“An Invisible and Insuperable Barrier”: Enculturation, Stratification, and Ritual Interaction[°]

Carter Butts[*]

Department of Social and Decision Sciences

and

Center for the Computational Analysis of Social and Organizational Systems

Carnegie Mellon University

Abstract

Previous work on the sociology of culture has suggested a linkage between societal stratification and cultural differentiation. Here, a simple model is introduced in which culture emerges from ritual interactions between individual actors embedded within a larger social system. It is shown that this process leads to path dependence in actors' institutional affiliations, and that this in turn reproduces an equilibrium distribution of cultural practices across institutions. A simple virtual experiment is considered which compares path dependence across conditions, and some preliminary implications for the sociology of culture are presented.

[°] The author would like to acknowledge the input and assistance of Tom Fararo.

[*] This material is based upon research supported under a National Science Foundation Graduate Fellowship, and was supported in part by the Center for the Computational Analysis of Social and Organizational Systems and the Institute for Complex Engineered Systems at Carnegie Mellon University.

Introduction

> "If such be the function of culture and if it be love of art which really determines the choice which separates, as by an invisible and insuperable barrier, those to whom it is given from those who have not received this grace, it can be seen that...their true function [is] to strengthen the feeling of belonging in some and the feeling of exclusion in others." (Bourdieu, 1968)

In the above passage, Pierre Bourdieu speaks of artistic perception and, in particular, the way in which institutions such as art museums act to perpetuate a division between those who (by virtue of their education and social class) are trained to "decode" artistic works and those who are not. This statement, however, might as easily have been applied to rules of etiquette, occupational skills and credentials, or standards of formal research: even as cultural knowledge serves to bind persons together ((Durkheim, 1933), (Carley, 1991)) by creating the basis for social interaction, so too does it form an "invisible and insuperable barrier" which prevents certain interactions from taking place. Such barriers, in turn, serve to perpetuate cultural differentiation itself, and hence reproduce a system of stratification which overrules - and outlasts - the efforts of any individuals passing through it. In the following work, a formal theory of cultural development will be presented which draws on these ideas, and results from a computational model will be shown. These results will serve to verify some of the intuitions regarding cultural stratification presented by Bourdieu (1968) and others (e.g., Collins (1975)), and will suggest some general principles governing the emergence of stratification in social systems. Because this model is founded at the micro level, it also relies heavily on insights from symbolic interactionism (Blumer, 1969), (Skvoretz and Fararo, 1996); it is hoped that work such as this will demonstrate the rich possibilities of computation for establishing a linkage between micro-level symbolic models and macro-level cultural theory.

The Actor Model

> "...when in the presence of others, the individual is guided by a special set of rules... Upon examination, these rules prove to govern the allocation of the individual's involvement within the situation, as expressed through a conventionalized idiom of behavioral cues." (Goffman, 1963)

As Goffman suggests, the social actor is constrained to act in accordance with various rules and procedures; some of these, it is presumed, are the result of frameworks of expectations regarding the behavior of others, while others are due to learned associations which motivate the actor's behavior directly. Following Skvoretz and Fararo (1996), we seek to model individual actors as rule-driven entities who nevertheless may change their behavior over time. In particular, we shall here require that our actor model satisfy the following four requirements:

1. Actors must be capable of learning via an enculturation or influence process.
2. Actors must be capable of communicating with each other.
3. Actors must be able to engage in substantive behavior.
4. Actor communication must be able to influence behavior.

Individually, these requirements are fairly trivial to satisfy: influence models (e.g., (Friedkin and Cook, 1991), (Anderson, 1959)) have been developed in numerous contexts for quite a number of years; multi-agent models in which actors exchange information directly have been developed by Carley (et al, 1992) and others; as game theorists will attest, whole fields of study have grown up around the notion of actors as engaging in substantive behavior; and signaling in multi-player games is these days considered mundane enough to qualify as textbook science (Binmore, 1992). To combine all of the above in a single model, however, is more difficult. Work by Carley (et al, 1992) has included communicative actors engaging in substantive behavior, drawing upon the plural-SOAR agent model. Other work in the distributed AI community, as well, has considered this sort of problem, but the frameworks in which these models have been applied have generally been very limited, containing few agents and relatively little opportunity for socialization. In order to construct a model of culture, then, we require a framework which implements the above requirements, but which is also simple enough to permit implementation at the population level.

One solution to this dilemma is to model each actor as a finite state automaton. In particular, we here conceive of the actor as possessing some number of behavioral "states", each of which is tied to some communicative or substantive action. Each state is mapped to other states via a series of contingent links: hence, the state to which an actor moves at a given time point is dependent *both* on the previous state of the actor and on inputs from the environment. This actor model corresponds to the type of rule-based, path dependent ritualized behavior described by Goffman (1963), Skvoretz and Fararo (1996), and satisfies our need for a complex agent capable of communicating, acting, and learning. Furthermore, automaton models have the added advantage of being fairly compact; implementing the hundreds of agents required for a simulation of culture is entirely feasible using such a system.

An example of a simple four-state automaton is shown in Figure 1 below. As indicated above, each state of the automaton (represented by a circle) contains an action; here, those actions are presumed to involve either sending one of two communication tokens (numbered 0 and 1 respectively) or making one of two choices (A or B). In the example, as in the model we shall be using, choice actions are final: once an actor commits him or herself to a choice, he or she may not renege on it. Communicative actions, on the other hand, provide the possibility for contingency. Note that each communicative node in Figure 1 is linked by two outgoing arrows to one or more other nodes; these arrows represent the *state transition* which occurs when an actor at a given state receives a given signal. As is clear from the example, these transitions need not allow all states to be reached, nor need they always proceed to a different state (reflexive ties are permitted). Likewise, there is no requirement that an actor have at least one of every possible action available on some state: the automaton below, for instance, has no state carrying the action "send token 1," even though this action is presumably possible. Different actors may have different capabilities, and, though learning, these capabilities may change over time.

