

Appendix SA – Simulation Model of the Social Protocell

Contents

Introduction.....	1
Global level.....	2
Space	2
Time.....	2
Resources	3
Community-level	4
Territorial behavior.....	4
Community fission.....	5
Community death.....	6
Agent-level	6
Ageing.....	6
Death	6
Reproduction.....	7
Behaving socially	7
Learning independently.....	7
Learning socially	Fel! Bokmärket är inte definierat.
Expressing tradition alleles.....	7
Institutional activities.....	9
Harvesting resources and the effects of traditions in general	10
Non-resource functions available as targets for apex loci	11
Social learning rate boost.....	11
Innate capacity for social learning.....	11
Bibliography.....	13

Introduction

We here expand and complement the main article’s brief description of the simulation model used for producing the results. We assume that the reader has first read the main article model description. See Appendix B for discussion about parameters and stability analysis.

The model describes a much-simplified setting that we argue may abstractly represent an early *Homo* environment (ca. 2.0-3.0 Mya), implementing the basic entities and processes of the Social Protocell Hypothesis (SPH; see Andersson & Törnberg, 2019; Davison et al., 2021). The model has proximate as well as more more ultimate purposes. Its proximate purpose is to explore the behavior of hypothetical social protocells to examine whether the behavior that has been proposed is in line with what we observe in the model. The more ultimate purpose is to serve as a platform that can be extended to test the SPH

in many more ways in the future, but also to provide ideas for simpler models (see also *Summary* in Appendix SB and *Conclusions* in the main article).

Figure A1 provides a visual map to the ontology of the model, and of the contents of this appendix.

