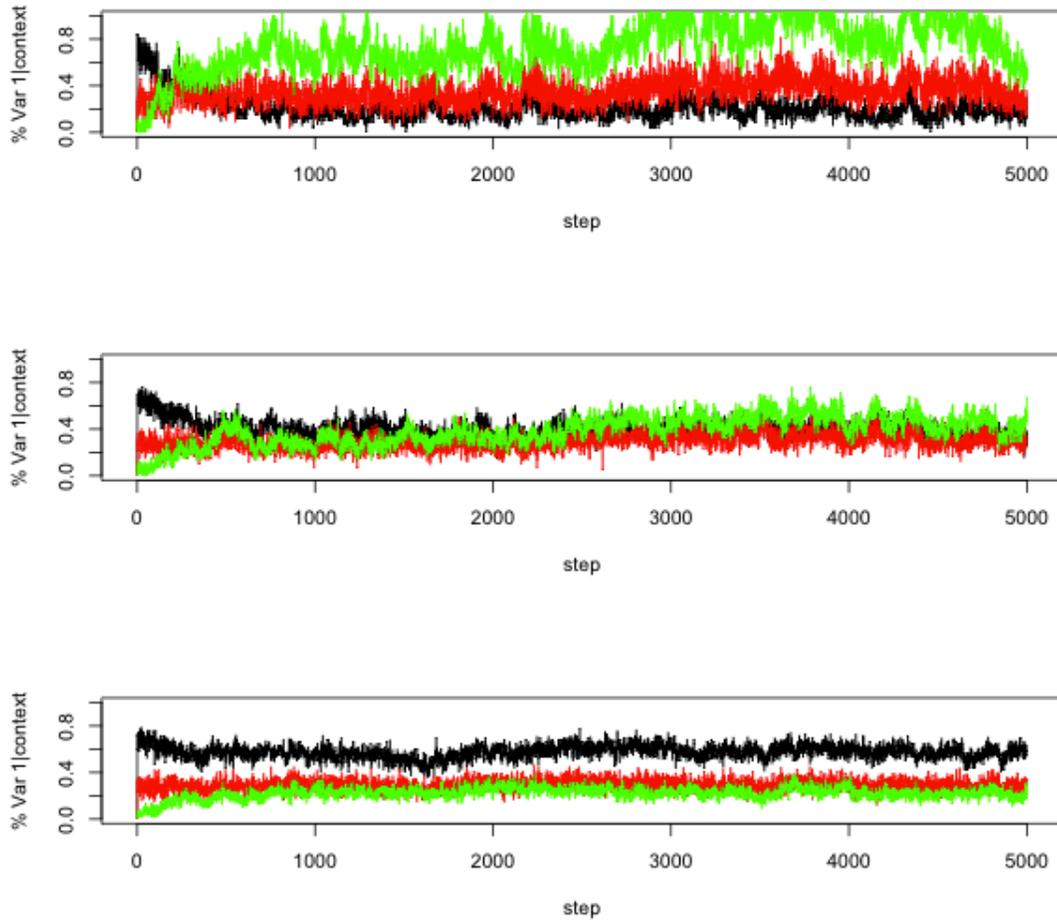


Figure 6. Patterns of context differentiation with exemplar updating for runs with $P(g_1) = 0.26$ (top), $P(g_1) = 0.47$ (middle), and $P(g_1) = 0.71$ (bottom), with $P(v_1|C_A)$ in black, $P(v_1|C_B)$ in red, and $P(v_1|C_C)$ in green.

Finally, given memory-averaging updating, we see the same context- and group-independent convergence discussed for unrestricted interactions above, only operating at a slower rate. Figure 7, to be compared with Figure 3, illustrates the convergence of v_1 frequencies in a case with initial $P(g_1) = 0.22$.

