Appendix SB – Stability analysis and parameter exploration

Contents

Introduction

Simulation modeling is a tremendously flexible and powerful tool for investigating hypotheses about how properties of entities in dynamical systems of give rise to some proposed higher-level phenomenon. Scientifically, the fact that we can do this charges us with doing it (Epstein, 2007), since it constitutes a kind of test that the hypothesis may fail. But it also is not some binary question of telling the difference between episteme and doxa. In the process of building and running the model we may discover additional assumptions that must be made and evaluated, needed adjustments of the model, interesting phenomena we did not anticipate, or that additional models must be designed. We may also be alerted to possible empirical or experimental tests. In other words, simulating hypotheses about complex dynamical systems is a critically important tool in the larger scientific toolbox.

However, simulation is also a risky tool. It tends to produce black boxes that are very hard to penetrate. While this is not wholly unlike mathematical models, the challenges are different. Mathematical models tend to be relatively simple for those equipped with the right specialized knowledge to decipher them. For simulation modeling, on the other hand, one of its main strengths – its flexibility – easily becomes a liability since the threshold for adding unneeded and confusing components is very low. Also, as soon as one begins to build the “model world” one realizes how many assumptions that need to be made simply to make the model work – nothing can be implicit. Simulation models therefore tend to overwhelm their analysts, as tangles of dependencies and design choices, interacting non-linearly in a massively parallel fashion.

In this Appendix we seek to verify that the behavior of the model we use in this paper is robust to variations in its parameters and in some salient modeling choices. This, in conjunction with the specification of the model in Appendix SA should go some way toward ensuring that the results that we report follow from the Social Protocell Hypothesis rather than from contingent modeling choices. In the end, only a multitude of models independently designed from the same specification can persuade us that this is the case. That must happen, however, over time rather than in the scope of a single article.

What we do here is to investigate the neighborhood around the point in the parameter space represented by the BASE case. We do this by varying the numerical values of parameters up and down (and in some cases by altering a design choice) from this point and see how the model react on the emergent level. We pay particular attention to sensitivity in *qualitative* outcomes, such as whether a reported phenomenon disappears. Quantitative variations are of course to be expected in most cases. The type of behavior we hope to see, since it signals that the BASE case is representative for the behavior of the model in a broad and reasonably chosen parameter space, is either low or smooth quantitative sensitivity that corresponds to arguable expectations.

We first go through the parameters and model features to be varied one by one. We describe the parameter or feature and its implementation and motivate the variations. We focus on the issue of whether we had to be selective about parameter values for producing the results. If the reported behavior breaks down if we vary parameters or change modeling assumption, then this affects the results negatively unless we can defend the choice of numerical value empirically, which is generally hard to do in this type of conceptual models. We also analyze whether the observed variation in model is expected, in the sense that there is a straightforward argument in terms of the model for why the behavior varies in the way that we observe. In most cases we observe the time evolution of complexity of the IGUT institution, arguing that deviations in model behavior should be visible in such a plot. In some cases, we complement with other plots. Finally, we take up particularly important points separately (problems and future opportunities for study), and then we summarize the findings.

Parameter stability analysis

All runs are performed over 200,000 updates unless otherwise stated, using the BASE scenario with variations as specified, and averages over 50 runs with different random seeds. The BASE case, in turn, exists in three different versions: The BASE case with only the opportunistic and the IGUT apex locus types (institutions), the BASE case with the Social Learning Rate (SLR) boosting institution in addition, and the BASE case with SLR and genetic evolution of a cognitive factor in the agents. We use version that is appropriate for the test. That we average over fifty rather than a hundred runs (as we do in the main text) in each set means that some plots look noisier than in the Results section of the paper. However, we think the quality is adequate for the purposes and our excuse is that over twenty thousand hours' worth of runs went into the analysis.

Model updates per "year"

How much time elapses during an update of the model? In other words, how much happens in the world during an update? At which temporal resolution do we simulate the system?

Ideally, the behavior of the model should be independent of the time scale of simulation, but in practice this can be true only within certain bounds. For example, if the temporal resolution is coarse, fluctuations will become more dominant. The agents and socionts will not be "tested" as many times, their qualities and shortcomings may not shine through when noise does not average out as much.

Agents in the model grow older according to an external clock time in units we have dubbed "years". The "updates per 'year'" parameter tells us the time scale of one update of the system. All rates and probabilities are specified per "year", and the effect of this parameter is to scale their rates to this unit. For example, if we *decrease* the number of updates per year, the likelihood of a dominant community taking over a cell per update will *increases* correspondingly.

However, as mentioned, certain scales are imposed by the model design. Consider the example of territorial competition. If a lot of clock time elapses during an update (low number of updates per "year"), a boundary community cell that dominates a neighboring cell should, as we just stated, have a higher likelihood to take it over. But if this likelihood approaches unit, the time scale cannot be further increased since only one cell may change owner per update. The spatial discreteness thereby imposes a limit via the spatial resolution (referred to as a "speed of light" in cellular automata terminology).

Similar effects apply to, for instance, learning and to the amounts of resources harvested per unit time. In summary, we may expect the model implementation to cause the adjustment of temporal and spatial scales not to produce wholly scale invariant outcomes.

What is the right rate to use? Using a too coarse scale in straightforwardly undesirable – it is easily realized that agents need to be updated *at least* once during their lifetime. In practice a good number of times. In a sense, there may be no such thing as a too high resolution, but, at some point, we may not expect to see much further change in the output produced. So, at some point, the increase in computational cost that higher resolution entails will make it impractical to explore the dynamics. There should be a "Goldilockean" range where the parameter value strikes a wise balance between these ills.

The question we ask here is whether the results break down as the value we have used is varied. In other words, did we have to fine-tune the parameter?

In the BASE case we simulate on a time scale of ten updates per "year", which means that there will be on the order of a few hundred updates during the lifetime of an agent. We deem that a substantial amount of change on the sociont level (such as territorial change) may then happen during a typical agent lifetime.

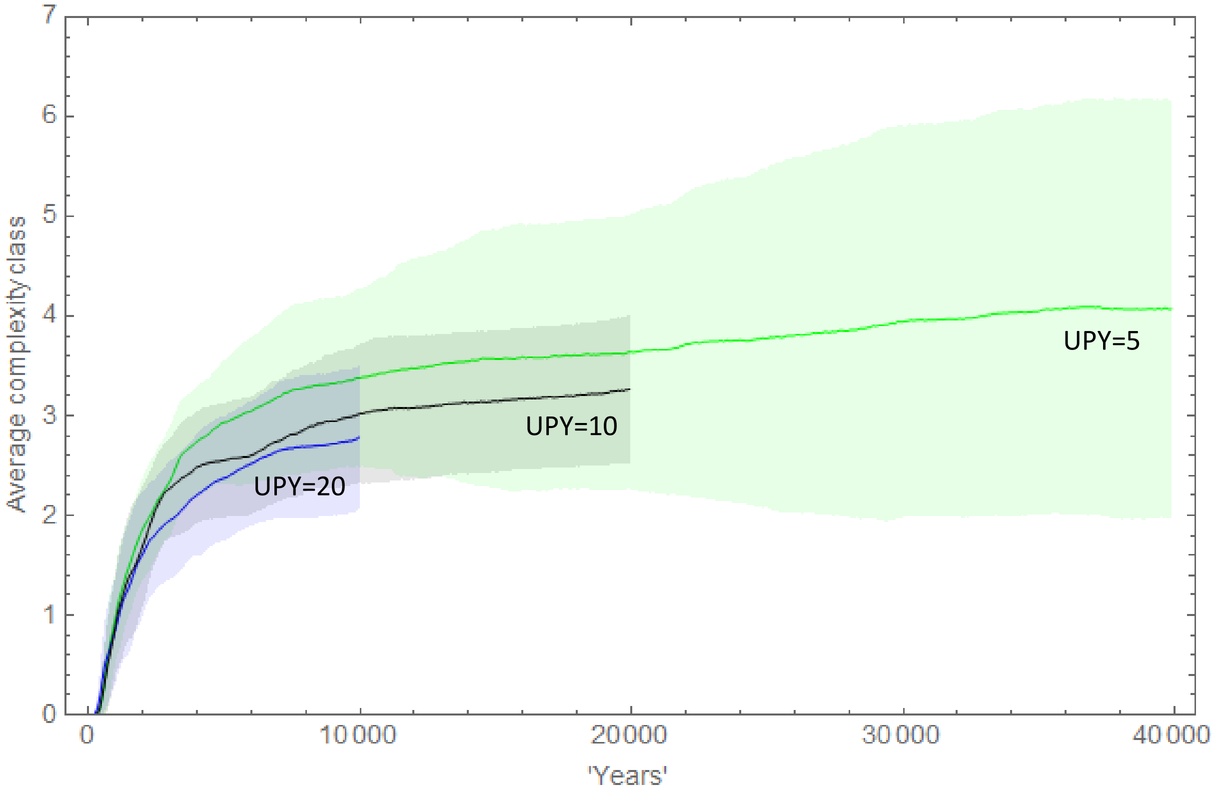

Figure B1 – We here see average complexity class of the IGUT institution in the BASE scenario, varying the number of Updates per "Year" (UPY). The runs were made over a fixed number of updates (200 000), which means that they vary in extent in terms of "years" (x axis). The time evolutions are comparable across this span of parameter values (UPY=20, blue. UPY=10, black. UPY=5, green). The black plot is the BASE scenario.

We test what happens if we adjust the temporal resolution by a factor two upward and downward (Figure B1). As we see in the plots, the behavior of the model varies smoothly in the neighborhood of the used parameter value. As may be expected, the standard deviation increases the coarser the time resolution we use.

Setting the parameter lower means that fewer updates happen during the lifetime of an agent. We can also see that the result is that the dynamics is faster: more "years" will elapse during the 200 000 updates that are made. The change in dynamics is however smooth and no qualitative effects can be seen, which leads us to conclude that the results reported is not sensitive to the used temporal resolution.

Community dispersal threshold – population

Social communities are not effective at any size, and we have assumed an arbitrary number of members below which the community will cease to function and disperse. The parameter value is bounded downward by zero – a community where no members remain. Upward it is bounded by the criterion for the splitting operation (in our model fixed at 50 members) since the model will not work if splitting yields communities below the dispersal threshold. Beyond this, it is hard to motivate the numerical value of this parameter based on empirical data, and we have not tried to do so. It is therefore desirable that the dynamics that we observe is not sensitive to our arbitrary choice.

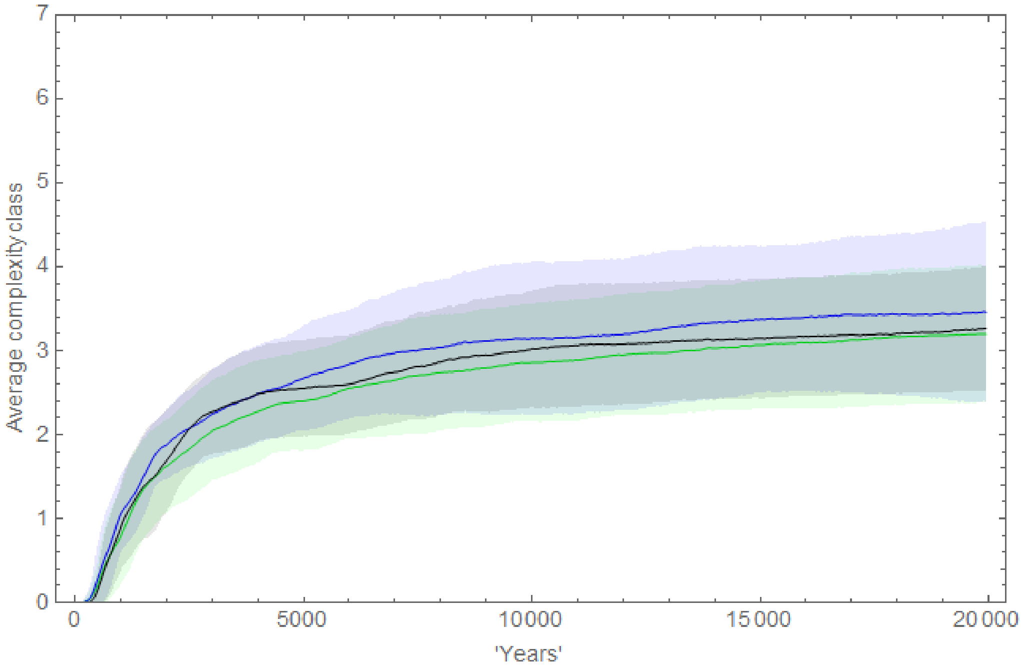

Figure B2 – We here see average complexity class of the IGUT institution in the BASE scenario, varying the community dispersal threshold population (DTP). The time evolutions are highly similar across this span of parameter values (DTP=20, blue. DTP=10, black. DTP=5, green). The black plot is the BASE scenario.