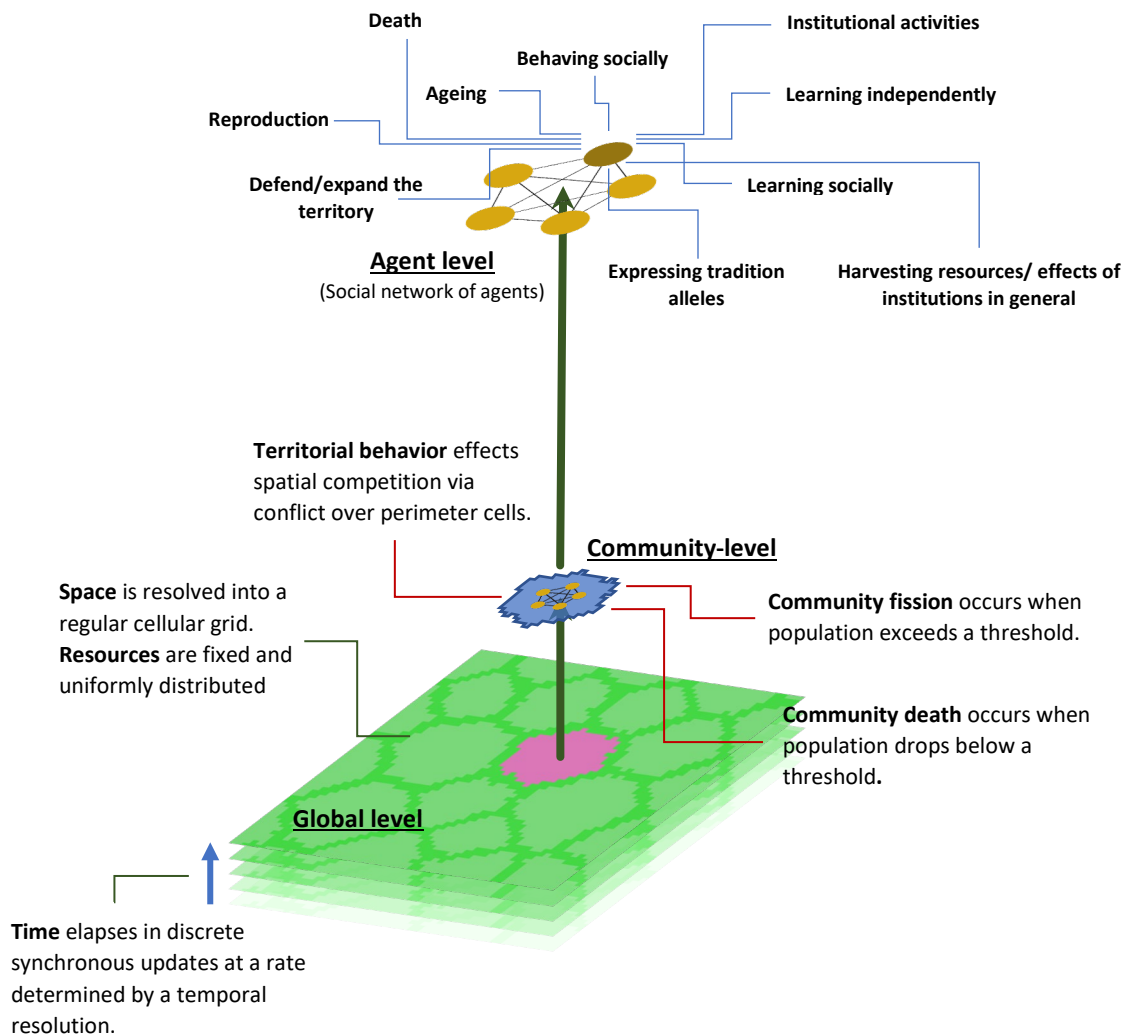


Figure A1 – Visual representation of the overall model ontology, references to processes (and subsections below) indicated in boldface.

Global level

Space

The model simulates a toroidal area (square with periodic boundary conditions) of unit size, resolved into a regular grid of *area elements*. The resolution is controlled by a parameter (see *Varying the spatial resolution* in Appendix SB). Territorial competition plays out over this area, which also contains *resources* but no other features in this setup. Area elements are either controlled by a community as part of its contiguous territory or it is uninhabited.

Time

The temporal resolution is controlled with a parameter (see *Model updates per “year”* in Appendix SB), and the effect is a scaling of all rates (e.g., in learning, territorial competition, and so on.) Time begins at the seeding of the configuration with a single community.

Resources

Resources are uniformly distributed across the area and are fixed rather than regrowing. We judge that this does not affect the qualitative dynamics and thereby do not implement dynamical resources since this would substantially increase the number of parameters and the computational cost of running the model (which is already high).

The resources are defined by their size, degree of monopolizability, and degree of access by harvesting, as a function of the efficiency of the harvesting strategy (see Appendix SB). For example, a somewhat complex strategy is needed for accessing nuts, but once the nuts can be cracked, additional sophistication yields little additional benefit. Nuts are also highly monopolizable as they can be immediately consumed by whoever cracked the nut. By contrast, for large carcasses (the IGUT/carnivory resource), complex strategies may gradually yield additional benefits almost indefinitely. Large carcasses are moreover not just hard to monopolize, they are also less meaningful to monopolize. A single agent can only consume fractions of a large animal alone anyway.

Monopolizability in the model is a real valued parameter in the unit interval. To an extent that depends on the resource's degree of monopolizability, resources obtained by one agent become available also to others in the community. The agent that expresses the apex allele of an institution used for accessing a resource first receives the monopolizable part of the payoff. The non-monopolizable part of the payoff is then distributed equally to all other agents belonging to the same community.

The basic resource

We pool together all opportunistic resources in a class of resources that can be obtained with little or no reliance on social learning. In modeling terms, this means that the resource is relatively large (the agent can get by accessing only this compound resource) and already at zero institutional complexity there is a substantial level of access to the resource. The level of access then increases rapidly until it levels off. For example, anybody can figure out how to pick berries, and while you may gain from learning a few tricks, additional tricks will soon not be overly worthwhile. This represents a state within the expected range of a Last Common Ancestor as a generic chimp-like species of great ape. This resource is not commented on in the article and is mostly there to enable agents to survive also without targeting the second resource which is in focus.

The IGUT/carnivory resource

The carnivory institution plays the lead role as a cultural adaptation for obtaining high-quality food, inherited via the emergent heredity channel that we have described. In the model, there is no access to this without any support by social learning. A single tradition locus without any supporting tradition loci (complexity level 1) provides a minimal access, but since the resource is large (see *Size of the IGUT (carnivory) resource* in Appendix SB), even a small level of access is of importance. The big difference from “the basic resource” is that access keeps increasing steadily if additional supporting tradition loci are added in the proto-institutional structure.