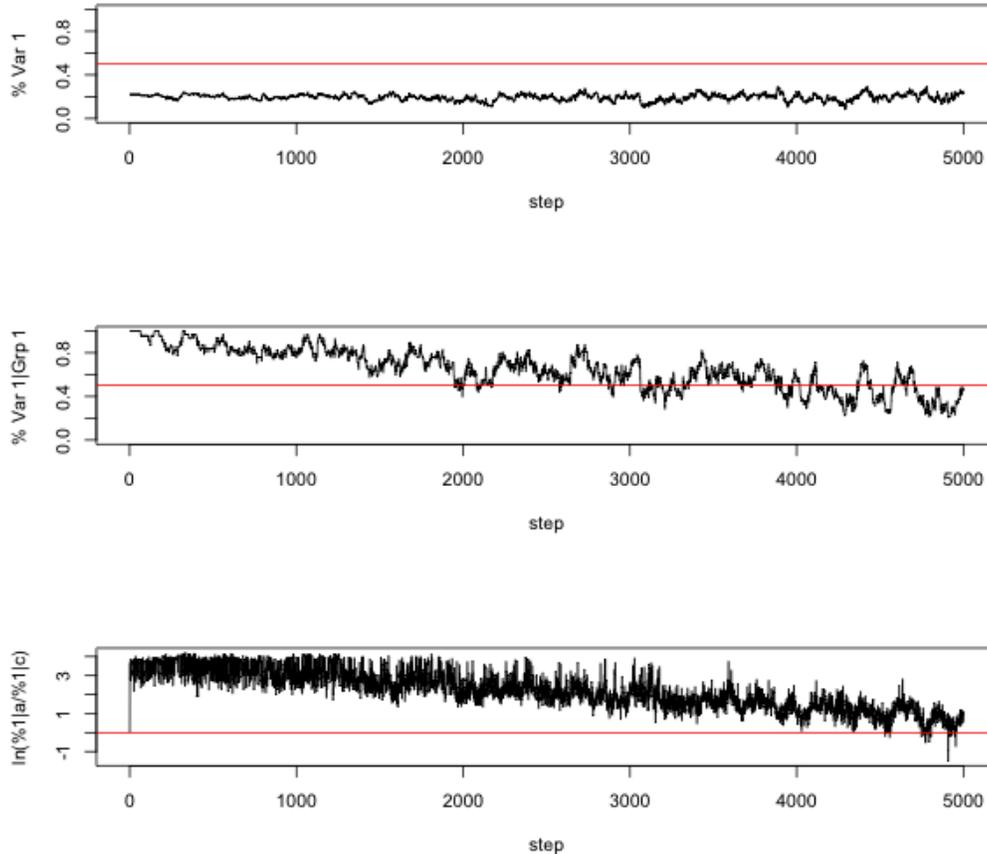


Figure 7. Convergence of variant frequencies to match group proportions with locality-restricted memory-averaging updating.

3. Model Behavior with Age Restrictions on Learning

The model may be further modified to simulate critical period effects in acquisition, by restricting updating to an agent's first 100 interactions. The model environment is not ideal for studying age effects, since new agents are introduced into the model with exact copies of their parents' updating information, but it is clear that the effects of age limitation on learning are not simply to slow down the model by delaying an agent lineage's updating. Instead, the simultaneous interaction of actively updating agents and fossilized agents often results in context-dependent stable variation similar to that generated by locality restriction. A notable difference is that this stable outcome extends to the memory-averaging update condition as well.

This result is most regular in the case of unrestricted interaction for payoff-dependent updating and memory-averaging updating, and in the case of interaction with locality restrictions for exemplar updating. It was found that the combination of critical period effects with locality restrictions for payoff-dependent and memory-averaging updating resulted in an almost completely static model, with agents deviating slightly from categorical group behavior and immediately correcting themselves. In the case of exemplar updating, the unrestricted case with critical period effects yielded an outcome with more randomness than was observed with locality

restrictions alone, but the combination of the two factors produced orderly context-dependent variability.

Turning to the cases of critical period effects which exhibited orderly context stratification, we first consider payoff-dependent updating with no locality restrictions. Figure 8 shows $P(v_1)$ for each group varying by $P(g_1)$ for 30 runs. The relationship is not linear, as it appeared for locality restrictions with no critical period, but instead follows a steep S-curve. As before, when one group forms a clear majority, it does not deviate much from its initial behavior. However, minority groups accommodate very little when group ratios are balanced, and come close to matching the majority group's behavior when they are strongly outnumbered. The resulting pattern in the extreme case comes close to complete replacement by the majority variant, but due to fossilization the minority variant continues to survive with low frequency of use. It remains true in all cases that $P(v_1|C_A) > P(v_1|C_B) > P(v_1|C_C)$. Figure 9 plots v_1 frequencies by context for cases with initial g_1 proportions of 0.25, 0.50, and 0.80 respectively.