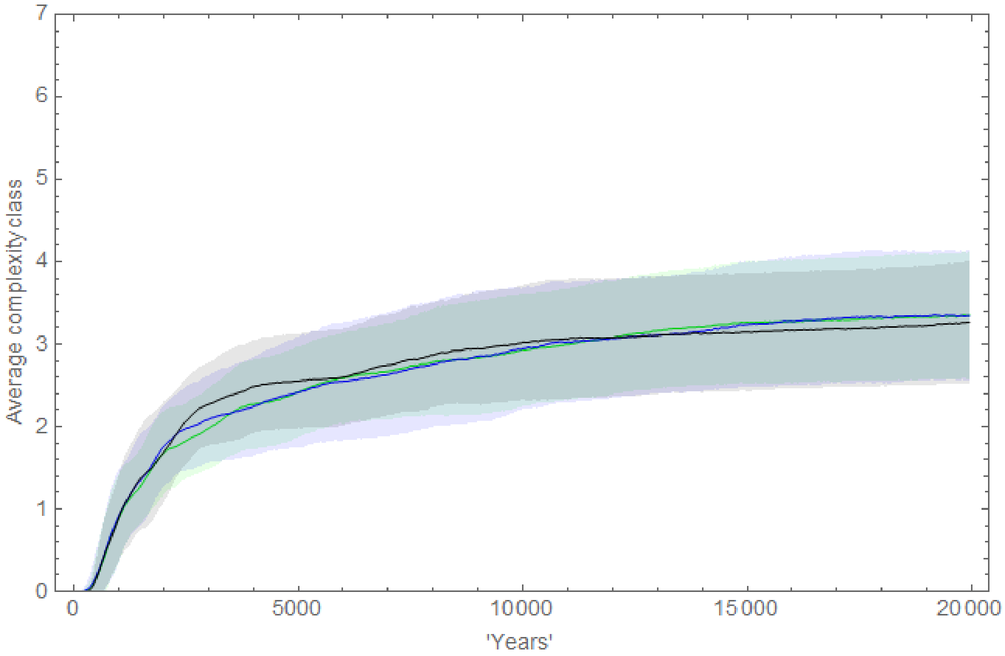

Figure B3 – We here see average complexity class of the IGUT institution in the BASE scenario, varying the community dispersal threshold area (DTA). The time evolutions are highly similar across this span of parameter values (DTA=0.005, blue. DTA=0.000625, black. DTA=0.00001, green). The black plot is the BASE scenario.

To test we vary this parameter upward and downward relative to the chosen value for the BASE scenario to observe the outcomes (Figure B2). The model is clearly robust to exact choices of this parameter value. The only difference appears to be that the standard deviation increases slightly with a higher threshold, which is not surprising since random fluctuations downward in community size will more frequently be punished (a newly split community has twenty-five members, which is only five more than the dispersal threshold).

Community dispersal threshold – area

It is deemed that territories cannot function if they are arbitrarily small also in terms of size. This threshold is mainly in place to prevent problems with resolution on a finitely sized grid; see *Varying the spatial resolution* below.

A too high value would cause the community to disperse at typical territory sizes that arise at equilibrium in the course of simulation. A too low value will mainly cause above mentioned resolution problems.

Choosing a higher and a lower value for this parameter produces neither qualitative nor quantitative differences in average histories, and we deem that the model is not sensitive to the choice of parameter value; see Figure B3.

Community splitting population threshold

The community splitting threshold must co-vary with the dispersal population threshold so that a divided community does not fall below the dispersal threshold. Varying this parameter has a number of effects that may be interesting in their own right. Let us first reason about how the model should react as we vary this parameter.

These have to do with the statistical effects of low numbers of discrete agents. Decreasing the threshold, the average size of communities in the system will decrease. Now, if the communities shrink, fluctuations (deaths, births, luck in spatial competition) will increasingly dominate their internal dynamics. Because of this, decreasing this parameter should weaken sociont-level selection, while increasing it would strengthen selection on the sociont level.

Splitting at a smaller size should also be expected to adversely affect sociont-level fidelity of cultural inheritance since the likelihood that some institution component locus will not remain represented by alleles in either of the daughters will increase (see Figure 2, main article). This will cause loss of cultural information and thereby reduce fecundity more and more the more complex the institutions become, eventually cancelling out the fitness benefits of further complexity increases. In other words, lower equilibrium sociont complexity should be expected.

Furthermore, even if a split is in principle successful (i.e., the institutions make it to both offspring) daughter socionts may still be fitness-depressed after a split. Say that the parent maintains ten full expressions of an adaptive institution (i.e., there are ten alleles of its least frequent loci). It is still not certain that the offspring together will express ten instances of this institutions after a split. That would demand that the alleles of the least frequent component locus are split evenly. For example, three alleles may end up in one daughter and seven in the other daughter, leading to only three effective expressions of the institution in daughter A.

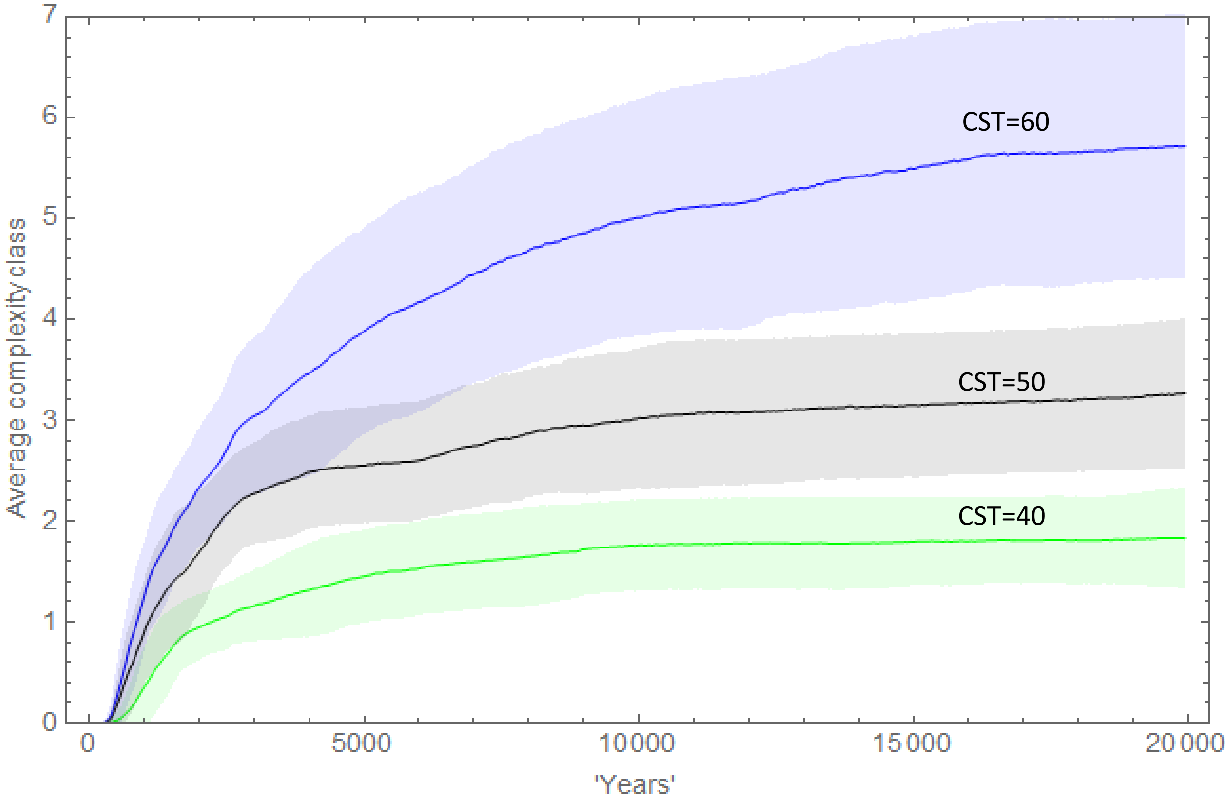

Figure B4 – We here see average complexity class of the IGUT institution in the BASE scenario, varying the threshold for community splitting (CST). The threshold affects the strength of selection but not the qualitative behavior of the system. (CST=60, blue. CDT=50, black. CDT=40, green). The black plot is the BASE scenario.

Even if these imbalances are eliminated over time, splitting will thereby come at a fitness cost to the sociont. So, the more of the time its underpinning social network can grow uninterrupted, the better. Increasing the size of the groups in this manner reduces the effect of fluctuations and their fitness-depressing effects, but at a diminishing rate, since they will diminish at higher population sizes due to the law of large numbers. We may surmise that an institution that permitted larger community sizes could be selected for because of this effect.

Analyzing Figure B4 we conclude that although the capacity of the system to undergo cumulative evolution increases when we increase the parameter value beyond the value (CST=50) that was used for producing the results, cumulative evolution remains also for the lower value used here (CST=40), albeit marginally.

In the BASE scenario, the split threshold is set to balance between the cost of computing[1] and the need to avoid domination of fluctuations in the manner just described. Community size does, it should furthermore be noted, not seem to be a proximate reason for community fission in *Pan*. The proximate reason appears to be conflicts between high-ranking males, which, in turn, may result from a number of reasons (such as skewed male:female ration; see J. T. Feldblum et al., 2018). However, from a theoretical point of view, it is hard to see how the performance of collective mechanisms for reconciliation possibly could scale anything but poorly with community size (e.g. Dunbar, 1992, 1995; Lehmann & Boesch, 2004; Moffett, 2013), which, again, makes community size important as an overall

---

[1] If we increase the size of the communities, we also need to increase the number of agents in the system to avoid having too few communities. This makes for more and more time-consuming simulations.

modifier of the likelihood of such conflicts spiraling out of control. We will return to how important assumptions about the process of splitting are.

Mode of community fission – random allocation of agents to daughter communities

We divide communities by first sorting the agents according to rank (see *Within the community*, Appendix SA), then moving agents with odd rank to one daughter community and agents with even rank to the other. This has the merit of preserving age and performance in both offspring, and is also in some accordance with how splits in *Pan* communities appear to happen (i.e. via conflicts between high-ranking males that subsequently become the center of the new social networks; see J. Feldblum et al., 2018).

This design favors the hypothesis in at least two ways. First, even if the number of individuals in the community is fairly large, the number of individuals possessing key traditions may not be very large, especially if the complexity of institutions is pushed near the limit of what the social learning rate can support. If it is plausible to assume that older and more successful individuals would be more likely to possess important traditions, then "rank preserving" fission would tend to divide traditions more evenly between daughter communities, thus increasing the fidelity of sociont-level inheritance.

Having motivated why we choose a more specific way of splitting communities, we will now investigate whether the results are sensitive to this choice (if so, it should have been part of the study). This is particularly important because the empirical argument for the choice is not strong, and because the choice favors the hypothesis.

We have therefor tried out an alternative implementation of the community fission mechanisms where we divide the individuals randomly; see Figure B5. As we see, while the problems that we outlined may well exist, the modeling choice we made was robust, at least as far as this test goes. We would not have gotten dramatically different results by using this simpler random mechanism instead.

The mode of splitting may be important as a constraint on reaching higher levels of cultural complexity at some point, but that is beyond the scope of this paper.