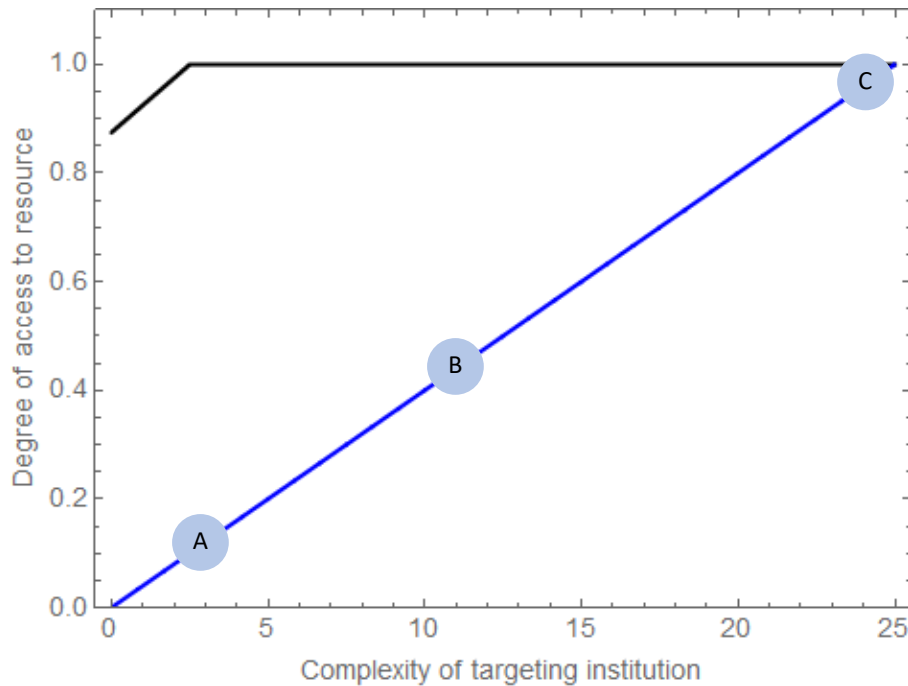


Figure A2 – (A) The blue plot represents level of access to the IGUT/carnivory resource afforded by institutions. Access is zero without the aid of culture, and then increases gradually as complexity increases, beyond the level that will be reached during simulations. The black plot represents the compound basic resource. That is, opportunistic resources that can be efficiently accessed without the use of social learning, but that gain slightly by a moderate addition of sophistication before leveling off indefinitely. Access says nothing about the size, so given that the IGU is very large, even a small level of access may rival the basic compound resource. The level of access to the IGUT/carnivory resource achieved by the traditions illustrated in Figure 3 (top row) in the main article (and Figure A5 below) are indicated by the circles. We may note that for the basic resource, all that increase in complexity would yield no increased access.

The carnivory resource is referred to as IGUT because it corresponds to the type of tradition that was described as needed by Andersson and Törnberg (2019) – i.e. an Important, Generative and Universal Tradition. This means a large resource that may be *important* already at a minimal level of access, which can then be *generatively* expanded almost endlessly, remaining thereby important and a steady source of possible improvements. *Universal* means that it is widely geographically distributed so that a large number of communities targeting it may exist. See *IGUT (carnivory) access function shape* in Appendix SB for more details on the access function.

Community-level

Territorial behavior

Communities collectively defend non-overlapping territories. This is one out of two coordinated activities that are emulated directly on the community level. While we avoid doing so in general, this is an unremarkable type of behavior, and simulating it partly on a higher level crucially saves a lot of computational complexity. Area elements at the perimeter of communities are the scenes of this territorial competition (see Figure A3a).

The count of agents above a threshold age is used as a simple measure of competitive power. This measure is aggregated in the geographic center of the territory, and then extended as a field that tapers off as the square root of distance from the center. The strength of this strength field in the perimeter cells around the community is then the competitive power of the community. The effect is that circular shapes are favored and that the strength of a community in its perimeter increases the more compressed the community becomes. Deaths are not modeled, only territory changing hands.

Each perimeter cell will attack one randomly selected outgroup perimeter cell per update. Competition happens by flipping the possession of these border elements with probabilities computed as a function of differences in this power field (see Figure A3b). Hence, if the population of a community increases, the pressure that it exerts will increase, and with it the likelihood of winning additional cells, which, in turn, decreases the pressure by expanding the territory; see also *Probability of attacking cell taking over a neighbor cell* in Appendix SB. Territories will typically be in an uneasy equilibrium in an approximately hexagonal pattern in space. The behavior of the model and its formulation represents standard techniques for simulating spatial competition (see e.g. Czárán, 1998, Chapter 6).