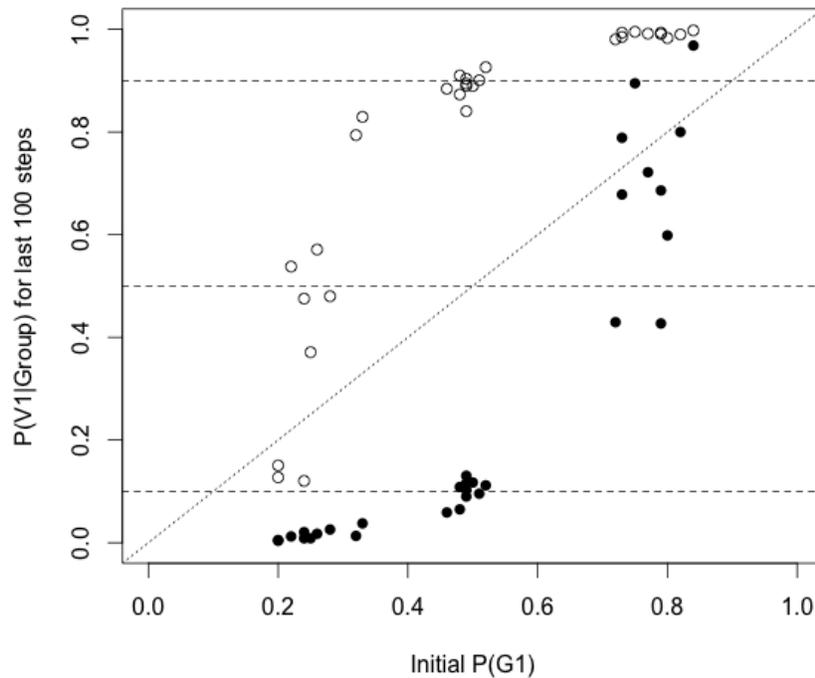


Figure 8. v_1 frequencies for g_1 (hollow dots) and g_2 (solid dots) for payoff-sensitive updating with critical period effects