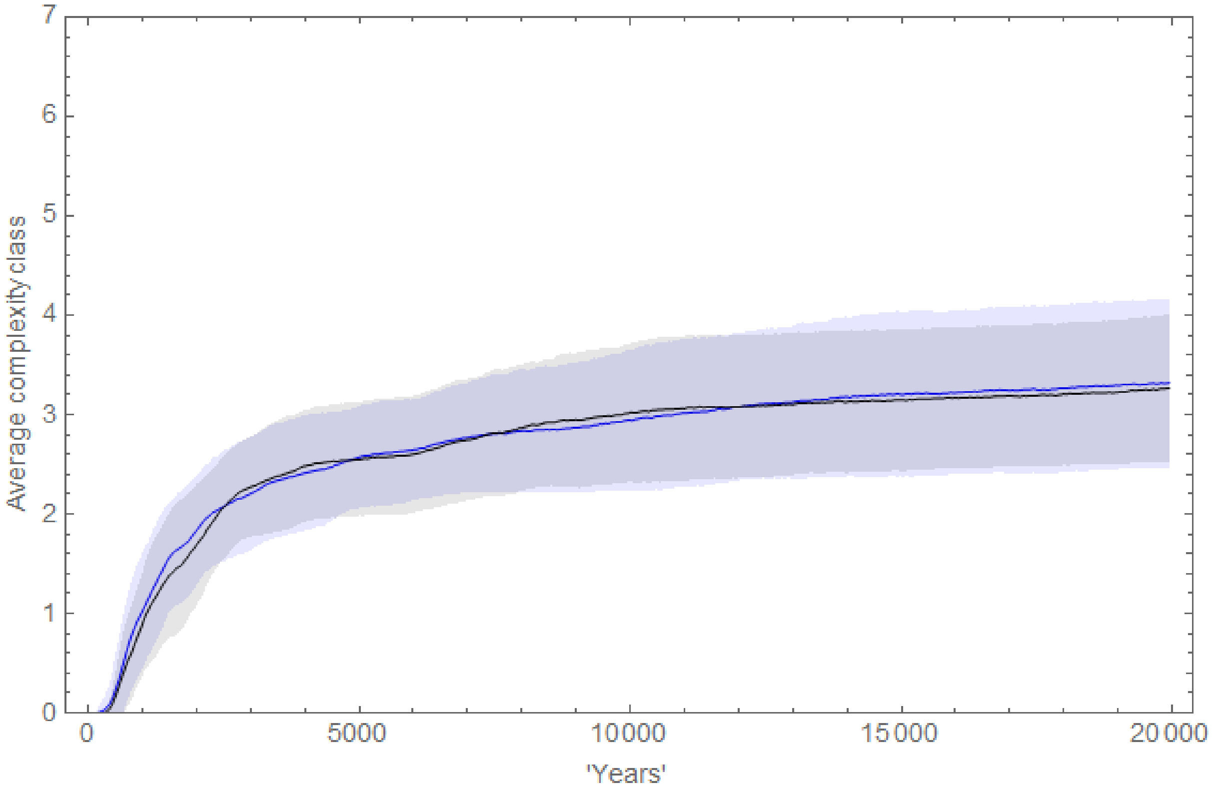

Figure B5 – We here see average complexity class of the IGUT institution in the BASE scenario, varying the mechanism for community fission. The choice of mechanism has very little impact on the dynamics, apart, possibly, for a slightly larger standard deviation in the random splitting case (blue). The black plot is the BASE scenario.

Criterion for community fission – random splitting at a fixed rate

Returning again to the question of how important our assumptions about the community splitting process are, we will now test an alternative, and very simple, criterion for *when* splitting occurs. The model initiates a split at a fixed threshold community membership, which in the BASE case is N=50. We have already tested sensitivity to the threshold value, but what if splitting is less sensitive to community size than we assume?

To test this, we try out a mode of splitting that is wholly random and not tied to community size at all. In this alternative criterion, splitting happens at a fixed rate rather than at a threshold, or as may plausibly be most realistic, at a rate that increases as a function of community size (i.e., that community size is one among several factors).

We find that the reported qualitative behavior is not adversely affected by shifting to this simpler alternative. This is in the sense that cumulative cultural evolution on the sociont level is not diminished (Figure B6). To the contrary, at least in this model, it *increases* in intensity.

The reason is that when communities may become arbitrarily large (see also Figure B4), this has the effect that sociont-level inheritance fidelity increases as well. (This is also the reason why we have limited the number of updates to 20,000 rather than 200,000 – high levels of complexity are reached much more rapidly.)

This furthermore indicates that assumptions about causes of community splits are not as important as they may seem. Competition will, after all, happen in any case, and that communities do split over time is uncontroversial.

It is fair to ask why communities would thereby not be predicted to have grown early and decisively under positive selection according to the SPH. One reason may be that other selection pressures favor smaller community sizes. For example, cognitive mechanisms for community cohesion may be inherently and fundamentally expensive. The social brain hypothesis (and hypotheses to similar effect; e.g. Byrne & Whiten, 1989; Dunbar, 1998; Whiten & Byrne, 1997) holds that group cohesion places a high strain on cognitive capacity which, in turn, is expensive since it requires increases in brain size. This may in particular be the case for *Pan* and *Homo* as they rely on a "political" process of group cohesion that does not reduce to biological kinship, which is the case in other primates (e.g. Read, 2012). It should be stressed here that we are not considering the complexity of individual practices but of cultural lifestyles as integrated systems of such practices. The SPH could, however, potentially add another factor that may have favored larger brain in *Homo*.

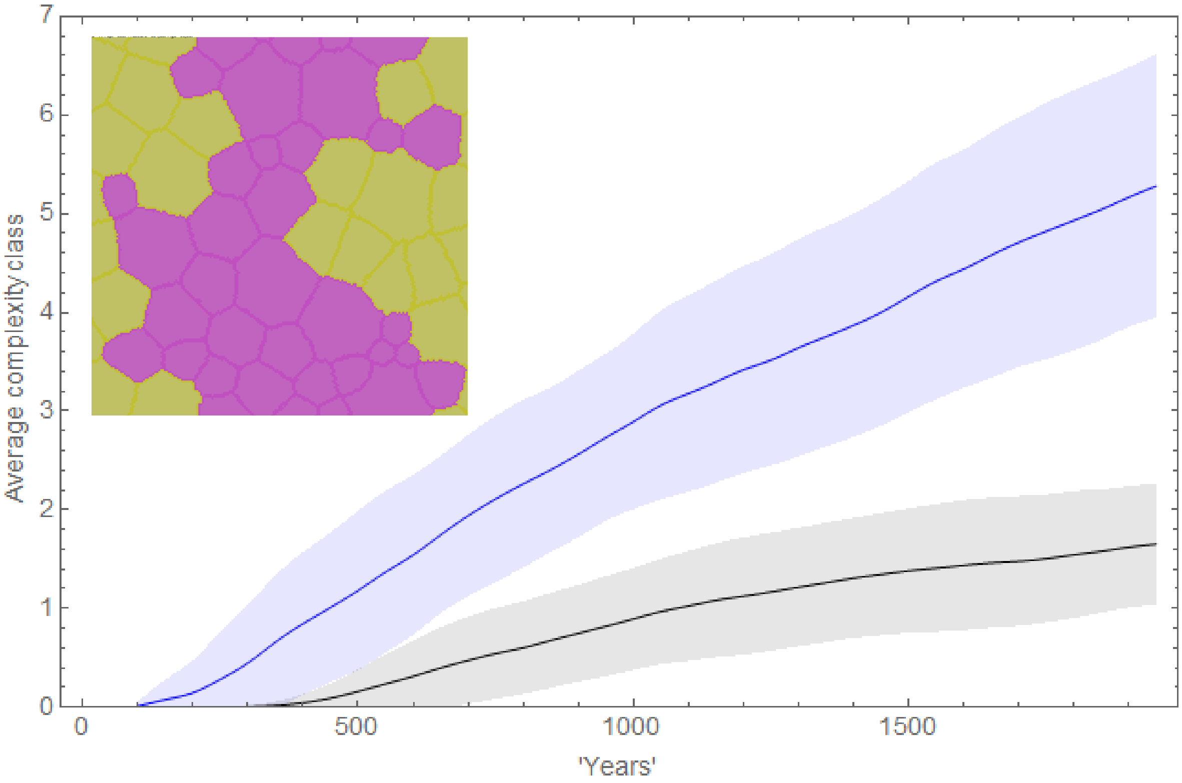

Figure B6 – We here see average complexity class of the IGUT institution in the BASE scenario, varying the criterion for community fission. Since groups can become much larger at the rate we have used, the fidelity of sociont-level inheritance increases, making for a rapid evolution of high cultural complexity (blue). The black plot is the BASE scenario. The inset picture shows a configuration illustrating the wide span in community sizes seen (compare Figure 4 in the article). Runs were made over 20,000 updates.

Varying the spatial resolution

As we have already remarked, the discreteness of the underlying space introduces a maximum "speed of light" phenomenon in processes where grid point states are updated in interactions within finite neighborhoods. In this model, the territories can expand maximally one cell per update, so the higher the spatial resolution, the lower the maximum speed with which the community can change (unless the change in resolution is interpreted as an increase or decrease in the area that is simulated, but that demands a simultaneous scaling of several parameters.)

To test how the choice spatial resolution affects the simulation, we here use a four times higher and four times lower spatial resolution than the value used in the BASE case. We see that the qualitative dynamics are not strongly affected (Figure B7). As the “speed of light” decreases (increasing the resolution), more and more action will also take place per fractional amount of territorial change. This did not seem affect the strength of selection but did act to reduce variance in outcomes slightly. For the lower resolution we may see that the strength of selection seems to attenuate slightly.

At the lower resolution we likely begin to see an effect that can be readily verified by running the model at even lower resolution. As socionts become increasingly efficient (later in a run) they will have smaller and smaller territories in terms of the number of geographical cells that they occupy.

This will for example make the splitting mechanism less exact. The reason is that it seeks to find a section across the territory that will divide it in half (see *Between communities* in Appendix SA), and this section is more likely to be asymmetric if the number of geographical cells is low. If daughter territories tend to deviate substantially from half the parent’s size, fecundity will be reduced since the population will still be divided in half. The likelihood thereby increases that daughters will have too many agents per unit area after a split, which gives them a disadvantage. In addition, chance events will in general have a larger influence when territories consist of less geographical elements (each change is a larger part of the territory).

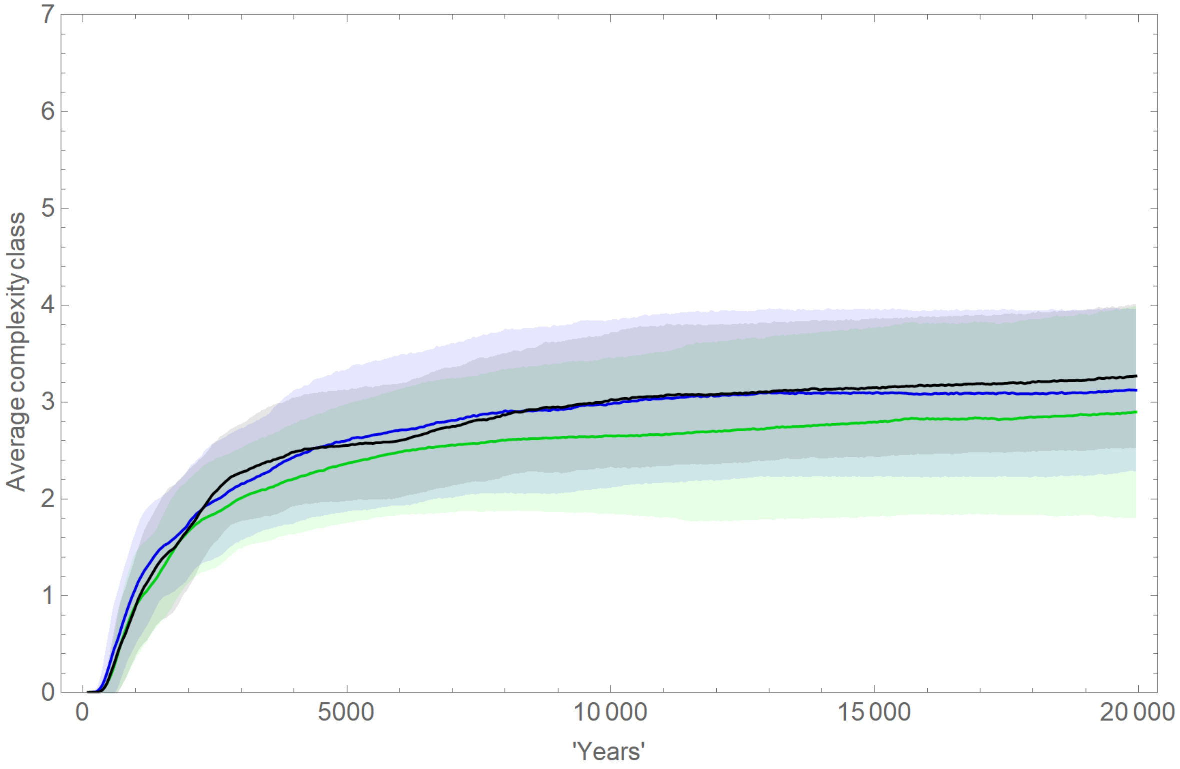

Figure B7 – We here see average complexity class of the IGUT institution in the BASE scenario, varying the spatial resolution of land. The dynamics is quite robust to variations in the spatial resolution in this range. (Resolution 400×400, blue. Resolution 200×200, black. Resolution 100×100, green). The black plot is the BASE scenario.