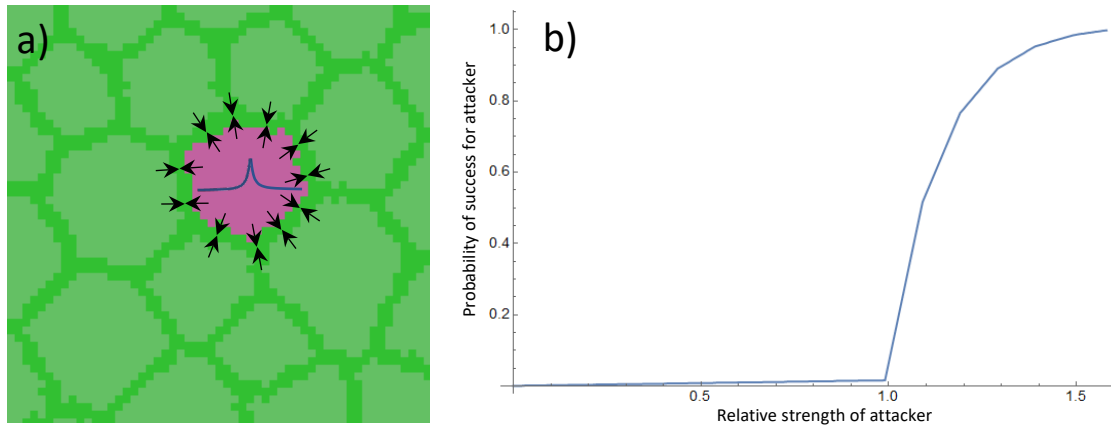




Figure A3 – The community controls a territory with agents contributing a unit of *strength* each to defense and opportunistic expansion. At the center, $pressure=strength$, and $pressure$ then decreases as the square root of distance to the center. Competition is modeled by a fictive attack by all border cells against one neighboring outgroup-controlled or empty cell per update (it may have several). The outcome is decided using a designed piecewise linear function of the pressure difference between attacker and defender. If the attack succeeds, the attacked cell changes ownership. If the cell is empty, there is an ambient pressure that represents competition with other species. Above are shown (a) an illustration of the process, and (b) a plot of the mapping between the relative strength of attacker and defender, and the probability of conquering the attacked cell.

Community fission


Communities undergo *fission* irreversibly at a rate that is a function of agent population size (see also *Community splitting population threshold* in Appendix SB). In this setup, we use a delta function where communities always split when exceeding a size of fifty members, otherwise they do not split. Here in place of an actual simulated quarrel between agents. The split produces two offspring communities and splits the network into two separate networks: half the agents go to each of the daughter communities (see also *Mode of community fission* in Appendix SB). The communities also split the territory evenly between them, which happens by choosing a random point on the interface of the community, and then finding an opposite point by “walking” half the perimeter length around the perimeter. The territorial competition dynamics will then relax their spatial extent over time; see Figure A4.



Territory of a community near the agent population threshold of undergoing fission.



Next update, the threshold is exceeded, and fission has occurred.



Five updates later, the shapes of the new communities and their neighborhoods have begun to relax and ease into the overall pattern.

Figure A4 – A community undergoing fission.

Community death

Communities disperse when their populations or sizes decrease below certain thresholds; see *Community dispersal threshold – population* and *Community dispersal threshold – area* in Appendix SB). When a community disperses, its agents die, and its territory turns into uninhabited land.

Agent-level

The inside of the community simulates the life of a social group of agents in a highly stylized way. Agents interact amicably within, but not between, communities. This social network topology is simulated by giving all agents within the community full access to the actions of other agents in the community, but no access between communities. No transfers are allowed between communities.

Agent life is resolved into synchronous update cycles during which a certain amount of time elapses, in an absolute unit (termed “year”) during which they:

- Age.
- May die.
- May reproduce.
- Defend/expand the territory (community-aggregated, see *Interactions between communities*.)
- Behave socially, which may lead to fission (community-aggregated, see *Community fission*.)
- Learn tradition alleles independently (i.e., invent new traditional species.)
- Learn tradition alleles socially (subject to learning biases.)
- Express their tradition alleles, invoking thereby their effects.
- Take part in institutional activities (emergent effect of expressing tradition alleles.)
- Harvest resources and generally become subject to effects of traditions used in the community.

See also Figure A6 for an overview of some of the processes involving traditions.

Ageing

Agents have an age that increases at a rate determined by the temporal resolution parameter. Until an (arbitrary) age of 8 “years,” agents do not reproduce and do not contribute to the strength of the community in territorial competition. They harvest resources but will frequently not have had time to gain alleles and their contribution will thereby be small. Age also affects the death rate.

Death

The likelihood of dying per unit time is generated from one multiplicative factor for nutrition-related causes of death, and one for age-related causes. Both these factors are generated via a piecewise linear function (same as in the mapping from relative strength to likelihood of conquering a neighboring cell, see Figure A3b) that maps agent energy level and agent age, respectively, to these factors. In *Agent*

survival as a function of age (Appendix SB) we discuss this parameter further and verify that the choice of the shape of these functions are not critically important.

Reproduction

Agents reproduce when a threshold surplus has been accumulated, subject to a minimum separation between births. Reproduction is sexual and begins at an arbitrarily chosen age of eight “years.” When reproduction happens, another agent over the age of eight is chosen randomly in the community population as the mate. Genetic inheritance involves, first, mutation of parents’ genes, and, second, particulate crossover between them so that the offspring inherits genes either from one or the other parent. However, only one feature is genetically inherited (cognitive capacity), and this feature varies only in some scenarios. Agents are otherwise genetically inert.