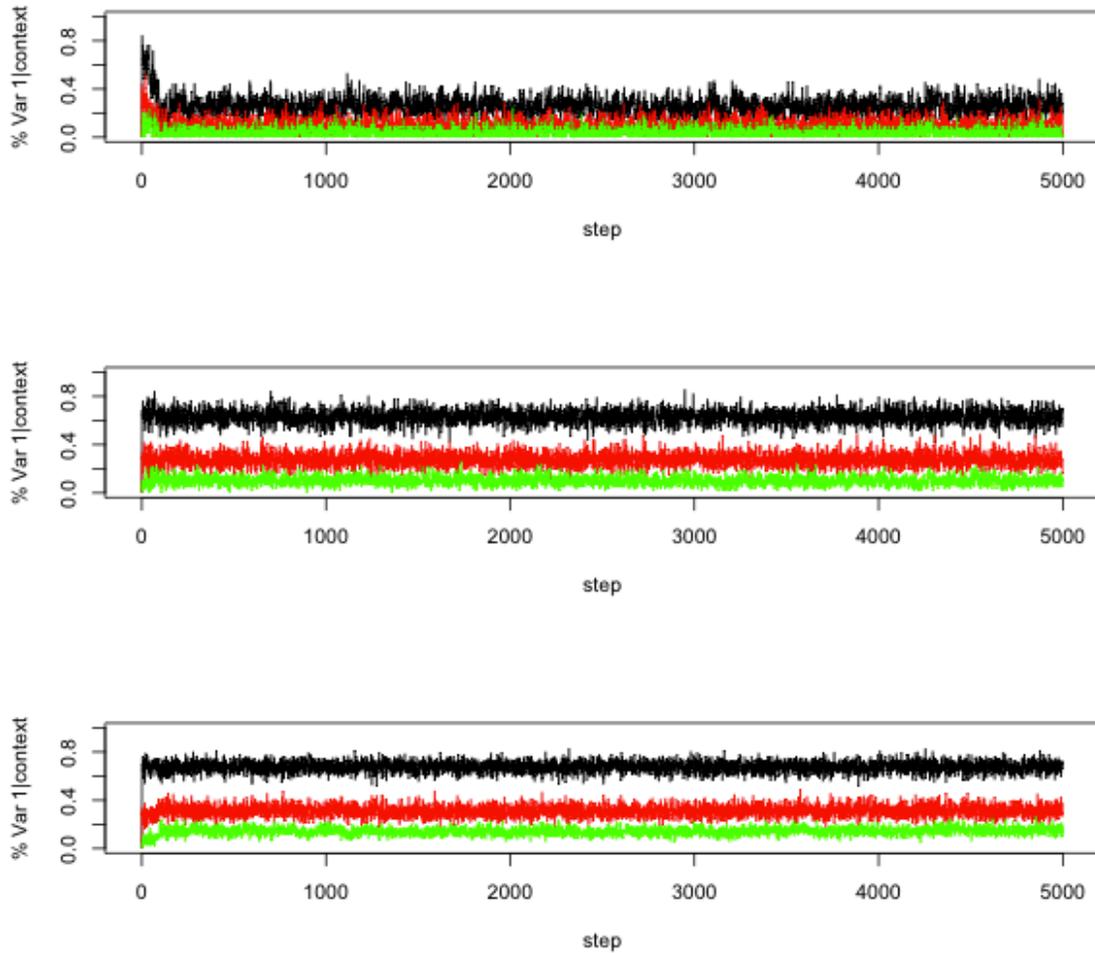


Figure 9. Context differentiation with critical period effects and payoff-sensitive updating for runs with $P(g_1) = 0.25$ (top), $P(g_1) = 0.50$ (middle), and $P(g_1) = 0.80$ (bottom), with $P(v_1|C_A)$ in black, $P(v_1|C_B)$ in red, and $P(v_1|C_C)$ in green.

For the case with memory-averaging updating and no locality restrictions, patterns of context- and group-dependent differentiation are both well established for all initial group proportions, in sharp contrast to the result when there are no critical period effects. Figure 10 plots v_1 frequencies by group and context for 30 runs under the relevant conditions.

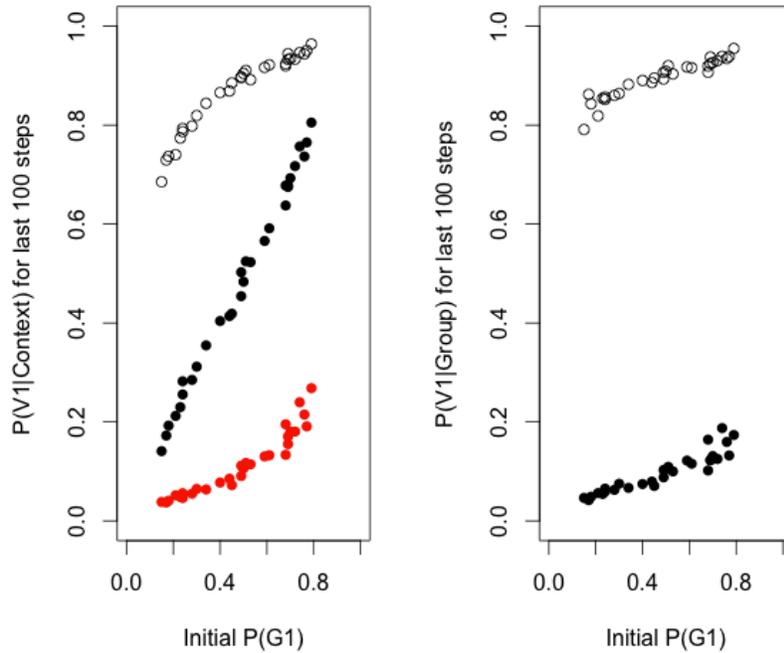


Figure 10. Proportions of v_1 by context and group for 30 runs with critical period effects and memory-averaging updating. On the left, $P(v_1|C_A)$ is hollow points, $P(v_1|C_B)$ is solid black points, and $P(v_1|C_C)$ is red points. On the right, g_1 is given by the hollow dots and g_2 is given by the solid dots.

Finally, the case of exemplar updating with both critical period effects and locality restrictions on motion follows a similar pattern only with more variability in the outcome, as illustrated by Figure 11.