Terrain resistance

Territorial competition results from “pressure” exerted by the populations of communities within their territories; in short, this “pressure” is the density of agents within the territory, but with abatement from the center toward the periphery to represent benefits of remaining in physical proximity and of minimizing the interface with other communities (see *Space and ecology* in Appendix SA).

The terrain resistance, in this context, is "null-pressure" exerted by uninhabited land, representing for example competition from other species and other sources of attrition. This comes into play in two circumstances. The first is the initial spread as the seed community divides and its descendants fill the empty configuration. The second is the infill by neighboring communities as a community disperses, leaving uninhabited cells behind.

If terrain resistance is too high, communities will not expand into empty areas. If it is too low, they will bloat and sprawl out. The parameter may be suspected to have little effect as soon as communities fill the grid, however, since uninhabited land will rarely be encountered.

This analysis is confirmed by the plots (see Figure B8): No discernable difference in dynamics can be observed. However, in many cases, runs using the higher parameter value failed to get going because the initial lone community would die due to fluctuations, leaving the configuration empty. That, however, did not affect the dynamics *if* it gets started, and using a lone community as a starting point has in itself no empirical significance.

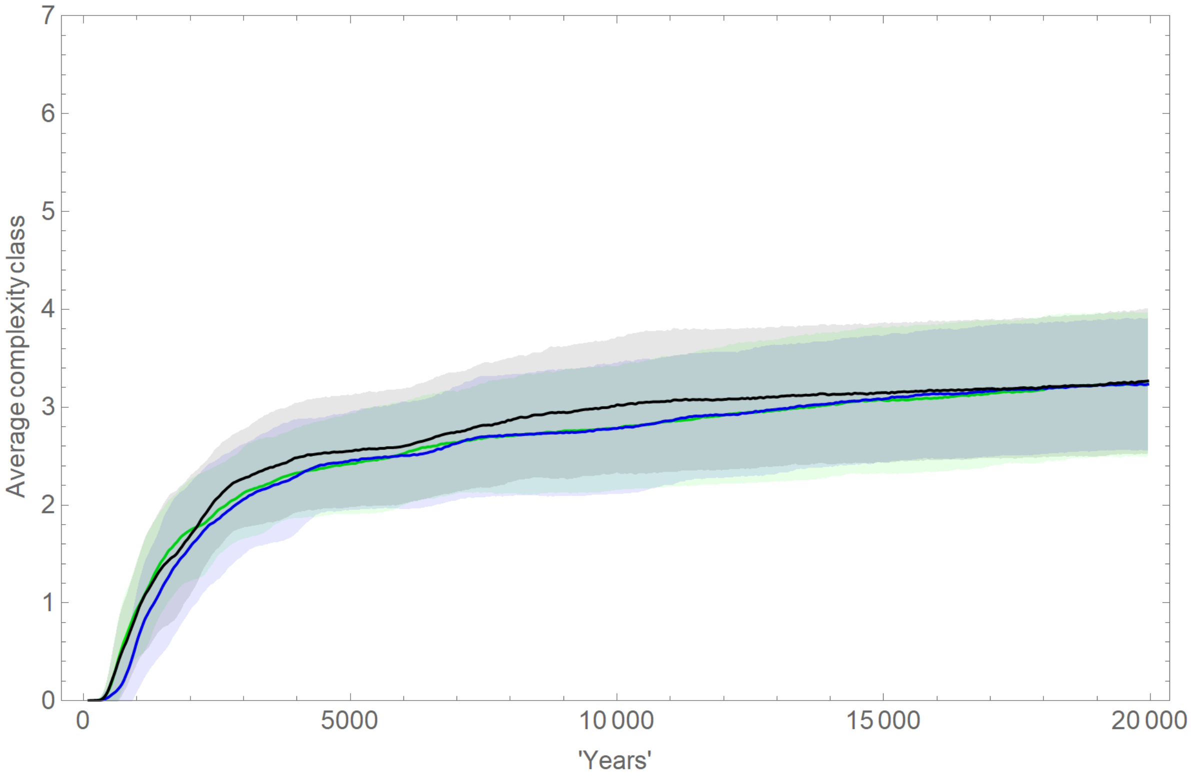

Figure B8 – We here see average complexity class of the IGUT institution in the BASE scenario, varying the terrain resistance parameter. The dynamics is robust to variations in this parameter. Not visible in the plots is, however, that the seed community died in a sizeable fraction of runs using the higher resistance. Once they began to grow, there was no difference, however. (Resistance=0.00005, blue. Resistance=0.00001, black. Resistance=0.000002, green). The black plot is the BASE scenario.

Agent energy cost per time unit

The cost of living per unit time must be small enough that, initially, available resources suffice to keep the agent alive. We assume that any realistic simulated species will be adapted to its biotope in such a way, i.e., that any realistic evolutionarily starting point will be tuned in this way.

If so, at initialization from an empty configuration, seeded with one community, spatial competition should be low, and this parameter should be set such that a surplus produces population growth. This leads to the growth and division of territories. This expansive phase then continues until the lattice is filled with territories, after which point the dynamics that we study begins to play out. This initial phase is included in the analysis, but it is, relatively speaking, very short. Once the lattice is filled, territorial competition ensures that the cost of living will precisely balance out with resource input and the death rate. Selection for higher efficiency variants now happens since such variants will have a birth rate that exceed their death rate as long as the communities are surrounded by communities with lower efficiency.

Increasing the cost of living in terms of energy (without tuning other parameters) will lower the growth rate. This leads to large areas at equilibrium since each agent will need more area to sustain itself. The dynamics will still be similar, but with sparser agents and communities. If the cost is increased so that not even the entire lattice can feed a community, this will simply cause extinction. Decreasing the cost of living will, correspondingly, decrease the equilibrium community areas. The "sweet spot" is when there will be a sizeable number of communities at the initial state (with no culture), such as at least a few tens of communities (see e.g., Figure 4 in the main article for an illustration of what a lattice populated by communities near equilibrium looks like.)

As efficiency increases due to cultural evolution in the simulation, the size of communities will progressively shrink from the initial point, per the same argument.

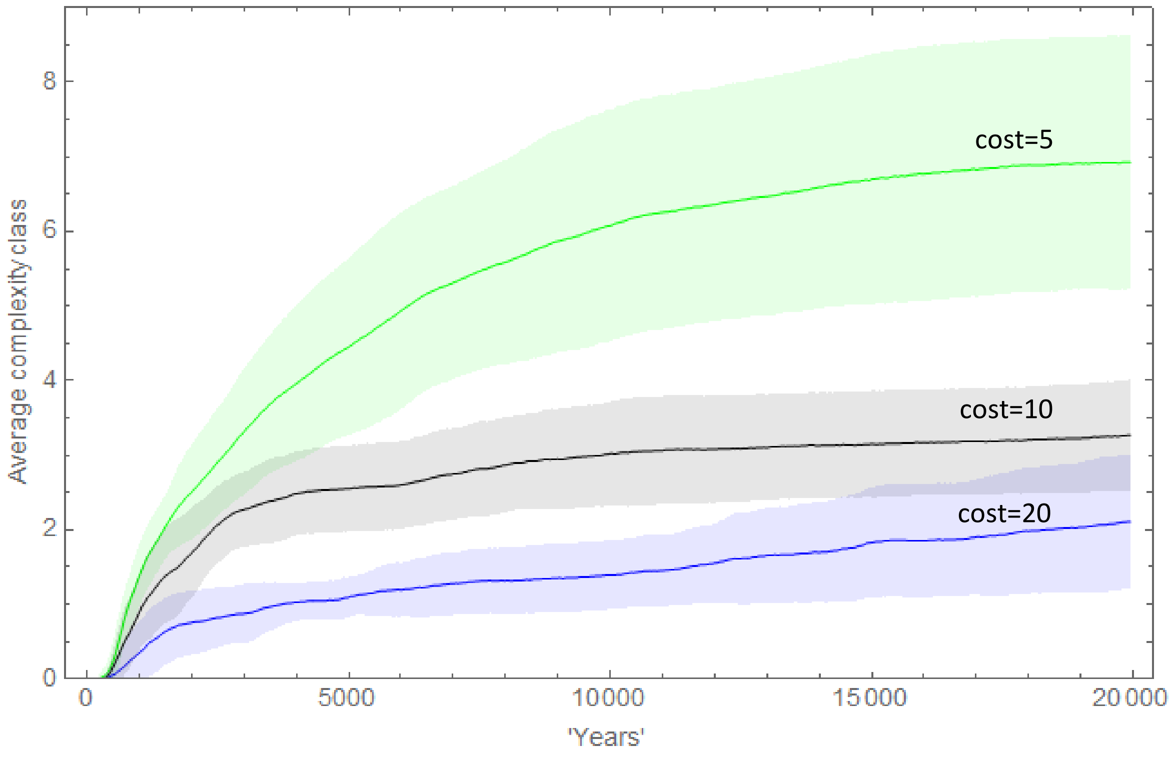

Figure B9 – We here see average complexity class of the IGUT institution in the BASE scenario, varying the cost of living for agents in terms of energy. The qualitative dynamics is robust to variations in this parameter, but it smoothly affects the pace and equilibrium level of cumulative evolution. (Cost=20, blue. Cost=10, black. Cost=5, green). The black plot is the BASE scenario.

As we see in the plots (Figure B9), changing this parameter has the predicted effects. The precise value used does not affect the qualitative dynamics reported in the results, but it does impact the rate at which evolution proceeds. It may be noted that the results may have been clearer if a somewhat lower energy cost had been used. On the other hand, to maximize analyzability, we want selective sweeps to be relatively separated in time, rather than overlapping (that we see a rapidly increasing complexity in the low-cost plot in Figure B9 indicates overlapping sweeps-in-sweeps.)

Energy cost per offspring

The effects of varying the cost of offspring may seem counterintuitive at first glance (Figure B10). Before discussing these effects, let us note that the qualitative behavior of the system varies smoothly with the parameter.

As for the counterintuitive reaction to the changes: Why would selection on the sociont level *increase* dramatically as the cost of offspring increases (indicated by the achievement of substantially higher complexity in the same period of time?) We may at this point only speculate about what this means, and whether it is mostly an artifact of how reproduction has been implemented in the model.

First, the implementation of reproduction works such that there is an energy buffer in each agent in reproductive age where surplus energy is stored. The cost of offspring is the threshold when an offspring is spawned, and the buffer is emptied.

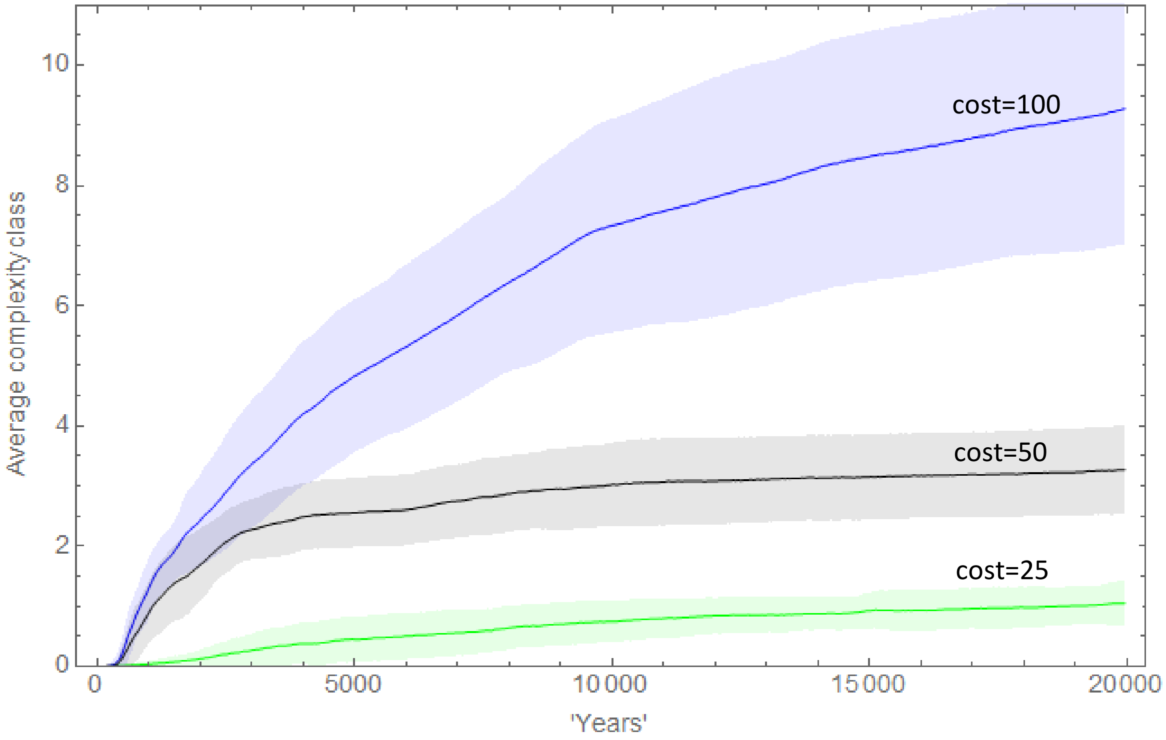

Figure B10 – We here see average complexity class of the IGUT institution in the BASE scenario, varying the cost of offspring for agents in terms of energy. The qualitative dynamics is robust to variations in this parameter, but it affects the pace of cumulative evolution in a way that is prima facie counterintuitive, namely that competition is lower if offspring is cheaper. (Cost=100, blue. Cost=50, black. Cost=25, green). The black plot is the BASE scenario.