Figure A5 – Visualization of institutions of increasing complexity from the left to the right. The large green circles are apex loci, while the grey circles represent component loci that together boost the function of the apex. The complexity classes are, from left to right: 3, 11 and 24. The topology of the trees could be subject to selection (by affecting function), but that is not the case in the present setup.

Defend/expand the territory

Agents that are “of age” contribute to an aggregated measure of community strength. We assume that this group behavior is stable and innate. See Section *Competition between territories*.

Behaving socially

Except for social learning (see below), social behavior is modeled in a highly aggregated fashion. Social interactions will lead to fission if the group exceeds a threshold membership size (see *Community fission*), and agents collectively defend the territory (see *Defending and expanding the territory*).

Expressing tradition alleles

Each update, the agents express all the tradition alleles they have learned during their lifetimes. This **(i)** produces their effects, and **(ii)** makes them available for social learning. Their effects in the community are accounted for by tallies of how many agents that have expressed alleles for each present locus. For example, if a locus #126 is expressed by twenty-nine agents possessing alleles of it, then we have twenty-nine effects of this locus logged.

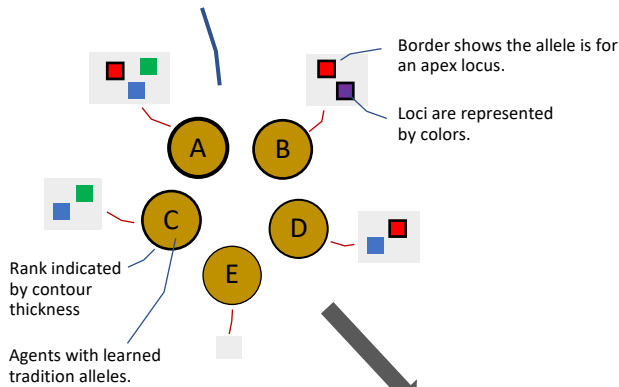
However, there is more to the function of loci than simply expressing the allele. They may have dependencies that must be satisfied. This must be recursively resolved, which we discuss next.

Learning independently


By independent learning we mean modifications of behavior invented by agents independently.

First, what we refer to as *innovation* gives rise to a new component locus that extends an existing institution by linking to some existing component or apex locus. Such new component loci get linked into the tree illustrated in Figure A5. Creativity can also alter the topology of these trees by changing the linkage of an allele (which gives rise to a new locus), but this has no effect in the model formulation used here since it is only the gross complexity of the tree of components under an apex tradition goes into computing functional performance.


Five agents in a community, expressing their tradition alleles.



Linkages between loci define two institutions, one of which is unlinked.



Finding complete expressions of institutions.



Social learning: Each expressed allele is a learning opportunity.

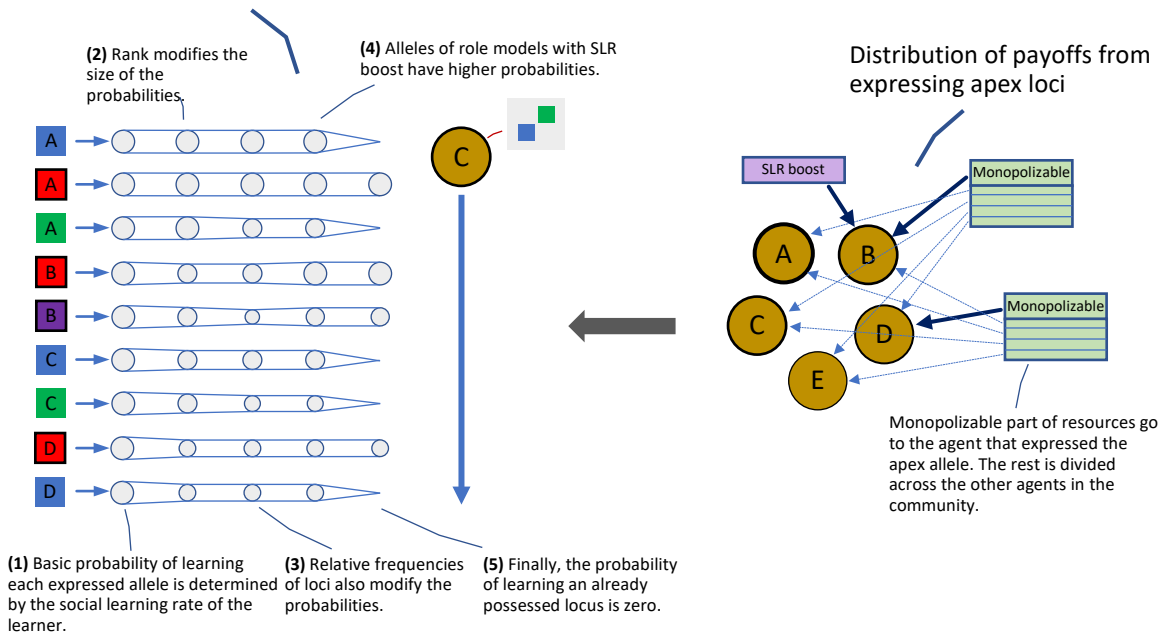


Figure A6 – Map over some of the operations performed on traditions in the model.