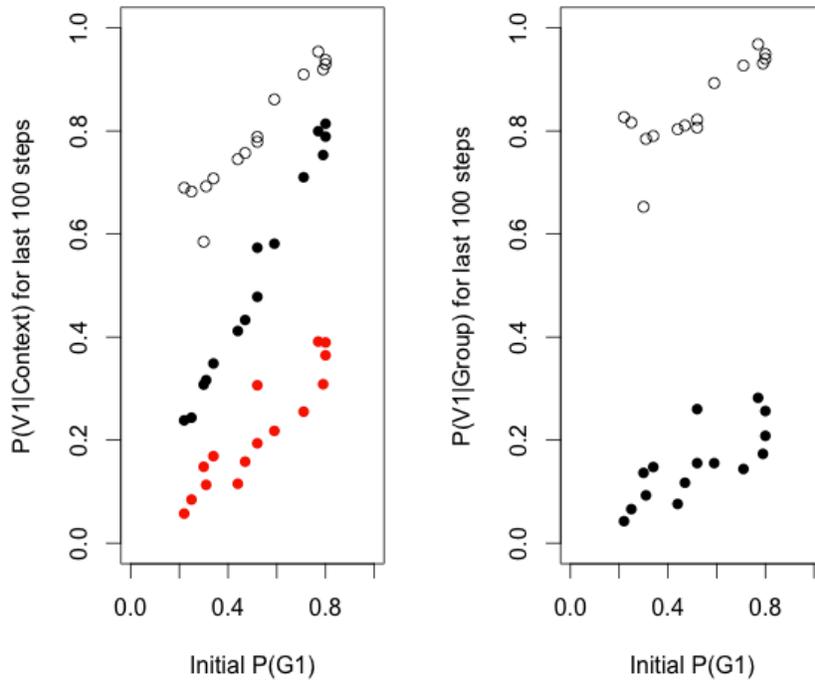


Figure 11. Proportions of v_1 by context and group for 30 runs with critical period effects and memory-averaging updating. On the left, $P(v_1|C_A)$ is hollow points, $P(v_1|C_B)$ is solid black points, and $P(v_1|C_C)$ is red points. On the right, g_1 is given by the hollow dots and g_2 is given by the solid dots.

Interpretation of Results and Discussion

The results described above indicate that the long-term outcome of the model is determined by only two factors: the update mechanism used by the agents and the degree of ‘impedance’ to grammar change contributed by other model parameters. The two impeding factors considered above were locality restrictions on interaction and age restrictions on learning, and ingroup bias has been observed to have a similar effect. The approach of combining local partner selection with motion if no local partner can be found simulates a dynamic small-world network within the model, and restricting learning to the first 100 age units simulates critical period effects in language acquisition.

The asymmetries corresponding to update mechanism are less clearly relatable to real-world conditions. Comparing them to other language models, we can see that the payoff-sensitive update method behaves similarly to the voting algorithm of Fagyal et al., which supported variation when there were ‘loners’ (as when the present model’s locality restrictions generate temporarily disconnected agents) but inevitably converged to the majority variant when loners were removed. This sort of approach, informed by evolutionary game theory, is common in models of language dynamics (cf. Castellano et al. 616 ff.).

The other two options, informed by exemplar approaches to language learning, are less well-tested in the agent-based modeling literature. Stanford and Kenny successfully used an update mechanism similar to the memory-averaging algorithm to model vowel shifts in a model population, without introducing any critical period effects. The failure of memory-averaging to produce meaningful long-term behavior here⁴ may be due to the fact that the variants in the present model are discrete and binary rather than gradient. Exemplar learning in linguistics is mostly supported in phonetics (cf. Bybee 2002), where objects exist in continuous and gradient dimensions. Thus the average of a phonetic memory produces an intermediate value within this spectrum, whereas the average of a memory of binary values produces a frequency -- which, given a large enough representative sample, approaches of the population frequency. So each step of the model generates closer and closer approximations of the population frequency within a certain margin of error, until every probability stored in the agents’ grammar is the same. It is not surprising, then, that this pattern is blocked by the addition of critical period effects, which serve to block individual agents from proceeding with the approximation process indefinitely and alter the population proportions through interaction with fossilized agents.