Having considered something of the anatomy of a typical automaton, let us briefly examine the way in which it operates. The automaton is presumed to start in an initial state (marked in Figure 1 with the "Begin" box) which is constant across executions. In the first time step, the automaton executes whatever instruction is contained within its state; here, that instruction is "send token 0," so our automatic actor promptly sends a "0" to whomever he or she is interacting with at the present. At the same time, presumably, that actor is also responding to his or her initial state. For the moment, let us presume that he or she sends a "1" to the sample actor. If this is the case, our actor follows the "receive 1" arrow, and winds up choosing B; at this point, the program halts (though in the implementation used here, actors who have chosen always send a default signal to their interaction partner). If, on the other hand, our actor receives a "0", things become a bit more interesting. Following the arrow, we note that the automaton is again sending the "0" token...now, however, the actor responds differently to external inputs. At any point after this, a "0" token causes a transition to the "choose B" node, and a "1" token causes the actor to return to the current node, sending "0" and starting over again. Needless to say, this "waiting" algorithm could run into trouble if its interaction partner sends only "1"s after the first time step...infinite loops are possible both within single actors and between multiple actors, and must be detected for by our simulation algorithm. (Fortunately, this is not difficult, as it can be shown that the maximum non-repeating conversation length between two automata is equal to the product of their numbers of states, minus 1[1].) This simple representation, then, is clearly capable of a wide range of rule-based, communicative behavior; what is missing, however, is a model of learning.

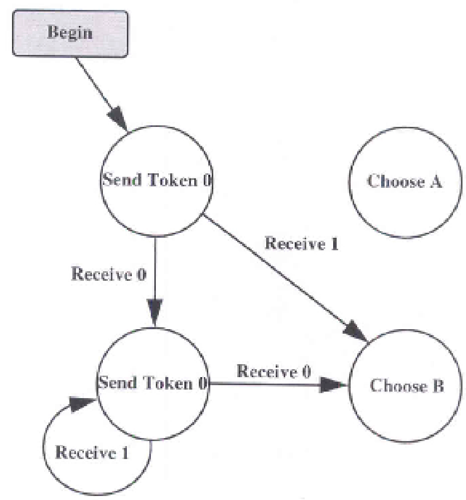

Figure 1: Example of a Simple Finite State Automaton

In addition to being capable of communication and substantive action, actors in the present model learn via a mimetic process of social comparison and emulation, and by a process of "trial and error" experimentation. At each time step, it is presumed that each actor is able to

[1] This follows from the fact that the two automata can be represented together as a single, larger automaton with a number of states equal to the product of the two smaller machines' states (i.e., the total number of possible configurations of the two machines). As state transitions within the joint automaton are equivalent to walks on the graph formed by its states and transition rules, it naturally follows that the longest non-repeating conversation is equivalent to the longest path length. Since the longest path length within any graph can be at most n-1 hops (where n is the number of nodes in the graph), it therefore follows that the longest conversation length between two machines is the product of their sizes minus 1.

compare him or herself with some other (randomly chosen) alter from the same institutional subpopulation. The basis for comparison in this case is the actor's "score" (which we shall describe presently), or more generally the overall success of that actor in interacting with other group members; actors are assumed to emulate alters only insofar as those alters are more successful than they (relative to the variance in the population as a whole) in ritual performance. While actors are somewhat systematic in when and how much they choose to imitate others, however, they are assumed to be less effective at selecting that which they imitate. Instead of picking up targeted behaviors, or lifting entire strategies wholesale, it is here assumed that actors emulate randomly selected aspects of alters' behavior patterns. While this may seem unreasonable at first blush, there are several points of justification for it. First, it is not in practice possible for actors to observe entire sets of behavioral rules possessed by alters; rules must be inferred, and the availability of the evidence which is required for such inference is heavily restricted by situational factors which the actor is unlikely to be able to control. Second, it is not clear that the emulation process which is posited here is a conscious one, and certainly not that it is in any way a rational process. The brain, while an excellent device for uncovering patterns, is likewise a reasonably haphazard one, and the particular elements of an alter which are noticed and incorporated into ego's behavior are likely to be highly variable. Even if we did presume some sort of rationality in the mimetic process, however, our third point would apply: an actor who observes an alter at some distance and/or in a limited number of interactions may not be able to discern the particular behavioral rules which are responsible for his or her success, and hence may not be in a position to select out only favorable behaviors. For all of these reasons, then, we model mimicry as a random affair, governed only by the relative difference in standing between ego and alter.

As mentioned above, emulation is only the first of two learning mechanisms which are considered to be available to social actors. While it is entirely possible for such actors to learn by incorporating the behaviors of others into their own routines, so too is it possible for actors to experiment with new behavioral rules "on their own", as it were. This process (typically identified with "mutation" in the adaptive agent context (Holland, 1992)) is assumed to take place through the introduction of random variation into actors' behavioral rules at each time step; the degree of such variation for a given actor is further assumed to be inversely proportional to his or her interactive success vis a vis others in his or her social group. Here, as before, a vague sort of intelligence is invoked without presuming rationality per se: actors are "smart" enough to change more when things are going badly than when things are going well. This is in accordance with the "win-stay, lose-shift" meta-strategy uncovered by Axelrod and others in computational dilemma research (Axelrod, 1987), and is compatible with a fairly conservative model of human search behavior.

