

MODEL DESCRIPTION

This description of the agent-based model introduced in Premo (2016) follows the ODD protocol (see Grimm et al. 2010).

Purpose

This model illustrates how the effective population size and the rate of change in the mean skill level of a cultural trait are affected by the presence of natural selection and/or the cultural transmission mechanism by which the trait is passed. The model, which combines aspects of Shennan's (2001) and Henrich's (2004) influential selection models, is used to address the following research questions:

How does census population size affect the rate of change in mean skill level of a cultural trait under different versions of directly biased oblique cultural transmission?

How does increasing the proportion of directly biased oblique transmission relative to vertical transmission affect the effective population size of a culturally transmitted trait in the presence or absence of natural selection?

Holding all else constant, does the presence of at least some vertical transmission (in addition to direct biased oblique transmission) require a larger population in order to support adaptive cultural evolution (i.e., a positive rate of change in mean skill level)?

Entities, State Variables, and Scales

Agents represent social learners. Each agent possesses four state variables: age, x , times-taught, and parent. age=0 signifies that the agent is a member of the naïve (or learner) generation. Naïve agents have not yet acquired a value for x , and thus naïve agents cannot serve as teachers. age=1 signifies that the agent is a member of the parent (or teacher) generation. Only members of the parent generation can serve as parents and/or teachers. x is a floating point value that can be thought of as the skill level of a culturally transmitted trait (as per Henrich 2004). times-taught is an integer representing the number of times that the agent served as a teacher when a member of the parent generation. Each agent also has a variable, cleverly called parent, which allows it to remember and identify its biological parent.

Global variables include: N , seed, natural-selection, xMax, z , direct-biased-version, α , and β . N is census population size. seed provides the integer used to initialize the random number generator at the start of each simulation run. When natural-selection=false, there is no natural selection in the model. When natural-selection=true, x is under natural selection. The maximum x value in the parent generation is recorded in xMax at the start of each time step. The probability that a naïve agent will acquire x via oblique cultural transmission rather than via vertical transmission is given by z . As explained in the text (and below), the model allows for two different versions of directly biased oblique cultural transmission. In Shennan's version, the naïve agent attempts to copy the highest x value displayed among 5 randomly chosen members of the parent generation, excluding the biological parent of the naïve learner. In Henrich's version, the naïve agent attempts to copy the highest x displayed among the entire parent generation. The Gumbel distribution is characterized by two parameters: α and β . α represents the difficulty of precisely copying the target value. β represents the variation (or spread) that results from

attempting to copy the target value. Please see Henrich (2004) for more discussion of α and β , including a helpful illustration.

Each tick represents one discrete, non-overlapping “generation” of social learning, whereby every member of the naïve generation has a chance to acquire its x value by learning from the parent generation. The model has no spatial scale. Although all agents have spatial coordinates (x_{cor} and y_{cor}), the coordinates are assigned randomly in all cases and at no point are they used in the model.

Process Overview and Scheduling

Each time step includes the following stages in the following order. Note that all agents complete each stage before any agents begin the next one.

1. Age: The N agents present at the beginning of the time step move from the naïve generation to the parent generation by changing their state variable, age, from 0 to 1.
2. Create naïve generation / vertical cultural transmission: Members of the parent generation (age=1) give rise to N naïve agents (age=0) either in the presence or absence of natural selection. During this process each offspring inherits its parent's x as its target value through vertical transmission.
3. Directly biased oblique cultural transmission: with probability z , a learner replaces the target value based on its biological parent with the x value transmitted from some member of the parent generation via either Shennan’s or Henrich’s form of directly biased oblique cultural transmission. Finally, the offspring adopts its x value from the Gumbel distribution “anchored” to its target value.
4. Update data: compile data from the current time step.
5. Cull parent generation: all members of the parent generation (age=1) are removed from the simulation.
6. Plot a simple line graph of the mean and maximum x values in the population at the end of the time step.

When using the experiments provided in BehaviorSpace, each simulation ends after 1000 time steps. Summary data are collected at that time.