A similar argument can be used to explain the randomness emerging from the single-exemplar updating method. Here, instead of generating a frequency, the exemplar memory

⁴ Unless long-term levelling to initial proportions is considered meaningful as a real-world outcome, which is a difficult argument to make. For example, one would not make the prediction that three centuries following the conquest of England by 10% Normans, all Englishmen would settle into speaking 10% French and 90% English in every situation.

completely determines the binary choice of each agent based on a randomized interaction. Despite the fact that distribution of variants is unequal over contexts (a fact weakly influencing the results in Table 3), a modest number of interactions with proponents of a minority variant can quickly transform it into the majority variant in a given context. This same effect may serve to speed up takeover of a majority variant, and given enough time, random fluctuation will result in categorical behavior. The fact that this behavior is differentiated by context follows directly from the structure of the model grammar, which allows no direct relationship between variant choice for different contexts. Thus, while unimpeded memory-averaging fails because it is too information-rich, unimpeded single-exemplar updating fails because it is too information-poor.

Returning to the notion of ‘impedance’ as a factor in determining model behavior, we note that in all cases for which context-stratification emerges, it follows a similar pattern. This pattern turns out to be related to the behavior of the model in the null condition for which updating never occurs and agents simply distribute themselves over contexts with $v(g_i) = v_i$. In this case, the distribution of variants is identical to the distribution of groups, and the distribution of groups over contexts is entirely predictable from the model. In particular, the probability $P(g_i|C_x)$ of a given agent in context C_x being of group g_i can be expressed in terms of $G_i = g_i/N$, the proportion of g_i in the population, and the constant probabilities $\alpha_i^x = P(C_x|g_i)$ and $\alpha_j^x = P(C_x|g_j)$. The result is the following function derived from Bayes’ rule:

$$P_x(G_i) = \frac{\alpha_i^x G_i}{\alpha_i^x G_i + \alpha_j^x (1 - G_i)}$$

With the constants α_i^x and α_j^x set to the model’s values as indicated in Table 1, the relationship between group proportions in the population and group membership in each context is as in Figure 12 below.

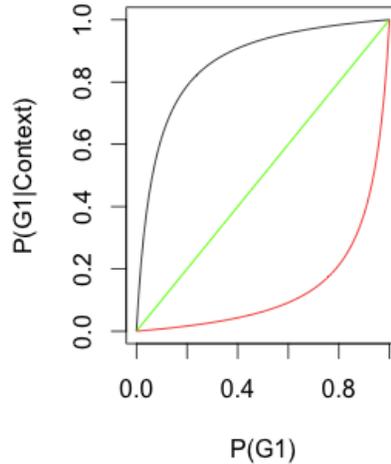


Figure 12. Group distribution over contexts (black = C_A , green = C_B , red = C_C) in the model.

Thus, in completely static conditions under which the groups never update, we should expect to see stable proportions of v_1 in each context corresponding to the proportions of g_1 above. We have seen a very similar relationship in Figures 10 and 11 above, and this is in fact the basic shape of stable context stratification in every condition which admits it.

This observation leads to the apparently inevitable inference that the emergence of context stratification in the model indicates that updating is not working and the model is simply static. This cannot be the case, however, for the several conditions under which interaction is occurring with high frequencies and final group-conditional frequencies are significantly altered from the initial categorical state, as for instance under locality-restricted payoff-sensitive updating or free interaction with critical period effects. In these cases, the model mechanism is instead acting to counterbalance the clear directional tendencies of the various update algorithms at each stage of updating, using the added parameters to maintain conservative patterns despite the dynamic behavior of individual agents.

The impedance effect can thus be seen to have some real-world implications for the study of language in society, given the reasonable assumption that groups are unevenly distributed over discourse contexts (such as Mühlenbernd and Quinley's 'farm' vs. 'dinner'). While one may remain agnostic as to the precise cognitive mechanism of language learning, any mechanism similar to those utilized in this model would be completely incompatible with the emergence of stable style or register stratification in an completely interconnected population with perfect learning. Some degree of register differentiation is, however, a near-universal feature of human language, just as functional diglossia is typical of multilingual communities (cf. Blumenthal and Kahane 1979). In order to maintain such differentiation, at the very least a critical period is necessary to regulate learning by individual agents, and a network structure with weakly connected individuals is also important. Both of these processes create individuals who are reservoirs of conservative behavior, but may continue to influence the trajectory of the population through interaction (cf. Fagyal et al. 88). A further implication is that the breakdown of social barriers may function to level register hierarchies and patterns of in-group norms, and favor instead convergence to a single norm by the population.