The Ritual Interaction Game

In addition to engaging in simple conversation, it is assumed that actors within this model participate in substantive interaction of some sort. The actual nature of this interaction is left unspecified - co-participation in task completion, involvement in a recreational activity, or even enactment of a formal ritual might qualify - but it is assumed that those involved in the interaction care in some way about the outcome. Following the extensive work which has been done in studying ritual interaction (e.g., (Collins, 1988), (Goffman, 1963), (Goffman, 1959), (Durkheim, 1912), (Mehan and Wood, 1975)), we may set out some basic requirements that a substantive interaction should satisfy:

1. Actors engaging in ritual interaction have an essentially infinite number of possible behaviors (within some set bounds).

2. Actors have no a priori reason to favor one ritual over another.

3. Some possible rituals "coordinate" with others in some way (e.g., when two persons agree on using "hello" as a greeting).

4. The degree of ritual coordination varies with the particular combination of choices.

5. All parties to a ritual interaction suffer insofar as their choices do not coordinate.

The framing of this problem above in terms of "choices" of particular behaviors which then have concrete consequences for actors suggests the formalism of the game. Furthermore, the sort of game being described appears similar in a many respects to the family of games known as "matching" or "coordination" games in the game theoretic literature (Schelling, 1962). Following this intuition, let us define such a game:

> **Definition:** Let the *ritual interaction game* be defined as the two player simultaneous-move game in which each player must select a strategy s from the real interval $[0,1]$ with payoff functions for the strategy choice pair (s_1,s_2) given by $\pi_1=\pi_2=1-|s_1-s_2|$.

The ritual interaction game defined above has a number of interesting properties. The number of possible strategic pairs, for instance, is clearly infinite: while there are obviously some constraints on the sorts of ritual in which actors may engage, there is nevertheless an infinite number of options within those constraints. The game also satisfies the notion of ritual coordination mentioned above, as it is in each player's interests to "match" strategies in some way. Furthermore, it is not readily apparent that any given strategy (ritual) is favored over any other. At least initially, then, the ritual interaction game would seem to be a reasonable candidate for fulfilling our substantive interaction requirements.

How can we obtain a deeper sense of how the ritual interaction game "works"? While the present work does not assume rational actors, it is nonetheless both interesting and useful in this regard to examine some game theoretic results concerning the properties of the ritual interaction game. It was claimed above, for instance, that the ritual interaction game incorporates our intuition that actors seek to "match" their rituals in some fashion, but that there is no non-social reason to favor one ritual over another. This claim can be expressed more formally in game theoretic terms; consider, for instance, the following theorem:

Theorem 1: Every pure strategy, s, of the ritual interaction game is a Nash equilibrium strategy.

Proof: Let the pair (s, s_2) represent the strategy choices of players 1 and 2 respectively. We then may represent s_2 as $s+\delta$, where δ is some real number in [-1,1] such that $s+\delta$ lies in the interval [0,1]. By the definition of the Ritual Interaction Game, the payoff for player 2 is given to be $1-|s-(s+\delta)|$; this simplifies to $1-|\delta|$, which is maximized when $\delta=0$. The best response of player 2 to a given strategy s, then, is $s+0$, or simply s, regardless of which strategy s may be. By symmetry, however, if player 2 chooses strategy s then s is also player 1's best response to player 2's strategy choice. The strategy pair (s, s), then, represents the mutual best response of player 1 and player 2, regardless of the value of s. Every pure strategy in the ritual interaction game, therefore, is a Nash equilibrium strategy. ■

Theorem 1 tells us that every strategy in the ritual interaction game is its own best response; hence, actors have an incentive to coordinate. These incentives are extremely robust, as we can see from our next formal result.

Theorem 2: Every Nash equilibrium of the ritual interaction game, (s^*,s^*), is also subgame perfect.

Proof: (by contradiction) Assume that the Nash equilibrium (s^*,s^*) is not also subgame perfect. By the definition of subgame perfection, therefore, there must exist some subgame such that $s^*+\delta$ is the best reply to s^*, where $\delta\neq 0$. From the definition of the ritual interaction game, however, we see that the payoff of responding to s^* with s^* is 1, and the payoff of responding to s^* with $s^*+\delta$ is $1-|\delta|$. (By symmetry, this is true regardless of who moves first.) In order for $s^*+\delta$ to be the best response to s^*, it then follows that $1-|\delta|>1$, which implies that $|\delta|<0$. Since this is impossible, it must be the case that every Nash equilibrium (s^*,s^*) is also subgame perfect. ■

Intuitively, theorem 2 tells us that the optimality of coordination behavior is not dependent on the simultaneity of strategy choice: if we were to let one actor state his or her strategy in advance, the other could do no better than to imitate him or her (and, in fact, any deviation would be punished). This is a more stringent requirement than that of theorem 1; the fact that it here is satisfied for all Nash equilibria (and, indeed, all reciprocated strategies) reinforces the generality of our formulation. Still, some questions remain: how do we know, for instance, that there are no "favored" equilibria? Given that players act in ignorance of each others' choices, we must consider the possibility that some strategies are generally more attractive than others, and hence that many of the possible equilibria will never be reached. In order to determine whether or not this is the case, we must look for dominant strategies[2]. This leads us to the following result:

[2] The result which follows is proved for weakly dominant strategies. Because all strongly dominant strategies must also be weakly dominant, the results of theorem 3 apply to these as well.

Theorem 3: Neither player has any dominant strategies in the ritual interaction game.