The likely explanation, which bears more scrutiny, is that increased efficiency gives a larger adaptive edge when offspring is expensive. The implementation has some possible side effects that may obfuscate the results (or be significant). One such effect is that juveniles do not contribute to obtaining resources and performing other duties, including guarding the territory. So, if an edge in efficiency is exploited very rapidly by the production of large numbers of juveniles (low cost per offspring), this can cause an overshoot that eliminates the edge of the more efficient sociont during a period of time. It will rapidly gain a high proportion of unproductive juveniles (this may be seen be studying the states of the simulation in detail as it runs). As stated, whether this is a potential result or simply an artifact of the implementation is beyond the scope of the stability analysis. For that purpose, we are content with noting that the response is smooth and does not change the qualitative behavior of the model.

Agent innovation rate

This parameter governs the rate at which changes occur in the topology of institutions, i.e., the addition, removal, and movement of component loci (see Appendix SA and illustration in Figure 3 in the main article). Increasing this rate will increase the rate at which adaptive novelty arises to expand existing institutions. That is, if the value is sufficiently high, we will see many selective sweeps-in-sweeps as one superior variant will not reach fixation before a new improvement arises in the expanding field. A too high rate will therefore lead to results that are harder to interpret visually.

Decreasing the rate, conversely, will lead to prolonged periods of stasis between sweeps that happen intermittently. The average dominance of specific levels of complexity is much lower, which means that the appearance of new sweeps is less regular. The sample histories (displaying representative histories for single runs) illustrate the difference.

The parameter is not expected to affect the dynamics with regard to whether cumulative evolution happens or not, and neither is it expected to affect the outcome strongly in a quantitative sense either. It is desirable, however, to use a value for which sweeps are not too overloaded, yet not unnecessarily far apart (demanding unnecessarily long runs).

Note in Figure B11b that we see the appearance of a problematic effect of the formulation of the model, that we will have reasons to discuss further later in the text: Higher and higher complexity classes appear and remain in existence, without gaining ground, but also without disappearing. This creates the illusion in Figure B11a that a higher rate of innovation (rate=0.5) leads to higher cumulative complexity than the rate used in the BASE case. Let us discuss why this is an illusion and actually an unwanted behavior of the model, but, luckily, one that does not affect the results.

The reason is that once a configuration becomes stable (for whatever reason) inert socionts can persist indefinitely. Losses in division events is what primarily causes socionts with too high complexity to suffer, but if they never divide, they can keep accumulating complexity subject only to the considerably lower rate of cultural loss that stems from non-division-related sources of cultural loss (fluctuations in births and deaths); see Figure 6 in the main article. The reason why these higher-complexity socionts do not grow and divide (they do not out-compete even much less complex socionts,) is that the increase in complexity also means fewer successful executions of the superior institutions. To the extent that they do divide, it is almost guaranteed that at least one of the daughters will not contain a functioning version of the institution (and rapidly be outcompeted) since they are soon far beyond the level of complexity that can be supported via sociont inheritance via division. The more we increase the rate of innovation, the more likely it becomes that this will happen. *Notably, this does not happen in the BASE case* (Figure B11c).

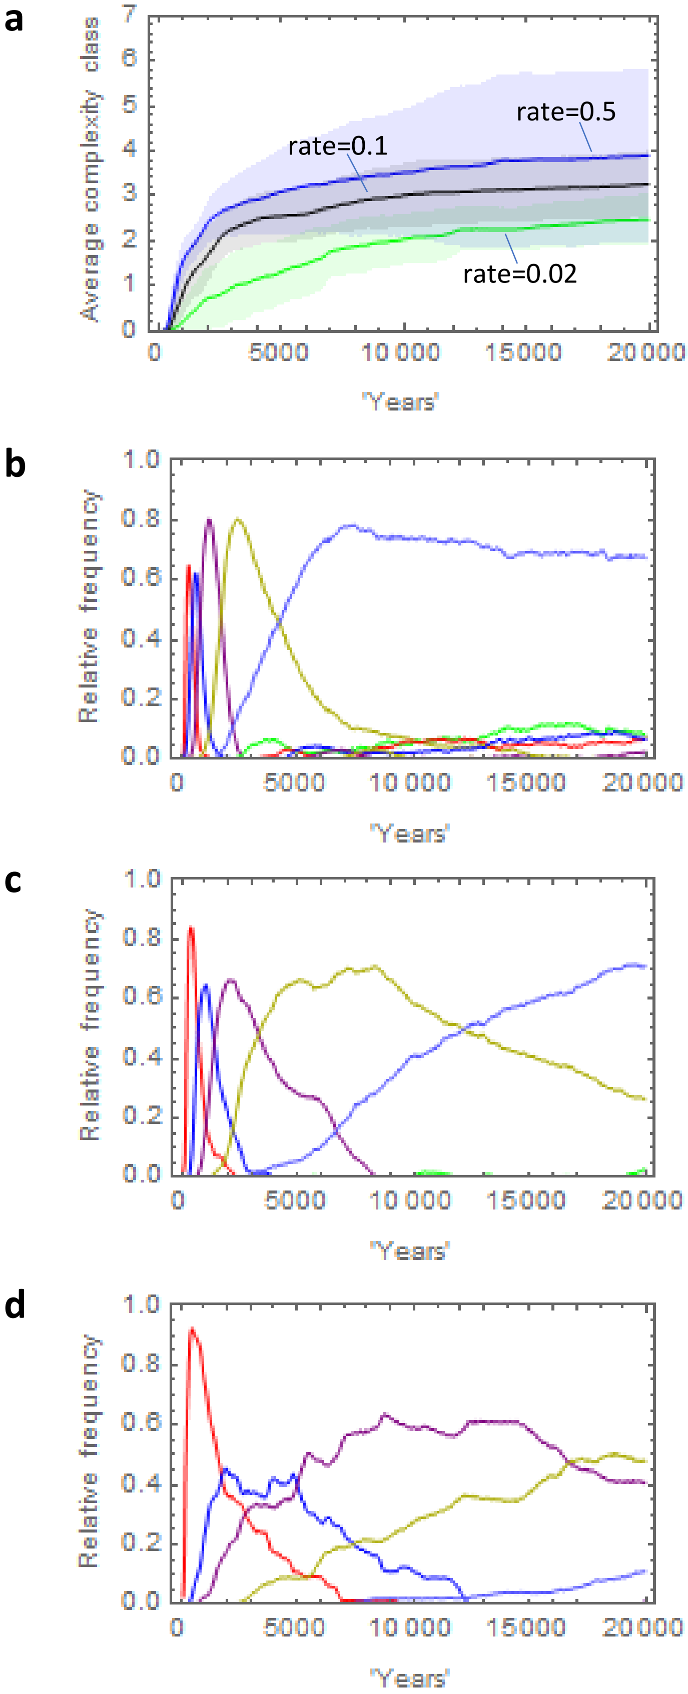

Figure B11– (a) We here see average complexity class of the IGUT institution in the BASE scenario, varying the rate of innovation. The qualitative dynamics is robust to variations in this parameter, but it affects the pace of cumulative evolution – not by increasing the capacity for cumulative evolution but by speeding up the arrival of adaptive novelty. (Rate=0.5, blue. Rate=0.1, black. Rate=0.02, green). The black plot is the BASE scenario. **(b-d)** Plots of the average relative frequency of socionts with different complexity classes as a function of time for the cases of rate=0.5 **(b)**, rate=0.1 **(c)**, and rate=0.02 **(d)**. Complexity classes are color coded cyclically from 1 and upward through red, blue, purple, yellow, light blue and green – beginning then anew for higher classes.

The reason why communities can persist in the model is that they never dissipate for external reasons – socionts in the model die only by being outcompeted by other socionts.[2] In other words, introducing a non-negligible rate of community failure for "other reasons" (say, diseases, disasters, bad luck, etc.) would be realistic and would fix this problem.

If we look closer at what happens in the high (Figure B11b) and BASE rate (Figure B11c) cases, we see that the same complexity class dominates in both cases – they do not actually lead to very different outcomes. The difference is that the higher rate yields this scatter of socionts with higher complexity, which shows up in the average. Figure B11d tells us that the green plot in Figure B11a (corresponding to the low innovation rate case) seems to keep increasing beyond the plotted interval on its way to a similar state as the BASE case. Finally, the higher rate of growth of novel complexity classes seen in the higher innovation rate cases (Figure B11b,c) happens because higher complexity classes will tend to arise several times independently, creating several growing clusters.

The interpretation of the parameter itself is somewhat degenerate in its empirical interpretations. On the one hand, an increased rate may be interpreted as a higher creative capacity. However, it could also be interpreted as a higher resolution of the model in terms of the size of the creative steps that are modeled, with a high rate corresponding to each addition corresponding to smaller increases in efficiency (this depends on how this parameter is matched against the function that maps complexity to efficiency).

Agent learning rate

This parameter sets the rate at which potential functions (such as resources) are noticed and targeted with a lone apex locus via learning (as opposed to social learning.) This affects how long time it takes before traditions arise at all. The value is set such that this happens relatively quickly in the simulation.

The parameter is relatively unimportant since the number of targets for cultural adaptation is small and since the traditions themselves lack internal organization. That is, if two different communities come up with a simple strategy for targeting a resource independently, then these will be identical in the model, whereas in reality they could have represented different ways of solving the same set of problems, which in turn would affect downstream history. This unrealism represents a simplification that is not of consequence under the analysis performed here.

The only effect we observe by tuning this parameter is indeed how long time it takes before the cumulative dynamics gets going, which is visible as a slower gain of the initial levels of complexity at the lower value of the parameter (see Figure B12). Higher values provide a slight head start, but otherwise does not affect the outcome.

One lesson is that the subdivision of rates between "learning" and "innovation" is somewhat artificial and could be combined.

---

[2] In principle the population of a community may also fluctuate below the dispersal threshold for no reason, but this is in practice very unlikely.

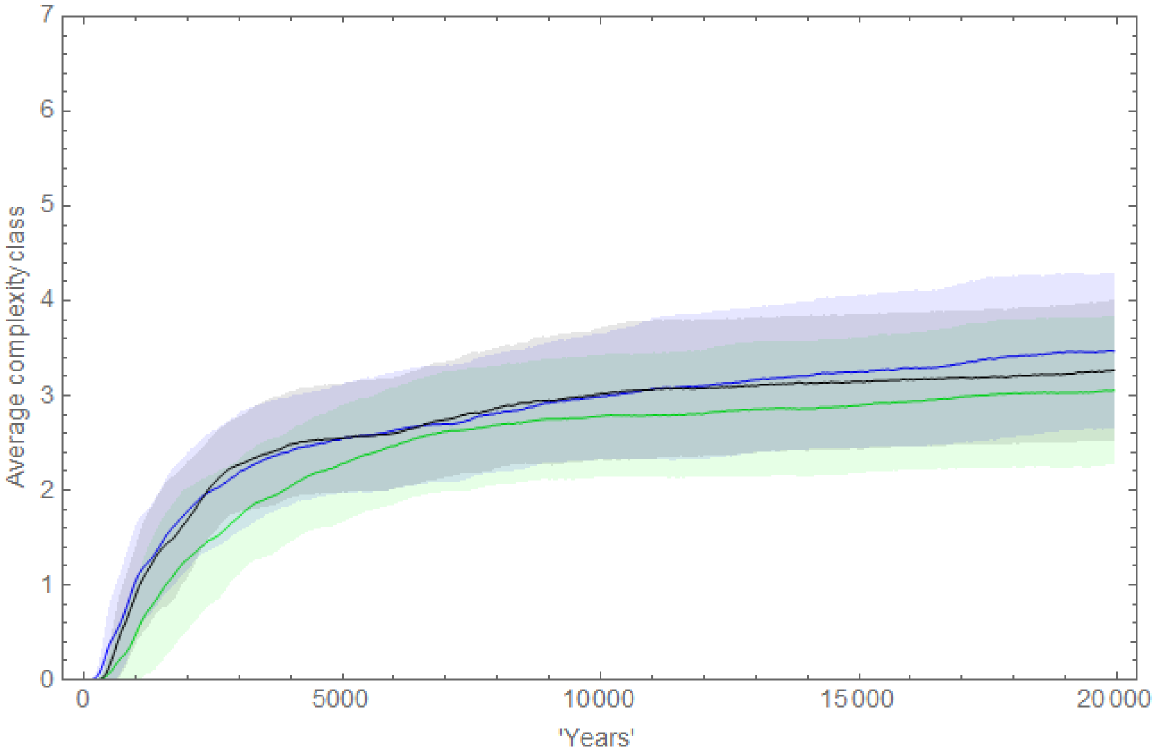

Figure B12 – We here see average complexity class of the IGUT institution in the BASE scenario, varying the rate of learning, i.e., the rate at which agents target resources and custom function with new apex loci. The qualitative dynamics is highly robust to variations in this parameter since it, in practice, affects only how long time it takes before the dynamics gets started. (Rate=1, blue. Rate=0.25, black. Rate=0.05, green). The black plot is the BASE scenario.