A more puzzling implication regards the split between lexemes and grammatical elements in acquisition. It is generally thought that lexemes are not as subject to critical period effects as grammatical variables, and may continue to be acquired into adulthood (Thomason and Kaufman 1988). This assumption underlies proposals such as the relexification theory of creolization (Mufwene 2006) and the observed ubiquity of loanwords in contrast to borrowed grammatical elements. However, if this is the case, our model predicts that context stratification will remain more stable for grammatical elements than for lexical items; words will be borrowed between groups more easily, but will also be less likely to be encoded as register markers. In fact, although grammatical style markers do exist (Labov 1968), words are also extremely common as style markers (cf. the entire basis of Mühlenbernd and Quinley's model). This asymmetry between our model's prediction and observed facts may be explained by the simplicity of the model's coordination game, which does not take into account the pragmatic function of lexemes as markers of context or group membership within a more complex signaling game (Clark 2012).

Concluding Remarks

The agent-based model developed in this study aimed to simulate the emergence of differentiated registers in a linguistically heterogeneous population, studying social and cognitive constraints on stable context-stratified variation. It was found that, given predictable and asymmetrical group assignment to contexts, the more conservative the model's long-term

behavior is, the more likely it is to support stable context stratification of variants. Elements contributing to conservative behavior included the stipulation of local/weakly-connected networks, rather than unimpeded communication across the entire population, and the enforcement of critical period effects on learning. It was further found that the learning mechanism of the model played a large role in determining the final outcome of the model, and that a payoff-dependent behavioral learning algorithm yielded results most consistent with previous work, in contrast with levelling or randomized outcomes associated with exemplar-based learning (although these mechanisms can also be made to produce stable context variation given conservative enough conditions).

Two issues which could not be addressed given the present form of the model, but would benefit from treatment in a similar approach, are the preservation of archaic registers and the feasibility of relexification from substrates. The first of these requires introduction of language change through ‘mutation’ in non-initial stages of the model, and could be examined in the simplest case using an initially linguistically homogeneous community with group asymmetries. The second requires a more complex grammar with asymmetrical learning algorithms for structural and lexical elements, which is excessively complex for implementation in NetLogo and would be better modeled in another environment. Our result for critical period effect asymmetries also suggests that a more complex signaling game would be needed to accurately represent the conditions needed to test the relexification hypothesis.

Works Cited

- Bell, A. 1984. Language style as audience design. *Language in society* 13(2), 145-204.
- Blumenthal, H., & Kahane, R. 1979. Decline and survival of Western prestige languages. *Language* 55(1), 183-198.
- Bybee, J. 2002. *Frequency of Use and the Organization of Language*. Oxford: Oxford UP.
- Castellano, C., Fortunato, S., and Loreto, V. 2009. Statistical physics of social dynamics. *Reviews of modern physics* 81(2), 591-646.
- Clark, R. 2012. *Meaningful Games*. MIT Press.
- Fagyal, Z., Swarup, S., Escobar, A., Gasser, L., and Lakkaraju, K. 2010. Centers, peripheries, and popularity: the emergence of norms in simulated networks of linguistic influence. *University of Pennsylvania Working Papers in Linguistics* 15(2), 81-90.
- Ferguson, C. A. 1959. Diglossia. *Word* 15(2), 325-340.
- Gilbert, N. 2008. *Agent-based models*. Sage Publications.
- Labov, W. 1968. *The Social Stratification of English in New York City*. Center for Applied Linguistics.
- Labov, W. 2001. *Principles of linguistic change Vol. 2: Social factors*. Oxford: Oxford UP.
- Mufwene, S. 2001. *The ecology of language evolution*. Cambridge: Cambridge UP.
- Nettle, D. 1999. Using social impact theory to simulate language change. *Lingua* 108(2), 95-117.
- Quinley, J., and Mühlenbernd, R. 2012. Conquest, contact, and convention: simulating the norman invasion’s impact on linguistic usage. *Proceedings of BRIMS*, 113-118.
- Stanford, J., and Kenny, L. 2013. Revisiting transmission and diffusion: An agent-based model of vowel chain shifts across large communities. *Language Variation and Change* 25(2), 119-153.
- Thomason, S. and Kaufman, T. 1988. *Language contact, creolization, and genetic linguistics*. Berkeley: University of California Press.