Second, what we refer to as *learning* gives rise specifically to new apex loci linked to one of the potential functionalities that exist in the system, such as a resource or some other function (in this setup, only social learning rate boosting).

Both of these sources of novelty are associated with rate parameters, see also *Agent innovation rate* and *Agent learning rate* in Appendix SB.

Learning socially

In *social learning* an agent obtains an allele of some locus by observing another allele being performed by another agent. Notably, there is *no notion of copying* the know-how underpinning the behavior in question. What we assume is that the functionality of the role model's behavior is reconstructed by the learner. All alleles therefore have the same unit level of complexity (Tennie et al., 2009; i.e., the level of their creative capacity, corresponding to some Zone of Latent Solutions; see 2020), as well as the same function and level of performance – either you can perform the function of the locus or you cannot. No improvement or deterioration is inherited in this process.

It should be noted that the model does not *require* this limitation to function. It is in place because we are testing whether cumulativity can happen on the higher level (sociant) even if it does not happen on the lower level. We therefore rule out cumulativity on the level of social learning.

Agents are willing to learn components only if they are linked to other traditions (component or apex) that are present in the community. We assume that “orphan” component loci are pointless (say, collecting raw material that is not used.) Apex traditions produce some tangible function that we deem is sufficient for motivating the learning to take place independently. As we discuss in Appendix SB (*No bias against component traditions*), the assumption that agents are willing to learn component traditions must be viewed as a strong assumption since these would not produce immediate rewards.

There is a basic rate of social learning (see Figure 7, main article) that is modified by biases for the rank of the agent expressing the tradition (see *Rank/prestige bias temperature*, Appendix SB), and the relative frequency of the tradition in the community (see *Conformity bias temperature*, Appendix SB). For example, without any bias, if fourteen agents express alleles of locus #126, and seven agents express alleles of locus #31, then the basic chance of learning #126 is double that of learning #31, since twice as many role models provide twice as many opportunities to learn. Conformity bias makes lineage #126 *more than* twice as likely to be learned. This likelihood is then further modified by the relative ranks of the agents behind the expression count.

The rank/prestige and conformity biases are computed via transformation ($e^{-\frac{1}{T}f}$) where the parameter T works like a temperature, controlling how much attention agents pay to the factor f in question. Many factors potentially go into the choice of role models, and the question is what the relative importance is of the rank/prestige measure *in this mix*. At a high temperature, other factors swamp the prestige factor, which, since we do not model these other factors explicitly, are approximated as noise. At a low temperature, the prestige factor dominates, which is modeled as a reduction in the influence of these other random factors.

The biases then multiplicatively modify the probabilities of learning. The rank factor is computed by sorting the agents according to their lifetime average extraction of food (which is a somewhat arbitrary measure). The conformity factor is computed as the relative frequency of the locus in question in terms of what fraction of the community population that possesses it.

Institutional activities

Tradition loci have specific functions and sets of dependencies. For clarity we recapitulate the main article, we differentiate, first, between apex and component traditions. Apex traditions represent behavior that directly produces an external functionality. By “external” we mean in contrast with the “internal” functions of component traditions, which do not produce any benefit in themselves. Component traditions belong to hierarchical trees of subfunction under an apex tradition (see illustration in Figure A5).

Imagine that the apex tradition's function is to obtain food from a carcass. Then component traditions under such an apex may then be imagined as things like the production of tools for doing this, finding carcasses, defending the carcass, obtaining raw material for producing the tools, transporting parts of the butchered carcass, and so on.

Next, we differentiate between *loci* and *alleles*, where the former are *types* of traditions, and the latter are *instances* of such types. Alleles correspond to a locus and an agent may possess one or zero alleles of each locus.

Apex traditions without dependencies to components are standalone in the style of animal traditions. They can be performed for effect without regard to other traditions. They have unit complexity in the model. If an apex locus has dependencies, we call the apex along with its tree of dependent component loci an *institution*.