Proof: (by contradiction) Consider the dominant strategy s^*. By the definition of dominance, it must be true that $\pi(s^*,s) \geq \pi(s_a,s)\ \forall\ s_a \neq s^*, s$. This is equivalent to the condition $1-|s^*-s| \geq 1-|(s^*+\delta)-s|\ \forall\ s, \delta{:}\delta\in[-1,1], s^*+\delta\in[0,1], \delta\neq 0$, which implies that $|s^*-s| \leq |s^*-s+\delta|$. A simple counterexample to this condition is provided by the case in which $s^*\neq s$, and δ is chosen such that $\delta=s-s^*$. In this case, the above condition reduces to $|\delta|\leq 0$, which is disallowed; therefore, s^* cannot be a dominant strategy. By symmetry, the above argument applies to both players, hence neither player has any dominant strategies in the ritual interaction game. ■

Theorem 3 establishes that there are no dominant strategies in the ritual interaction game: the best strategy for either player depends on the other player's choice. This result is important in verifying our intuition that the ritual interaction game is not biased in favor of any particular outcome, and hence that it reflects a purely social process. An important extension of this finding is given by theorem 4:

Theorem 4: The ritual interaction game contains no evolutionarily stable strategies.

Proof: Consider the Nash equilibrium strategy s^*[3]. In order to qualify as an ESS, it is necessary for s^* to satisfy the condition $\pi(s^*,s) > \pi(s,s) \forall s\neq s^*$ (Binmore, 1992, p425)[4]. This statement is equivalent to the requirement that $1-|s^*-(s^*+\delta)| > 1-|(s^*+\delta)-(s^*+\delta)|\ \forall\ \delta{:}\delta\in[-1,1], s^*+\delta\in[0,1], \delta\neq 0$, which in turn implies that $1-|\delta|>1$. Because this requires that $|\delta| <0$ (which is impossible), s^* cannot be an ESS. As this applies to all Nash

[3] Evolutionarily stable strategies are necessarily Nash equilibria (Binmore, 1992, p425). This follows from the fact that the presence of an alternative best response would obviously permit invasion.

equilibria of the ritual interaction game, the game contains no evolutionarily stable strategies (see footnote). ■

This result indicates that not only are there no dominant strategies, but also that there are no evolutionarily stable strategies in the ritual interaction game. Informally speaking, this means that there is no preferred point of convergence for a coevolutionary process in which the strategies of the ritual interaction game are implemented as replicators. In the present work - where actors are modeled as adaptive agents - the result of theorem 4 is especially important. If one or more evolutionarily stable strategies did exist, then over time our agents would tend towards it. This would defeat the purpose of the ritual interaction game, and of the actors' communicative process, by biasing the outcome in favor of certain specific rituals. The fact that no ESS exists, then, combines with the above results to demonstrate formally that the ritual interaction game fulfills the basic requirements set out previously to apply to substantive interaction among actors within out model.

The Societal Model

At present, we have suggested a framework for modeling communicative actors; we have given those actors the ability to learn from each other, and from experience; and we have given them a concrete situation of interaction in which they can succeed or fail. What we have not established at this point is the "big picture": in what structures are actors embedded? From whence do they come? To whence do they go? While these are ultimately questions to be answered using models such as this, we must start somewhere. Here, we start with the following assumptions:

[4] This also follows from the need to prevent invasion; in this case, the condition requires that an ESS have a

1. Actors follow a demographic process of birth, institutional membership, reproduction, and death.

2. There are multiple "tiers" of societal institutions through which one advances as one progresses through the life course (e.g., educational institutions, work organizations).

3. At each tier, there may be multiple institutions into which actors may flow.

4. Entry into institutions is dependent on the presence of vacancies within those institutions, as well as the relative ability of actors to interact positively with current institutional members.

5. While within institutions, actors engage in ritual interaction with other institutional members.

The above list of assumptions ties the micro-interactive processes previously described with the evolution of the simulated population over time. By defining a structure of "tiers" and "institutions," it explicitly determines the context for direct interaction (within institutional boundaries) and indicates the way in which actors move through the system. This demographic process (which is similar in some respects to Harrison and Carroll (1991)) is implemented through the following algorithm:

1. A fixed percentage of the mature population (second tier) is removed from the population (death)

2. The surviving mature population members are given the opportunity for reproduction (birth)

3. Children are created via random crossover

4. The second tier vacancies are filled via an interview process (producing first-tier vacancies)

---

higher fitness when matched against an alternative strategy than that strategy has against itself.

5. The first tier vacancies are filled via the interview process

When a vacancy is created in an institution (via departure or death), we assume that measures are promptly taken to fill it. This requires the selection, via some means, of an actor from a lower tier for assumption of the vacant position; here, this selection process is enacted through a series of *interviews*. In an interview, a fixed number of randomly chosen members of the vacancy-containing institution undertake ritual interactions with all eligible lower-tier actors (in this case, all actors within the tier are assumed eligible). The actor receiving the highest mean "score" (that is, the actor whose interaction patterns are the most compatible with the interviewers) is then inducted into the organization, spawning a new vacancy at his or her old position.

The above system allows for both a macrostructural organization of society and a crude demographic process, without giving up the fundamentally interactive basis of the cultural model. Because institutions recruit via personal interactions, the micro model which governs ritualized contact between persons also controls the flow of persons from one institution to another. While one can imagine other, more constraining systems, these alone are sufficient to produce stratification. To understand how this can come about, let us proceed to an examination of the model's behavior.

Model Behavior

Systematic evaluation of the culture formation model requires cross-conditional comparisons via virtual experiments; often, however, it is useful to examine a single simulation run in detail in order to get a feel for the sorts of processes which are going on behind the scenes before delving aggregate statistics. With this in mind, we shall briefly consider in "case study" fashion a number of results from a fairly typical execution of the culture formation model. This particular case included three first tier ("educational") institutions and three second tier ("occupational") institutions[5] with a total population of 200, a flow rate of 10% per time step, and a interview size of 10; actors were allotted 10 states each, with 5 possible communication tokens and up to the full 100 rounds to negotiate during the ritual interaction game. The simulation was run for 200 time steps total, and data was collected at each time step.