Conformity bias temperature

Conformity bias temperature (see Appendix SA) governs how the likelihood of socially learning a specific tradition allele, from a specific role model, is impacted by the relative frequency of the locus of that allele in the population. Note, however, that since agents will have one opportunity to learn a tradition locus from each individual possessing an allele in an update, there is also an inherent linear bias for tradition loci with many alleles in the model. This parameter concerns an extra non-linear bias.

It is far from clear whether, and, if so, how much, pre-*Homo* hominins would have been biased in their social learning in this way, so it is particularly important to verify that the model is not sensitive to this parameter in a way that impacts the results, or, if it is, to clearly state that the SPH somehow hinges on the presence of such a bias. We use temperature as a model of role model choice (see *Within the community* in Appendix SA).

In Figure B13 we see that a higher temperature (low conformity bias) increases the rate of spread, making for a more rapid increase in institutional complexity. The shape of the complexity curve is still essentially the same, however, so no qualitative effect can be seen. It may be interjected that we see so little of the plot for the high temperature in Figure B13 that we could be looking at a logarithmic curve. However, there is nothing in the underlying process that suggests that we would get a logarithmic curve. The likely explanation is that with an almost complete lack of conformity bias, we maximize the spread of alleles in the population, which will increase the sociont-level fidelity (see Figures 1 and 2 in main article).

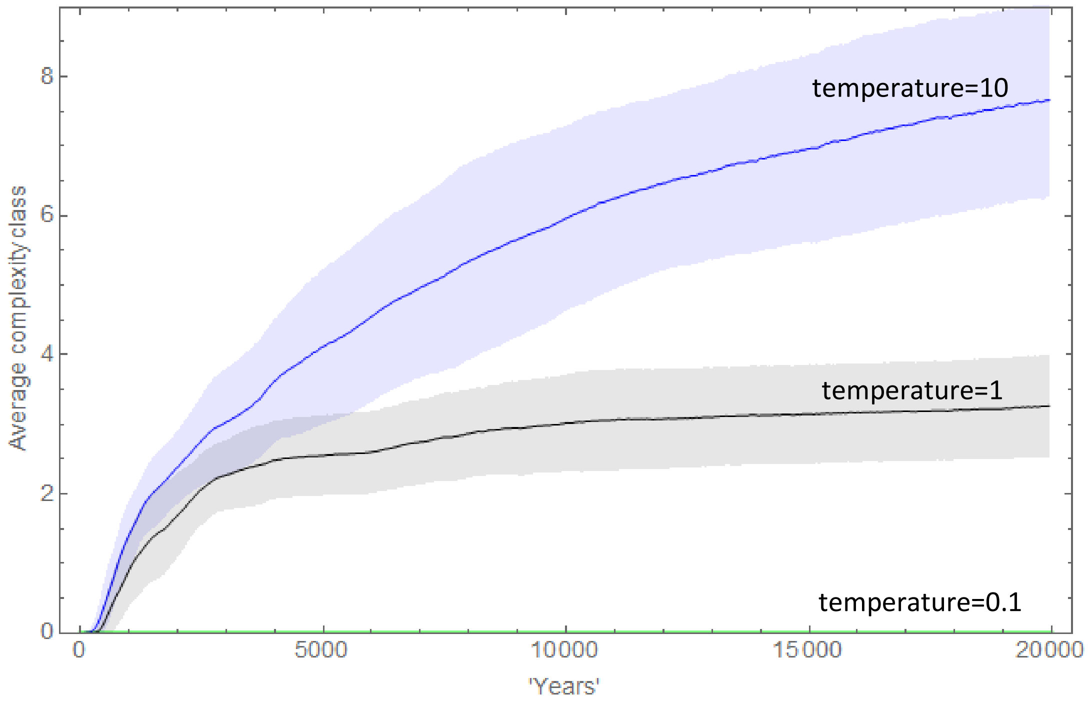

Figure B13 – We here see average complexity class of the IGUT institution in the BASE scenario, varying the strength of conformity bias in social learning. The qualitative dynamics is robust to variations in this parameter, but if increased too much it will eventually cause social learning not to happen at all since novel traditions will always be rare in the population, and the likelihood of learning them will be close to zero. (Temperature=10, blue. Temperature=1, black. Temperature=0.1, green). The black plot is the BASE scenario.

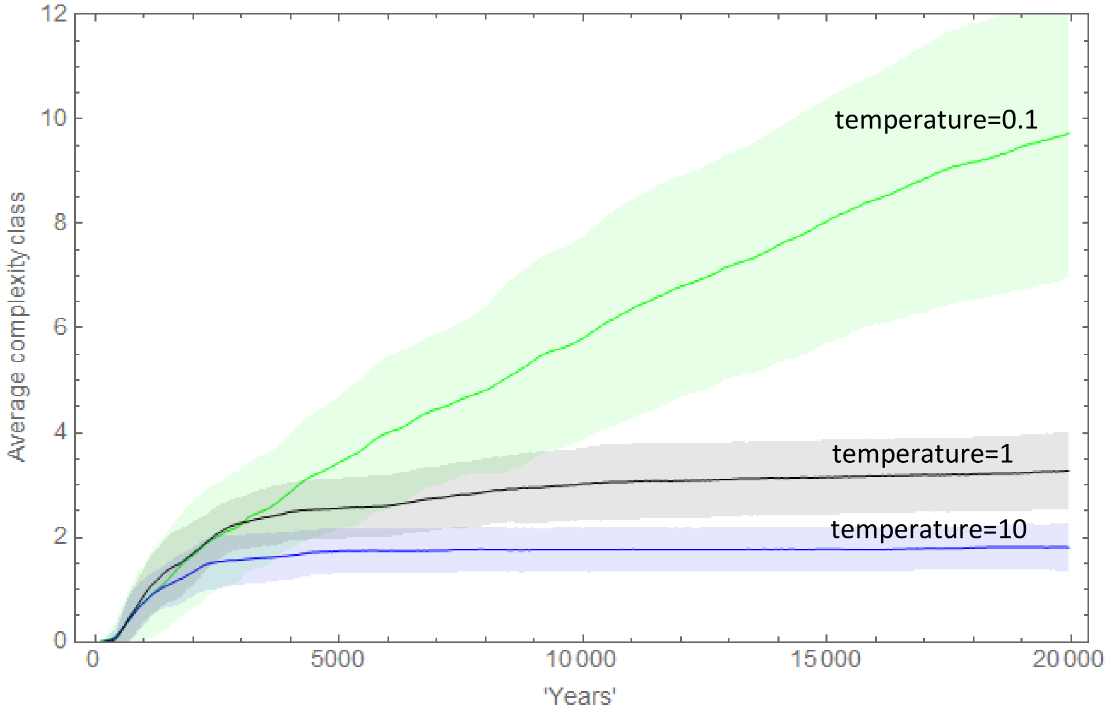

Figure B14 – We here see average complexity class of the IGUT institution in the BASE scenario, varying the strength of rank/prestige bias in social learning. The qualitative dynamics is robust to variations, but cumulativity is quantitatively strongly favored by a low temperature (strong rank bias). (Temperature=10, blue. Temperature=1, black. Temperature=0.1, green). The black plot is the BASE scenario.

Note that we make no statement that this says anything about reality. That would warrant a careful dedicated investigation. For example, while decreasing conformity bias increases fidelity in the model, this could be outright false in a real setting since conformity bias may play an important role in *increasing* fidelity by conserving traditions that many find worthwhile to keep using, while preventing the adoption of new and untried variants.

A low temperature (high conformity bias) will on the other hand make it hard for novelty to spread from its origin. Indeed, the lower value investigated translates into a bias that is so strong that it completely prevents traditions from spreading. Nothing happens at all. The reason is that the likelihood of learning from a lone agent is too low for *any* spread to occur; see Figure B13.

We can conclude that we have not used a conformity bias that specifically favors what we were looking for, namely strong cumulativity. Had we not implemented this bias at all (which corresponds to using a high temperature), we would have arrived at the same qualitative results. What would have corresponded to a very low temperature would have been not to have social learning at all.

Rank/prestige bias temperature

The rank bias temperature (see Appendix SA) governs the strength of the social learning bias for high rank. Alternatively, rank may be thought of as "prestige," since the interpretation depends on what measure we use. Rank is here computed on the basis of success in obtaining resources during lifetime, which is aligned with the idea of prestige as a way of integrating information about successfulness (Henrich & Gil-White, 2001).

If the temperature is low (bias is high), then agents will be unlikely to learn from anybody but the highest ranked individuals. If it is low, then it will not matter as much who they learn from.

In our sensitivity analysis we find a potentially interesting sensitivity to the strength of the rank/prestige bias (see Figure B14). At a low temperature (strong bias) learning proceeds much more rapidly and distinctively than it does if we use the parameter value we use for the results. At high temperature (weak bias), selection is less effective. It should be remarked here that the pattern of increase in the low-temperature case resembles that of the spurious behavior reported in Section *Agent innovation rate* (see Figure B11b) – a pattern that is even more strongly expressed in the Section *Size of the IGUT (carnivory) resource* below. Observing the individual histories that went into the average shown in Figure B14, we see however that this is not the case. New complexity classes reach fixation. However, the rate at which they appear is intermittent, which is the likely answer to the question of why the curve appears straight for a long stretch before the rate of increase begins to abate slightly.

We therefore deem that we see a case of very strong cumulativity that is due to a very efficient spread of traditions within the community, resulting in high fidelity on the sociont level. This appears to be in line with Henrich (2004), who argued that while conformity could guard against behavior that undermines group selection, prestige bias would ensures that innovation is not entirely shut down and will emerge from individuals that have proven their capability. High-prestige individual would then trump the bias against minority variants, while others would be able to transmit only traditions that are already widely adopted.

In this model we have not included spurious alleles, so the conformity bias is not needed to guard against such variants, so our setup is not ideal for an investigation. It could, however, rather easily be adapted for doing that.

Concerning stability, we may observe that we have not selected a parameter value that maximizes the system behavior that we test for. Although the model is sensitive to parameter choice, the dependence is smooth and cumulativity does not break down if we alter it. However, we should note that at very low degrees of bias (high temperature), cumulativity is substantially hampered, which means that this bias may be important for the function of the social protocell.

Tradition cost for agents

Increasing the cost of possessing tradition alleles (in terms of energy) means that the output generated by holding the traditions must be higher. In the end, performance in terms of obtaining energy from resources pays the way. Traditions that do not directly contribute to that must do so indirectly, such as by enabling traditions that do.

Hence, the cost of traditions and the returns of institutions must be balanced such that the dynamics will not be trivial. That is, between (A) a cost that is too high so that traditions will not be selected for at all, and (B), a cost that is too low, so that accumulation will happen by drift. The cost of possessing traditions can be though of, for example, as an opportunity cost when learning and exercising the traditional knowledge.

In between these extremes, the cost of maintaining traditions will affect the achievable equilibrium complexity. This equilibrium, as we discuss in the main text, is due to a break-even where losses (in the form of "failed" offspring that do not inherit a complete set of traditions) reduces the net benefit conveyed by more complex institutions. If traditions are more expensive, this equilibrium will move lower.

As we can see in the plots (Figure B15), halving and doubling the cost of traditions has this effect, which varies smoothly.