Institutions have a compound complexity that may lead to higher efficiency than a standalone tradition can (such as for harvesting a specific resource). How functional efficiency increases with complexity varies between different functions, such as between different resources, and is specified as a parameter using a piece-wise linear function (see Figure A2).

The set of external functions available for apex loci to target are pre-identified and includes types of resources, as well as custom functions that affect the lives of the agents in other ways. In this case, we have one such additional function, namely a type of institution that affects the rate of social learning.

The tallies produced via expression of alleles (see *Expressing tradition alleles*) are used as follows for detecting full expressions of institutions:

Alleles without upstream dependencies are satisfied already by being expressed – they do not need alleles of other loci to function. Alleles with upstream dependencies must be resolved, however. For example, if an expressed allele of locus #126 demands an expression of an allele of locus #31 to function, we see whether the tally of satisfied expressions of lineage #31 is larger than zero. If it is, the expression of this allele of locus #126 is also satisfied, and we decrease the tally of expressions of locus #31 by one since it is now “consumed.” This is carried on until all expressions that can be matched have been matched.

The right-most example in Figure A5 has a complexity of twenty-four, which is a very high number in the context of the present setup. This may produce a high efficiency, but it also produces brittleness. If the tally for *any* one of its twenty-three components is zero, it has failed and will produce nothing.

Harvesting resources and the effects of traditions in general

Harvesting is partly a community-level activity since the model assumes that food is shared within the community as a tight social network, but we treat it as an agent-level process since it is only the distribution that takes place on the community level. The mechanisms for sharing that we have in mind are simple, such as tolerated theft and reciprocity (e.g. Jaeggi & Gurven, 2013) and the degree to which distribution occurs is assumed to vary with the degree of monopolizability of the resource. The logic is that the importance of tolerated theft increases if large parcels of food are harvested where it is not only hard to keep the food for oneself, but also less important to do so since the parcel is sufficiently large to feed everybody in any case (Blurton Jones, 1984; Hawkes et al., 1993; Winterhalder, 1996).

Resource availability is proportional to territory area per agent, which is computed simply by dividing the territory area of the community with the number of agents in the community. As we have mentioned in the context of resources, the agent that happens to express an apex allele (green circles in Figure A5) obtains the benefit derived from the institution as a whole. This agent keeps the monopolizable part of the parcel, and the rest is then distributed to the community (see *Monopolizability of the IGUT (carnivory) resource* in Appendix SB). This does not mean that we assume that everybody is ensured

an equal share (which would suppose some mechanism), but rather that we do not assume any accounting by the agents of who did what.

Non-resource functions available as targets for apex loci

Any number of non-resource functions can be added by customizing the model to use the fitness computed for apex alleles. In this setup, we use only a single such additional function, namely a social learning rate boosting institution. Potential other such functions could include punishments of cheaters, ability to defend the territory, and so on. These mentioned two potential functions have been tried in preliminary versions but have not been included in the setup used in this paper.

Social learning rate boost

Social learning rate is controlled by a parameter (see Figure 7 in main article). In the agents individually, this parameter is then modified by learning biases before being applied to learning opportunities as expressions of alleles in communities (see *Learning socially*). However, to test whether institutions with more *indirect* fitness contributions can evolve, we have also used an institution that increases the rate at which social learning takes place. More specifically, this institution is intended to represent *evolution of evolutionary individuality* as proposed by Davison et al. (2021). That is, an institution whose effect is to improve the efficiency of evolution on the sociont level, which is central to the idea of a cultural ETI.

The boost is implemented such that any tradition allele held by the agent expressing the apex locus of such an institution will have an increased likelihood of being learned by other agents. In other words, it functions as an additional learning bias. It is indicated when this institution is active in the Results of the main article.

In reality we do not thereby imagine that institutions dedicated solely to improving learning would have emerged. It seems more plausible that improvements of social learning would initially have been side-effects of institutions mainly focused on other tasks. What we wish to test is, however, whether such functions of institutions may be selected for, and assuming those to be separate units does not constitute a weaker test.

Innate capacity for social learning

Along with the institution for boosting social learning (see *Social learning boost*), we also use a varying genetically transmitted potential to benefit from such an institution (see *Reproduction*). In such a setup, agents will only benefit *up to a point* by learning from an agent with culturally boosted performance as a role model. As an illustration of a limit case, consider how human institutions presuppose a high level of cognitive capability in humans. Modern human institutions would be meaningless to, say, a chimpanzee. The point is that there must be a co-evolution between an innate capacity and institutions that harness such a capacity. We hereby enable the model for testing whether such co-evolution can take place in principle in a simplified case (see Figure A7).