With that background in mind, then, let us turn to the data. Figure 2 presents the most general statistics collected during this run: the cross-institutional stratification indices. As these are the three statistics which will be used during our virtual experiment, it is worth paying some attention to their definitions and properties. The first of these three statistics is the educational history percentage; this measures the overall percentage of actors who are accounted for by the modal educational history of all actors in their workplace; as such, it provides a simple measure of the degree to which firms "hire" from the same educational source. The higher this mesure, then, the more channeling there is of actors from particular first tier organizations to particular second tier organizations. Starting from a natural level near 33%, the measure climbs rapidly and then levels off, bouncing randomly around the 90% mark. This means that, on average, 90% of the work force of any given occupational organization come from the same educational organization. This represents a striking pattern of cultural stratification. There is still some occupational flexibility when averaging across organizations, but the vast majority of actors are channeled directly into particular work organizations as a result of their educational backgrounds.

Secondly, a less extreme, but no less interesting story is told by the parental history percentage (also seen in Figure 2) which measures the overall percentage of actors for whom at

[5] For purposes of this simple demonstration, we shall refer to the first and second tiers as if they represent educational and occupational institutions, respectively. While this is a valid interpretation of the model, it should be noted that it is not the only possible one.

least one parent occupied at some point the same institutional slot (job or education) which they currently hold. This can be thought of as a measure of intergenerational stratification: given that there are three instances of each institutional type of organization, a high parental history percentage indicates that a surprising fraction of actors are following in their "family footsteps." This occurs because (in a differentiated environment) a child is far more likely to be compatible with his or her parents' educational background than with the alternatives. Such sorting at the educational level, in turn, leads to sorting at the occupational level, where the cycle repeats itself. As can be seen from Figure 5, however, this intergenerational stratification is neither total nor constant. In particular, the measure levels off rapidly at approximately 60%. While not extreme, this level *is* higher than the approximately 56% which would be expected given a purely random allocation of actors to organizations – an important finding, given that the model incorporates no homophilous or selective mate choice mechanisms. As we shall see, this moderately high aggregate parental history percentage reflects the fact that certain organizational memberships tend to be highly hereditary, while others are much less so. Clearly, then, there are some members of our artificial society who are able to break out of their familial heritage, but not many: *a class structure has been born before our very eyes!*

Finally, the third measure presented in Figure 2 is the conversation mode percentage, or the percentage of the population within each organization accounted for by the modal conversation pattern, averaged across all organizations in both institutional tiers. During the simulation run, a formal method is used to construct a dictionary of conversations between actors. The incidence of conversations within each organization is then counted, and the modal conversation is identified. The number of actors who share this conversation are thus "accounted for" by the mode, and in Figure 2 we can see that this fraction rapidly approaches 100%. Despite a high rate of transition through the population, it seems, actors are able to quickly learn a common language and to use that language to achieve coordination (as we shall see). It should be emphasized, of course, that the high value observed for this statistic does not imply that all actors within the population share the same conversations, or the same strategies – they do not – but rather that almost all actors *within organizations* communicate in the same fashion as their peers.

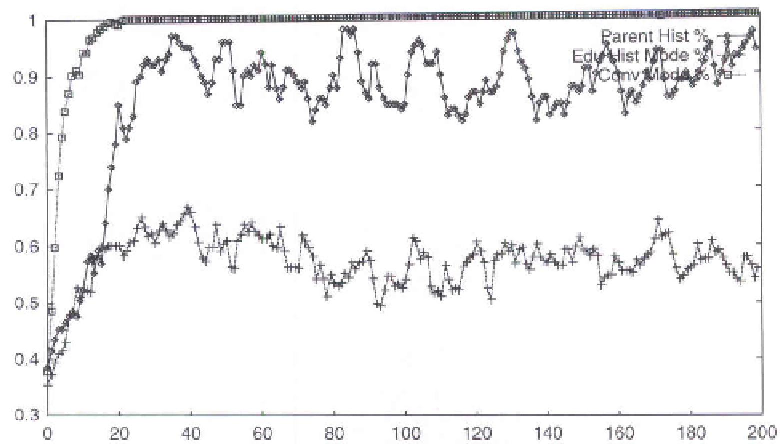

Figure 2: Aggregate Modal Percentiles

Continuing our investigation of the behavior of the culture formation model, Figure 3 shows the modal education histories and percentages so accounted for by occupational organization. As can be seen, two out of the three work organizations rapidly converge to a stable state in which they admit workers from only one educational institution; the third switches back and forth between a primary and a secondary hiring source, with between 50% and 100% accounted for by the mode at any given time. Here we can plainly see the "invisible and insuperable barrier" of which Bourdieu spoke. Actors who arrive for whatever reason at educational institution 3, for instance, cannot hope to be hired into job number 2, for no other reason than their relative cultural incompatibility. Even if such an actor originally *was* compatible with job 2, he or she would be rapidly changed by the need to survive in his or her local environment and would no longer be as "good" a match as before. This, of course, would compound itself: for every time step in which our hapless actor is passed over for job number 2, he or she becomes less isomorphic with 2's culture, and hence even less likely to get selected during the next round of interviews. Though driven purely by local interactions, this process creates global stratification.