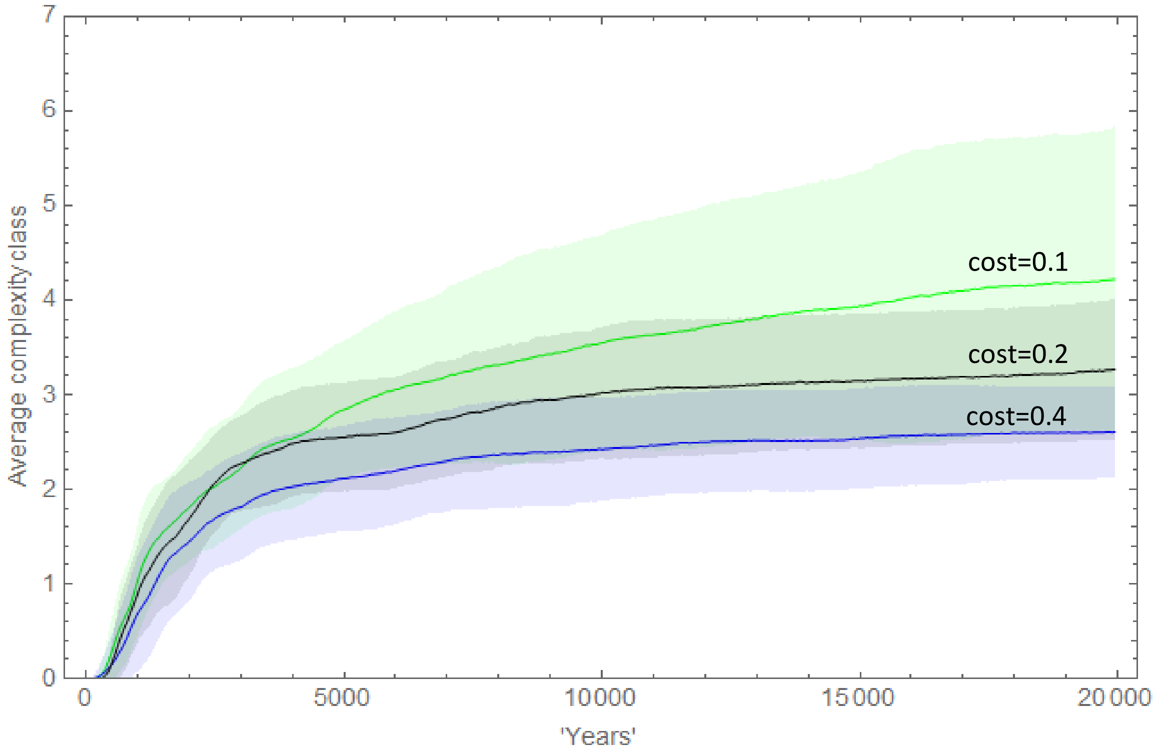

Figure B15 – We here see average complexity class of the IGUT institution in the BASE scenario, varying the cost for agents to possess tradition alleles, in terms of energy. The qualitative dynamics is robust to variations in this parameter. (Cost=0.4, blue. Cost=0.2, black. Cost=0.1, green). The black plot is the BASE scenario.

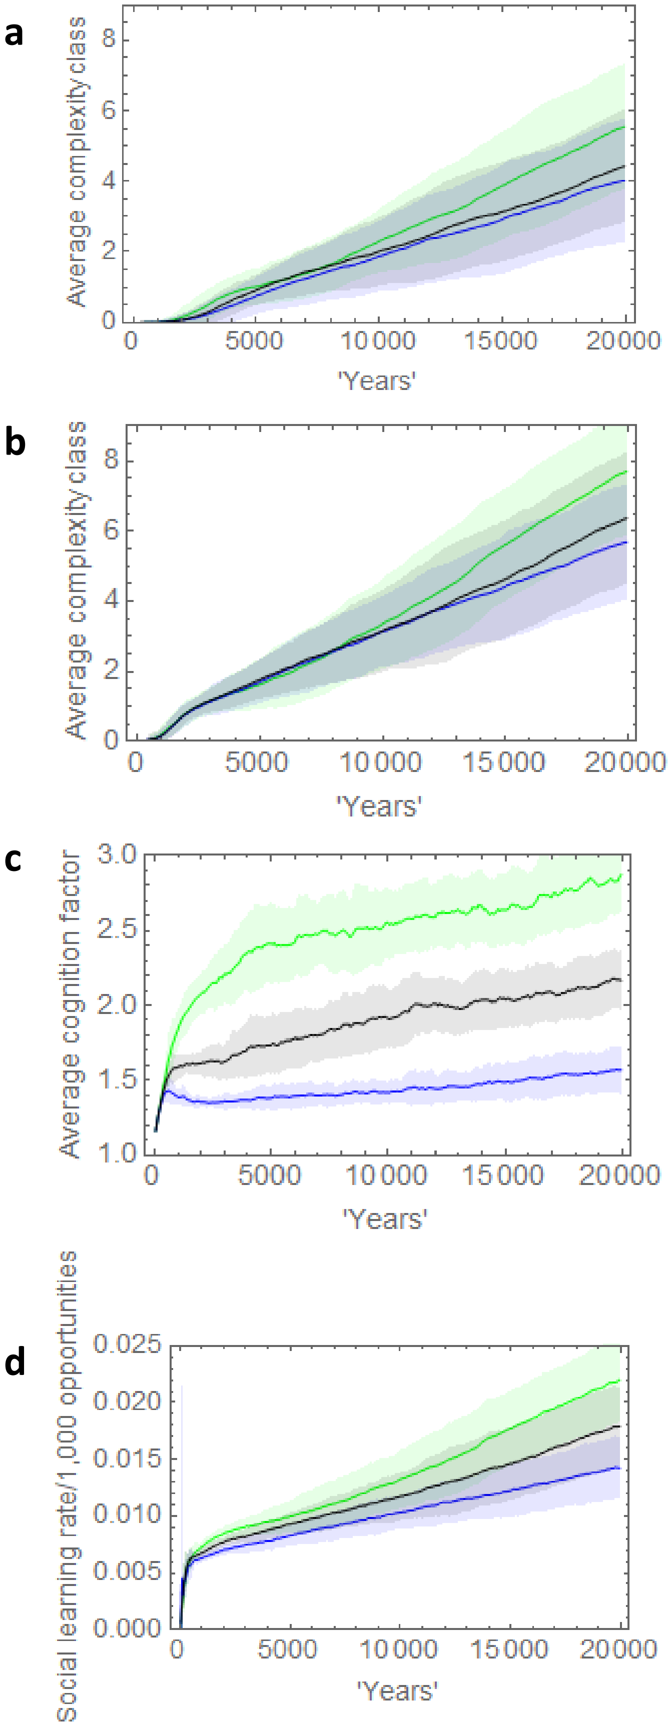

Figure B16– **(a)** We here see average complexity class of the IGUT institution in the BASE scenario variant with the SLR boosting institution, varying the cost of cognitive capacity for benefitting from said institution. The qualitative dynamics is robust to variations in this parameter (Cost=0.4, blue. Cost=0.2, black. Cost=0.1, green). The black plot is the BASE scenario. **(b)** Same plot but for the SLR boosting institution. **(c)** Average cognition factor is here shown for the same parameter variations as above. The response is as expected. **(d)** The rate of social learning, as expected, follows the increase in the complexity of the SLR boosting institution together with the increase in agent cognition.

Cognition cost for agents

Cognition cost applies only when the SLR boost institution is active (see *Results* in the main article). It determines the cost of possessing genetic cognitive adaptations, as a function of the level of cognition delivered by those adaptations (i.e. the brain). Cognitive performance has only one effect in the model. It determines the SLR as a function of the complexity of the SLR boost institution, including the case of no boost (see *Innate capacity for social learning* in Appendix SA). The effect of this mapping is that for every given level of cognitive performance, there is an optimal level of SLR boost. Increasing the complexity of the SLR boosting institution to this point yields increasing SLR, but beyond that point, the SLR will go down again. This provides a simple way of implementing a necessity for co-evolution between cognitive capacity and culture.

If the cost of cognition is high, cognitive capacity (in the sense used here) will be more strongly selected against on the agent level. This makes it harder for group benefits to punish those reaping short-term fitness benefits by having varied to lower and thereby cheaper levels of cognitive capacity. Such capacities would in reality have various other and more direct adaptive advantages also on the individual level, but we add no such extra benefits in the model.

Cognition level may for this reason be expected to fluctuate. When high cognition gets established, it will constantly be challenged by agents that unilaterally “defect” via variation to lower cognitive power, and thereby lower costs. On the individual level, assuming no other benefits, paying for high cognition is not an evolutionarily stable strategy in itself. Occasionally, these defectors will prevail since they are adaptive in the short term. However, socionts underpinned by communities where defectors have spread will be weeded out in higher-level selection – if the sociont benefits sufficiently from higher cognitive capacity. The question is whether this happens quickly for spread to occur on the higher level, and for high-cognition socionts to have sufficient time to grow and divide to take over on their level.

If the cost of cognition is higher, then this challenge from “defectors” will be stronger. Eventually, cognition will not be selected for, when the questions posed above has a negative answer. If the cost is lower, then cognitive capacity will be more easily selected for as the motivation to defect is lower.

The effect of halving and doubling the cost, as we can see in Figure B16, is smooth and predictable. At a lower cost, slightly higher rates of social learning, cognitive ability and cultural complexity are achieved. At a higher cost, the opposite happens. The selected parameter value for the results (BASE) is not finetuned to achieve the reported behavior.

Agent survival as a function of energy level

We use a custom and quite detailed piece-wise linear function to map agents’ energy levels to a multiplicative component of their death rate. This function is not parametric and thereby hard to vary systematically. It is designed to achieve specific properties of its shape. The shape sought was that when the energy input is lower than the cost of living, the agents should suffer an elevated mortality that increases as energy levels fall.

Mortality, notably, is not *net* zero when needs are met ($energy \geq 1$). It is only the component of mortality that stems from lack of energy that is zero at that point. This factor is then combined multiplicatively with a death rate that depends on age (see Section *Within the community*).

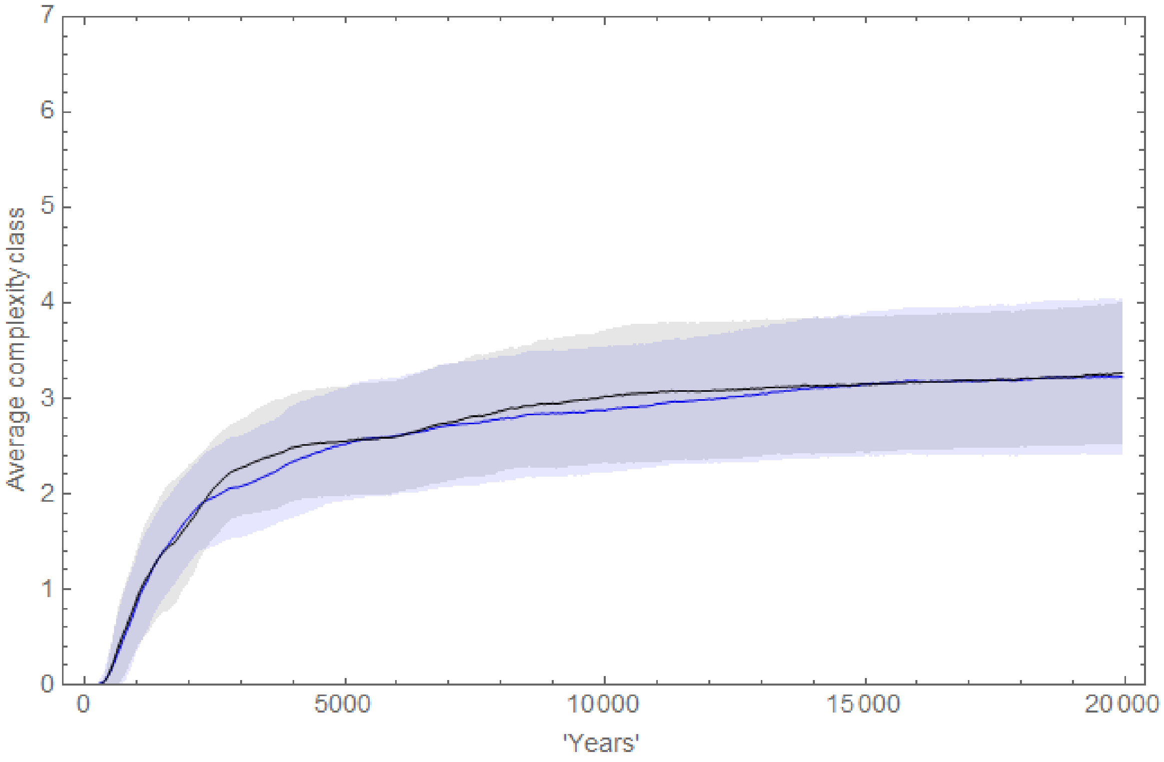

Figure B17 – We here see average complexity class of the IGUT institution in the BASE scenario, varying the shape of the function that contributes the energy-level related component of mortality. The model is robust to variations in this function's form, as long as it conforms to the stated boundary conditions. (Linear mortality function, blue. Custom built mortality function, black.) The black plot is the BASE scenario.