This cognitive capacity is genetically inherited and comes at a cost for the agent (see *Cognition cost for agents* in Appendix SB). In other words, higher cognitive capacity in itself is maladaptive to agents in this setup. It is only adaptive to the extent that possessing such a capacity will benefit the agent via an ability to partake in institutional organization that *in turn* improves social learning. The utility derived is then whatever else the agent may thereby learn by from having such capacity in the presence of a social learning booster institution of commensurate complexity.

In reality, there are agent-level benefits that may further favor such evolution. For example, if high intelligence provides high performance, and high performance then yields high prestige, and high prestige, in turn, translates to better access to mating. In order not to stack the model in favor of the hypothesis we do not model such potential adaptive uses.




Figure A7 – The genetically transmitted cognitive adaptation determines the *actual* boost that an agent experiences when faced with a tradition allele displayed with a specific institutional boost. Institutional boost without the effect of cognition is here displayed in the blue plot. The solid red line shows $COG_{learner}=1$, the dashed line $COG_{learner}=2$ and the dotted line $COG_{learner}=3$. The function used is designed to trace the institutional boost, but with a deviation downward that grows faster the lower the cognitive capacity is, eventually peaking and decreasing as the institution is beyond what the agent can make sense of. Two scenarios are shown. By **(A)** we see a situation where an increase in the institutional complexity class in the role model provides a large boost while an increase in the cognitive capacity of the learner makes little difference. The institution can here be elaborated within the capacity of the agents. In situation **(B)**, the institution strains the capacity of the agents, so an even more complex institution would have a negative effect, while an increase in the cognitive capacity yields a large positive effect. The dynamic that we seek is that increasing the complexity of institutions creates opportunities for higher cognitive capacity such that the only way to reach high levels of performance is via co-evolution between the cultural and the genetic adaptation.

The actual social learning boost of the learner is modeled using the map

$$SLR_{learner} = SLR_{rolemodel} - \frac{2^{SLR_{rolemodel}} e^{-2 COG_{learner}}}{\text{Log } 2}$$

Where $SLR_{rolemodel}$ is the institutional boost provided to the role model (see Figure A5), $COG_{learner}$ is the level of cognition of the learner, and $SLR_{learner}$ is the actual social learning boost in the learner. This boost is applied as a factor to the learning probabilities attached to all tradition alleles present in the community.

Bibliography

- Andersson, C., & Törnberg, P. (2019). Toward a Macroevolutionary Theory of Human Evolution: The Social Protocell. *Biological Theory*, *14*(2), 86–102. <https://doi.org/10.1007/s13752-018-0313-y>
- Blurton Jones, N. G. (1984). A selfish origin for human food sharing: Tolerated theft. *Ethology and Sociobiology*, *5*(1), 1–3. [https://doi.org/10.1016/0162-3095\(84\)90030-X](https://doi.org/10.1016/0162-3095(84)90030-X)
- Czárán, T. (1998). *Spatiotemporal models of population and community dynamics*. Springer Science and Business Media.
- Davison, D. R., Andersson, C., Michod, R. E., & Kuhn, S. L. (2021). Did human culture emerge in a cultural Evolutionary Transition in Individuality? *Biological Theory*.
- Hawkes, C., Altman, J., Beckerman, S., Grinker, R. R., Harpending, H., Jeske, J., Peterson, N., Smith, E. A., Wenzel, G. W., & Yellen, J. E. (1993). Why hunter gatherers work. an ancient version of the problem of public goods. *Current Anthropology*, *34*(4), 341–361. <http://www.jstor.org/stable/2743748>
- Jaeggi, A. V., & Gurven, M. (2013). Reciprocity explains food sharing in humans and other primates independent of kin selection and tolerated scrounging: a phylogenetic meta-analysis. *Proceedings of the Royal Society B*, *280*(1768), 20131615. <https://doi.org/10.1098/rspb.2013.1615>
- Tennie, C., Bandini, E., van Schaik, C. P., & Hopper, L. M. (2020). The zone of latent solutions and its relevance to understanding ape cultures. In *Biology and Philosophy* (Vol. 35, Issue 5). Springer Netherlands. <https://doi.org/10.1007/s10539-020-09769-9>
- Tennie, C., Call, J., & Tomasello, M. (2009). Ratcheting up the ratchet: on the evolution of cumulative culture. *Philosophical Transactions of the Royal Society of London. Series B, Biological Sciences*, *364*(1528), 2405–2415. <https://doi.org/10.1098/rstb.2009.0052>
- Winterhalder, B. (1996). A marginal model of tolerated theft. *Ethology and Sociobiology*, *17*(1), 37–53. [https://doi.org/10.1016/0162-3095\(95\)00126-3](https://doi.org/10.1016/0162-3095(95)00126-3)