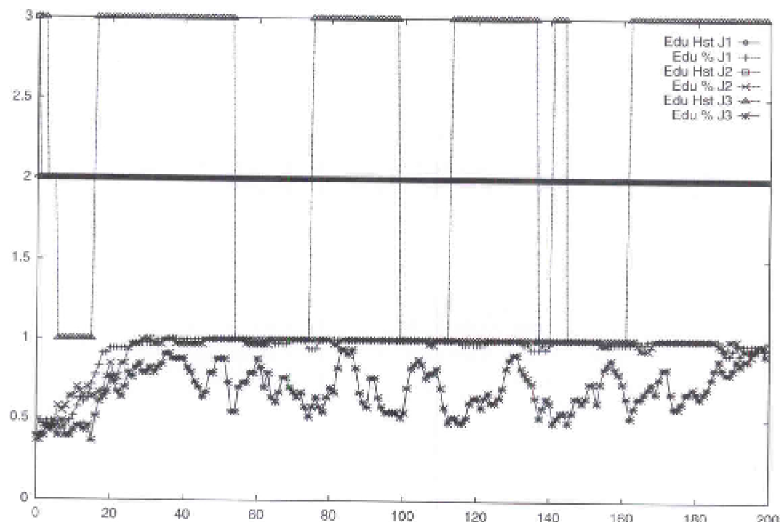

Figure 3: Modal Education History and % Accounted for by Mode, by Job

If Figure 3 presents a somewhat depressing lesson in the strength of cultural barriers, Figure 4 may suggest something of a modification. Underneath the apparent jumble of lines in this Figure lie the proportions of actors with one or more parents having occupied the same position, by institution. As can be appreciated, these proportions vary widely: further, they appear to oscillate within an essentially fixed range after an initial period of convergence, and then divergence. Jobs and educational institutions, then, run the gamut between the exceedingly hereditary and the wide open. This indicates that some, but not all, cultural stratification in the model operates through heritable mechanisms, and that the degree to which this is the case for a particular institution tends to remain reasonably constant across time.

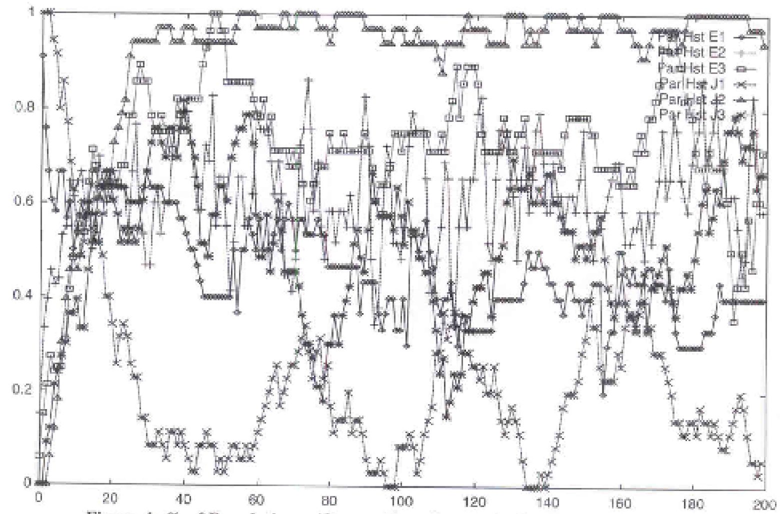

Figure 4: % of Population w/One or More Parents in Position, by Institution

Earlier, we saw how extremely rapid convergence in modal conversation took place across institutions; now, Figure 5 shows us that the same phenomenon is encountered within institutions as well. While there is some initial variation in speed of convergence, it appears that convergence to the mode is rapid and lasting across the board. Again, this should not be interpreted as indicating convergence to the same conversation for *all* actors, but rather the same conversation for actors *locally*. Likewise, the meaning of the conversation (in terms of the subsequent plays in the ritual interaction game) is not guaranteed to be constant.

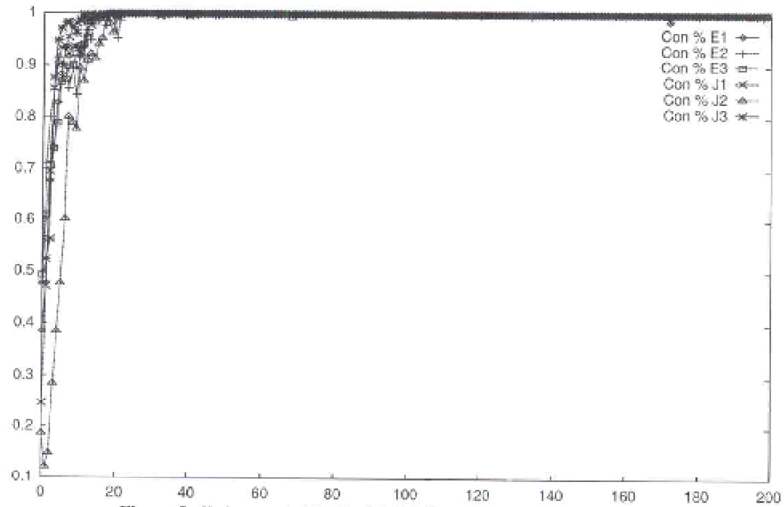

Figure 5: % Accounted for by Modal Conversation, by Institution

From questions of talk to those of action: Figure 6 presents the average strategies played in the ritual interaction game by institution. Here, in particular, it is possible to see the effect of cultural stratification in action. Note, for instance, how the sets of strategies quickly settle out in the initial rounds of the simulation with strong synchronization between job 2 and educational institution 2, synchronization which leads to a locking in of local rituals and a constant flow of recruits from the latter to the former. Some other institutions, such as E1 and J3, are close but unable to synchronize fully: J3 appears to be flipping back and forth (as we saw earlier) between hiring E1 and E3 employees, and as such is never able to converge to either. In the long run, one suspects, chance events will produce enough recruits from of the two sources to alter the institution substantially, and a convergence such as that surrounding job 2 will occur. Such interesting, nonintuitive findings are one of the many advantages of using simulations to formalize social theory.