Design Concepts

Basic principles

Naïve individuals acquire a continuous trait (x) via cultural transmission from the experienced, or teacher, generation. A proportion $1-z$ (on average) of naïve agents obtains x via vertical transmission while a proportion z (on average) of agents obtains x via one of two forms of directly biased oblique cultural transmission. Census population size

remains constant throughout the simulation run. Noise in cultural transmission (both vertical and oblique) is defined by the Gumbel distribution, α , and β .

Objective

The basic objective of each agent is to acquire x through social learning. Under some forms of cultural transmission, such as directly biased oblique transmission, the objective is to copy the highest x value among those known to the agent.

Sensing

Under vertical transmission, each member of the naïve generation learns from its parent. In this case, each agent can only sense its biological parent as a potential teacher. However, under directly biased oblique transmission, each naïve agent learns from an agent chosen out of a set of potential teachers. Under Shennan's version of directly biased oblique transmission, each member of the naïve generation has a set of potential teachers composed of a randomly chosen subset of 5 members of the parent generation, excluding the naïve agent's biological parent. Under Henrich's version of directly biased oblique cultural transmission, each naïve agent learns from the member of the parent generation that displays the highest x (i.e., the potential set of teachers includes the entire parental generation).

Interaction

Agents with $\text{age}=0$ interact with their parents and teachers. Naïve agents ask their teachers to increment times-taught to reflect cultural transmission of x . A naïve agent who learns obliquely from a teacher rather than vertically from its parent asks its parent to decrease by 1 its times-taught to reflect the fact that the naïve agent's target value is no longer its parent's x value. Agents do not interact with the physical environment.

Learning

Naïve agents ($\text{age}=0$) learn socially by attempting to copy the x value displayed by a teacher. Transmission is imperfect. The noise in learning is defined by the Gumbel distribution and two parameters: α and β .

Stochasticity

There are many stochastic processes in the model. The Gumbel distribution provides stochastic “noise” in social learning. Each naïve agent compares a real value drawn from a uniform distribution bound by 0 and 1 to z , which represents the probability (per agent) of engaging in directly biased oblique transmission instead of vertical transmission. Under Shennan's version of directly biased oblique cultural transmission, 5 possible teachers are chosen randomly from the parent generation (excluding the biological parent of ego). In the presence of natural selection, each potential parent (a member of the parent generation chosen at random and with replacement) compares its relative fitness to a real value chosen randomly from a uniform distribution bound by 0 and 1. In the absence of natural selection, parents are chosen randomly and with replacement from the parent generation. As a result, although all members of the parent generation have the same expected number of offspring (1) in the absence of natural selection, the actual

number of offspring per member of the parent generation is Poisson distributed with a mean and standard deviation of 1.

Collectives

There are 3 different agent collectives in the model. The parent generation is comprised of agents with age=1. The naïve generation is comprised of agents with age=0. Under Shennan’s version of directly biased oblique cultural transmission, each naïve agent draws its teacher from a third type of collective, called possible-teachers. possible-teachers refers to a subset of the parent generation composed of 5 randomly chosen individuals excluding the naïve agent’s biological parent. Note that this means that the size of the possible-teachers collective is just 4 rather than 5 when $N=5$. For $N>5$, each naïve agent randomly draws 5 non-biological-parent members of the parent generation. Members of the parent generation can be included in more than one naïve individual’s possible-teachers collective.

Observation

Data are “tallied” and displayed at the end of each time step (some of these data are then collected by BehaviorSpace after 1000 time steps). At the end of each tick, the following data are collected: the number of individuals in the parent generation that served as a teacher at least once during the time step, the mean number of times that each member of the parent generation served as a teacher, and the variance in the number of times that each member of the parent generation served as a teacher. Running “tallies” or sums are kept for each and for each an average is calculated by dividing the running sum by the number of ticks elapsed. It is the sums averaged over 1000 ticks that are collected at the end of each simulation run. To calculate the rate of change per generation in mean skill level ($\Delta\bar{x}$), I divide the population mean of x after tick 1000 by 1000.

Initialization

Each simulation run is initialized with a population of N social learners, each of which displays $x=0$, age=0, times-taught=0, and a randomly assigned spatial coordinate. Table 1 provides the parameter values used in the paper.

Table 1. Parameter values used in the experimental design employed in Premo (2016).