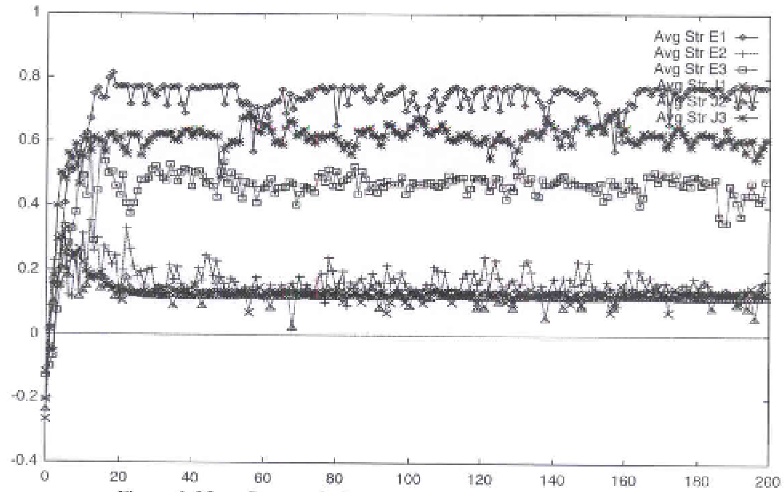

Figure 6: Mean Strategy in Ritual Interaction Game, by Institution

Virtual Experiment

By examining a particular simulation run in detail, it is possible to draw inferences regarding the fundamental processes at work in a model; to gain an understanding of the general behavior of that model, however, one must employ systematic, cross-conditional investigations. One member of this category of analytical approach is the virtual experiment. In a virtual experiment, as in a real experiment, a number of treatments are defined which are combined to form conditions; behavior is then observed in each condition, and statistical techniques are then employed to assess the effects of treatments and their interactions.

In the present case, the behaviors which we shall be interested in observing are the three global summary statistics produced by the culture model at the end of its execution[6]. As overall estimates of the level of stratification within the system, these variables reflect information which is of general interest and which, at the same time, is well enough understood at the micro level to permit interpretation of the experimental results. For experimental treatments, the population size per institution, the number of states per actor, the number of signal tokens actors had available for communication, the number of first and second tier institutions ("educational" and "occupational" institutions, respectively), the flow rate of the population through the model, and the number of interviewers (or "chances") per interview. While computational constraints prevented all of the desired data from being collected (and motivated a very restricted investigation to begin with), the resulting information provides enough data to attempt a simple regression analysis of the simulation output. Although this is not the most sophisticated way to approach this task, it is a naively reasonable one; it is hoped we may nevertheless learn something of interest, despite our methodology.

Table 1: Experimental Conditions

Treatment	Values
Pop/Institution	20, 100
#States	5, 10
#Tokens	3, 5
#1st Tier	2, 3, 4
#2nd Tier	2, 3, 4
Flow Rate	0.02, 0.05, 0.1
# Interviewers	5, 10
Limitations	Due to computational constraints, the 100/10 series was not completed.

Our first variable of interest is the global parental history measure. As indicated previously, this measure indicates the fraction of actors in the population at least one of whose parents had held an equivalent position previously. A simple perusal of this data series suggests that its distribution is fairly reasonable, and hence a straightforward model fit may be attempted. In lieu of optimal selection, we shall here fit only the complete model to the experimental data:

[6] Here, we are considering only "final" values (presumably after some stabilization has occurred). For purposes of this experiment, the simulation was set to run for 200 time periods before reporting its status.

given the other limitations present on this analysis, it would seem prudent to seek only the most robust results.

Fitting a linear model to the parental history variable produces the information given in Table 2. As can be seen, all regressors are insignificant save for the number of first and second tier institutions: both of these are negatively related to the parental history index. The logic behind this seems clear enough; increasing numbers of schools and businesses reduce the probability of generational overlap, and hence the parental history index falls. On the other hand, this was not altogether obvious a priori. One might imagine, for instance, that as the number of institutions increases, so too does the pressure for differentiation. In this case, stratification could induce higher than expected levels of correlation between parent and child histories. As it happens, however, this does not appear to be the case. Overall, the simple linear model explains approximately one third of the variance in parental history. This suggests that an alternative model would likely provide a better fit to the data[7].

[7] Obviously this is not the best or only way to assess fit. But it is convenient.

5. The first tier vacancies are filled via the interview process

When a vacancy is created in an institution (via departure or death), we assume that measures are promptly taken to fill it. This requires the selection, via some means, of an actor from a lower tier for assumption of the vacant position; here, this selection process is enacted through a series of *interviews*. In an interview, a fixed number of randomly chosen members of the vacancy-containing institution undertake ritual interactions with all eligible lower-tier actors (in this case, all actors within the tier are assumed eligible). The actor receiving the highest mean "score" (that is, the actor whose interaction patterns are the most compatible with the interviewers) is then inducted into the organization, spawning a new vacancy at his or her old position.

The above system allows for both a macrostructural organization of society and a crude demographic process, without giving up the fundamentally interactive basis of the cultural model. Because institutions recruit via personal interactions, the micro model which governs ritualized contact between persons also controls the flow of persons from one institution to another. While one can imagine other, more constraining systems, these alone are sufficient to produce stratification. To understand how this can come about, let us proceed to an examination of the model's behavior.

Model Behavior

Systematic evaluation of the culture formation model requires cross-conditional comparisons via virtual experiments; often, however, it is useful to examine a single simulation run in detail in order to get a feel for the sorts of processes which are going on behind the scenes before delving aggregate statistics. With this in mind, we shall briefly consider in "case study" the previous model, no actor-level variables proved to be relevant to the overall outcome. This may indicate that the external social forces of the environment are more important than minor effects produced by actor behavioral limitations, or it could serve to signify that actors are not using the majority of their machine space. A deeper examination will be required to determine which is the correct answer.

Education History Regression

Regression Statistics	
Multiple R	0.38572736
R Square	0.1487856
Adjusted R Square	0.1287233
Standard Error	0.1288238
Observations	305

ANOVA

	df	*SS*	*MS*	*F*	*Significance F*
Regression	7	0.861530383	0.123076	7.416181	3.19532E-08
Residual	297	4.928884879	0.016596		
Total	304	5.790415262			

	Coefficients	*Standard Error*	*t Stat*	*P-value*	*Lower 95%*	*Upper 95%*
Intercept	1.1062643	0.065084901	16.99725	9.98E-46	0.978178332	1.23435027
Base Pop	0.00012161	0.000232769	0.522439	0.601754	-0.00033648	0.00057969
#States	-0.0031327	0.00350614	-0.89349	0.372316	-0.01003274	0.00376731
#Sigs	-0.0026985	0.007456714	-0.36189	0.717694	-0.01737317	0.0119762
#Eds	-0.0322972	0.009278398	-3.4809	0.000575	-0.05055693	-0.0140375
#Jobs	-0.0376419	0.009042234	-4.1629	4.12E-05	-0.05543686	-0.0198469
Flow	-0.8454599	0.223821446	-3.77739	0.000191	-1.2859367	-0.4049832
#Chances	0.00662882	0.002950705	2.246521	0.025405	0.000821885	0.01243576

Table 3: Educational History Regression

Examination of our presumptive third experimental variable, the global conversation stratification measure, revealed that it was essentially degenerate. As our single-run analysis suggested, all institutions converged quickly onto a single modal conversation, leaving insufficient variance to attempt to fit a regression model. Even this, however, was informative: we now have evidence for the generality of the conversation convergence phenomenon, which we could not have gleaned had we not employed an experimental methodology.

Conclusion

While this is only the barest outline of a full analysis of the cultural creation model, it has nevertheless served to suggest several avenues for future research. More systematic investigation of the circumstances which are required for local stratification, and of the interaction of conversation and strategic play, would seem to be logical next steps. Furthermore, extension of the model to support multiple tiers, ongoing familial structures, and other forms of stratifying influences (such as preferential organizational relationships, homophilous mating, etc) could provide more information regarding the interaction between different demographic and cultural structures. Though this model is currently a simple one, and though our investigation has been limited, however, we have clearly shown one thing: that the "invisible and insuperable barrier" proposed by Bourdieu can emerge from a combination of interactive and structural forces, without the assumption of complex power relationships or bourgeois conspiracies. In this emergent model of culture, there are no "natural elites," no powerful businessfolk, no commoners and no aristocrats...only ordinary people, who are placed into societal niches by forces beyond their comprehension or control. It is the social system, not its members, which produces (and reproduces) cultural stratification.

References

Anderson, N. H. (1959), "Test of a Model for Opinion Change," *Journal of Abnormal and Social Psychology*, 59, 371-81.

Axelrod, Robert (1987), "The Evolution of Strategies in the Iterated Prisoner's Dilemma," pp32-41 in L. Davis (Ed) *Genetic Algorithms and Simulated Annealing*. London: Pitman.

Binmore, Ken (1992), *Fun and Games: A Text on Game Theory*, Toronto: DC Heath and Co.

Blumer, Herbert (1969), *Symbolic Interactionism: Perspective and Method*, NJ: Prentice-Hall.

Bourdieu, Pierre (1968), "Outline of a Sociological Theory of Art Perception," *International Social Science Journal*, 20(4).

Carley, Kathleen (1991), "A Theory of Group Stability," *American Sociological Review*, 56(3), 331-354.

Carley, Kathleen, Kjaer-Hansen, Johan, Prietula, Michael, and Newell, Alan (1992) "Plural-Soar: A Prolegomenon to Artificial Agents and Organizational Behavior," pp87-118 in Masuch M. & M. Warglien (Eds.), *Artificial Intelligence in Organization and Management Theory*. Amsterdam: Elsevier Science Publishers.

Collins, Randall (1988), *Theoretical Sociology*, San Diego: Harcourt Brace Javanovich

-- (1975), *Conflict Sociology: Towards an Explanatory Science*, Academic Press.

Durkheim, Emile (1933), *The Division of Labor in Society*, New York: The Free Press.

-- (1912), *The Elementary Forms of Religious Life*, London: Allen and Unwin.

Friedkin, Noah, and Cook, Karen (1991), "Peer Group Influence," *Sociological Methods and Research*, 19, 122-143.

Goffman, Erving (1963), *Behavior in Public Places*, New York: The Free Press.

-- (1959), *Asylums*, New York: Doubleday.

Harrison, J. Richard, and Carroll, Glen (1991), "Keeping the Faith: A Model of Cultural Transmission in Organizations," *Administrative Science Quarterly*, 36, 552-582.

Holland, John (1992), "Adaptation in Natural and Artificial Systems," Boston: MIT Press.

Mehan, Hugh and Wood, Houston (1975), *The Reality of Ethnomethodology*, John Wiley and Sons.

Schelling, Thomas (1962), *The Strategy of Conflict*, Cambridge: Harvard Press.

Skvoretz, John and Fararo, Thomas (1996), "Generating Symbolic Interaction: Production System Models," *Sociological Methods and Research*, 25(1).