β	1
α	1, 2, 3, 4, 5, 6
z	0, 0.2, 0.4, 0.6, 1
N	5, 15, 25, 35, 45, 55, 65, 75, 85, 95, 105
natural-selection	True, False
direct-biased-version	Shennan, Henrich
seed	1-20

Input Data

The model does not make use of input files.

Submodels

1. vertical-cultural-transmission

This submodel is responsible for creating the naïve generation of N social learners and passing the target values of x through vertical transmission at the same time. Members of the parent generation reproduce asexually until N new agents have been created. While there are fewer than N naïve agents, a member of the parent population is sampled (with replacement). This agent is given the opportunity to reproduce. Note that reproduction can occur in the absence of natural selection (natural-selection=false) or in the presence of natural selection (natural-selection=true).

In the absence of natural selection, the sampled individual reproduces with probability 1. After reproduction, the sampled individual is returned to the population and the process described above is repeated until N naïve social learners have been created. Note that the same agent can be chosen to reproduce more than once and that not all agents will be chosen due to the sampling error introduced by randomly drawing from a finite population.

In the presence of natural selection, the probability that the sampled agent reproduces is given by its relative individual fitness. Individual fitness is scaled relative to x_{Max} , the maximum x value displayed in the parent generation during that time step. Let x_i represent the sampled agent i 's x value. Following Shennan (2001), the relative individual fitness of the i th agent is $\exp(x_i - x_{\text{Max}})$. If a real number chosen randomly from a uniform distribution bound by 0 and 1 is less than the agent's fitness, then the agent reproduces. If the real number is greater than the individual's relative fitness, then the agent does not reproduce. In either case, the sampled individual is placed back into the population and the process described above is repeated until N naïve social learners have been created.

Each offspring is given an age=0, a random spatial coordinate, and memory of its biological parent. Each offspring also automatically inherits its parent's x value as its target value (at least for the time being).

2. directly-biased-cultural-transmission

This submodel is responsible for the transmission of the target values of x via directly biased oblique cultural transmission. Recall that z provides the probability (per naïve agent) of acquiring x via directly biased oblique transmission rather than vertical transmission. If a real number chosen randomly from a uniform distribution bound by 0 and 1 is less than z , then the naïve agent will learn x via directly biased oblique transmission.

If direct-biased-version=Shennan, then the naïve individual must go about selecting its teacher from the parent generation. According to Shennan's (2001) version of directly biased oblique transmission, each naïve individual tries to copy the highest x value displayed among a randomly chosen subset of 5 members of the parent generation,

excluding the naïve individual's biological parent. After first removing its biological parent from consideration, the naïve individual randomly chooses 5 agents from the $N-1$ remaining members of the parent generation to serve as members in its subset of possible-teachers. The naïve individual then selects the individual that displays the highest x value among the members of its possible-teachers collective to serve as its teacher.

If direct-biased-version=Henrich, directly biased oblique transmission is even simpler. Under Henrich's version of directly biased oblique transmission, naïve learners are able to attempt to copy the member of the parent generation who displays the maximum value of x .

The final component of this submodel deals with copying error—the imperfect inference or “noise” involved in cultural transmission. Note that copying error is operationalized the same regardless of whether the target value is obtained via directly biased oblique transmission or vertical transmission. In both cases, the x that is ultimately adopted by the naïve learner is a function of the target value it is attempting to copy (i.e., the target value is the x displayed by its teacher) and the Gumbel distribution characterized by α and β (again, see Henrich [2004]). For example, if the target value is 1.876, the x value the naïve learner will ultimately adopt is generated by:
 $(1.876 - \alpha) - (\beta * \ln(\text{random-exponential } 1))$.

References Cited

Grimm, Volker, et al. (2010) The ODD protocol: a review and first update. *Ecological Modelling* 221:2760-2768.

Henrich, Joseph (2004) Demography and cultural evolution: how adaptive cultural processes can produce maladaptive losses: the Tasmanian case. *American Antiquity* 69:197-214.

Premo, L. S. (2016) Effective population size and the effects of demography on cultural diversity and technological complexity. *American Antiquity*, in press.

Shennan, Stephen J. (2001) Demography and cultural innovation: a model and its implications for the emergence of modern human culture. *Cambridge Archaeological Journal* 11:5-16.