The motivation for using a designed rather than a parametric function is that there is no reason to believe that the phenomenon is governed by a simple process for which a good simple mathematical model can be formulated as a function of some single factor. In other words, seeking a function that approximates the shape we need is to design it anyway. The detailed shape of this function is, however, not likely to be important as long as the mortality increases suitably when agent energy level goes below the cost of living (which is unit,) and approaches unit when energy approaches the cost of living.

We here investigate whether the results are sensitive to the shape of this function. The shape can be varied endlessly, so we have chosen to compare with a very simple linear function, where the contributed mortality is zero when the agent has full energy and unit what it has no energy. As can be seen in Figure B17, it makes no difference if we choose a much simpler shape of this function.

Agent survival as a function of age

Life for the energy level of agents, we use a custom and quite detailed piece-wise linear function also for the other multiplicative component of the death rate of agents. Essentially, an elevated death rate should apply for juveniles, then a plateau in adult age, and a rapidly increasing death rate for higher ages. Similar argument applies here as in Section *Agent survival as a function of energy level.*

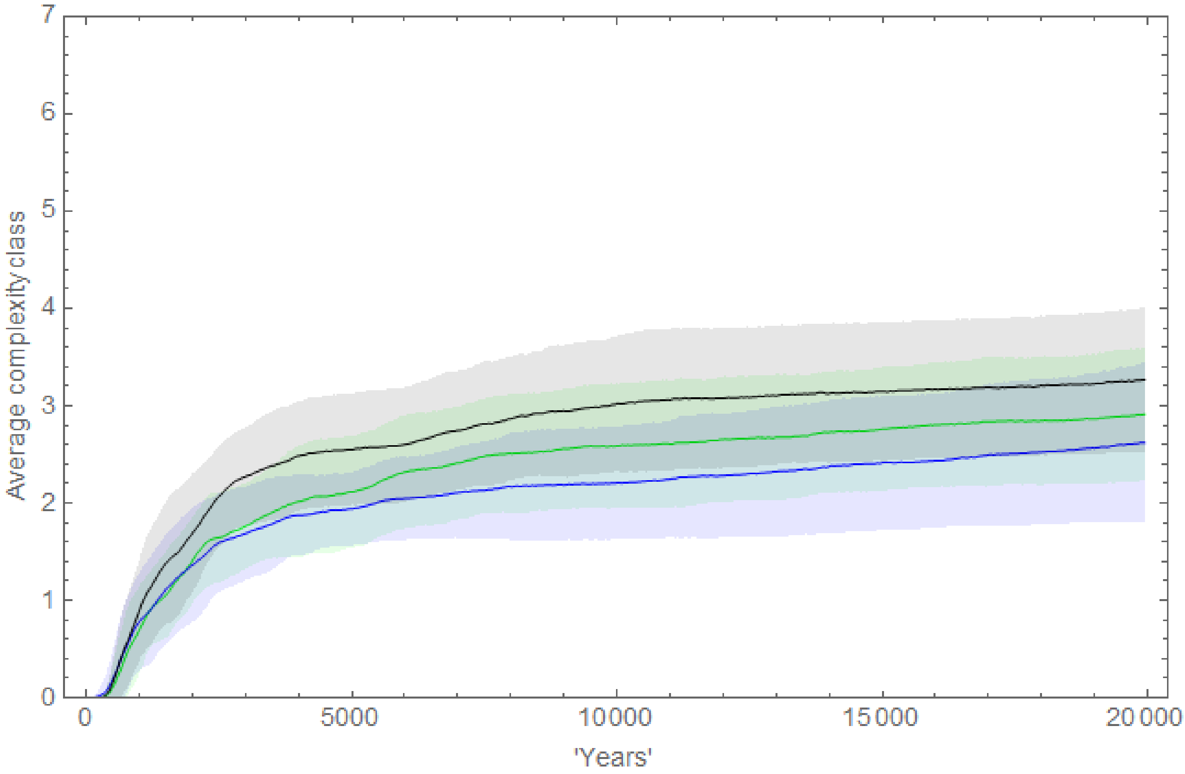

Figure B18 – We here see average complexity class of the IGUT institution in the BASE scenario, varying the shape of the function that contributes the age-related component of mortality. The model is robust to variations in this function's form. (Threshold mortality function, blue. Chimp-like mortality function, green. BASE version of the function.)

The shape can also here be varied endlessly so we compare first with a simple threshold function whereby all agents are guaranteed to survive until age 50 (if they have sufficient energy), at which point they drop dead as if struck by lightning. That is, the function returns unit from birth until an age of 50 "year," and after that it returns zero. As can be seen in Figure B18 (blue plot), it makes no qualitative difference if we choose this simpler shape of this function.

In this case, however, there is data available that we may compare to. We therefore also test a piece-wise linear function that is designed to match the shape reported for chimpanzees by Gurven and Gomes (2017, p. 185). In Figure B18 we see that also this variation produces no qualitative impact on the outcome. The quantitative differences compared with the BASE case, in both cases, are likely due to the different functions implying different average lifespans, with the chimp-like scenario yielding the shortest.

Probability of attacking cell taking over a neighbor cell

The rate at which attacking cells take over neighboring cells in the spatial competition model (see *Space and ecology* in Appendix SA) is also governed by a designed stepwise linear function. This function maps the relative strength of an attacking cell to a probability that it will occupy an attacked neighbor cell.

The function is designed such that a weaker attacker will almost never succeed, while a stronger attacker will progressively more often succeed. While the reasoning, in outline, is robust almost by definition (reversing the logic is clearly empirically unmotivated,) several other assumptions must be made. Not least, *how much* does strength affect the likelihood of taking over?

We have varied the function by using **(A)** a simple step function where the probability of winning is unit for the attacking cell if it is stronger and zero if it is weaker, and **(B)** a substantially weaker bias in favor of the stronger part.

As we can see in Figure B19, the step function produces practically speaking identical results to the more detailed function used in the results. The same is true for the weaker function where the dominant part has a smaller advantage. We see only a slight quantitative deviation in the expected direction, namely that the step function would produce stronger selection, and thereby more cumulativity, while a weaker advantage for a superior attacker would weaken selection.

In other words, the setup used for the results of the main article was more detailed than needed in this regard. This run also tells us that the SPH does not seem to hinge on strong assumptions of very high degrees of territoriality. Even weaker competition for territory seems to work, although this calls for closer and systematic investigations.

Size of the IGUT (carnivory) resource

Resource size sets the total amount of a resource that may be extracted by agents, and resources are distributed homogeneously across the system in this setup. Consequently, the smaller a resource is (*ceteris paribus*), the more area must be controlled in order to obtain some specific amount of it.

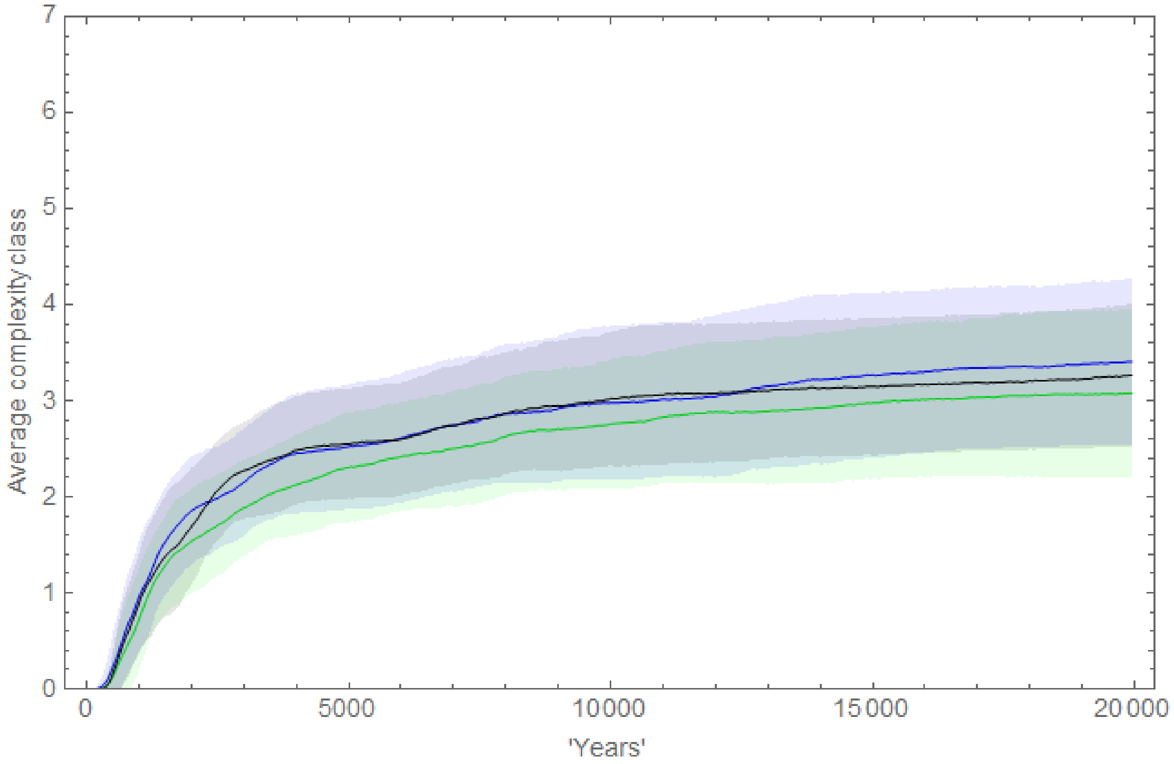

Figure B19 – We here see average complexity class of the IGUT institution in the BASE scenario, varying the shape of the function that governs the likelihood with which an attacking spatial cell takes over a neighbor, as a function of their relative strengths. The model is robust to variations in this function's form. (Step function, blue. Weaker advantage for strong cell, green. BASE version of the function.)

The parameters are tuned such that the "basic resource" itself provides enough resources for initial agents (without culture) to survive at some density that corresponds to a population of some tens of communities as a basis for evolution to take place. Generally, at any given level of efficiency in extracting the IGUT, the smaller the IGUT resource is, the smaller will the population of socionts be, since they will need more area to maintain themselves. Recall that once more efficient strategies are established in the population of socionts, they become mandatory since if you lose them, your competitors will still be using them.

The dynamics of the model should be relatively independent of the size of the IGUT, as long as it qualifies as “large.” What that means is essentially that it can potentially feed a substantially higher population density than what is possible without it.

However, since the resolution of the system is finite, a too large IGUT will mean that socionts will be associated with territories so small that they approach the dispersal threshold for area faster, which means that the simulation cannot be usefully run for very long. Also, we will more rapidly approach the point where the model begins to break down due to poor spatial resolution (see Section *Varying the spatial resolution*).

The value used in the BASE case has been tuned such that the IGUT is competitive, and that territories do not become too small until a substantial complexity has been reached (allowing us to investigate the dynamics).

When we vary the size of the IGUT we also discover a number of things about the model that does not necessarily have much to do with the size of resources as such. Before going into that, however, we note in Figure B20 that there is no qualitative difference between the outcomes of the resource size used in the results of the paper (black plot), and a size half of that size (green plot).

As we double the size of the IGUT resource (blue plot), we see that the configuration will rapidly be packed (as predicted,) which, in turn, leads to a situation with tiny territories where fluctuations of many kinds dominate (too low spatial resolution causes the model to break down). One such source of fluctuations has to do with asymmetric splitting, as described in Section *Varying the spatial resolution*.

Running the model with a tiny IGUT size of 50 000 units (red plot), we bump into a very strong case of the unintended behavior we saw when investigating innovation rate (see Section *Agent innovation rate*). In this case, an almost inert state appears because a small IGUT means that increasing the complexity of the institution in question will just barely pay off. Consequently, variants with higher complexity will survive stably, but they will just sit there and, for the reasons discussed above (Section *Agent innovation rate*), this causes complexity to just keep growing. As stated, this would be easily amended by introducing a random process that prevents indefinite lifespans of communities, which would also be more realistic.

Fortunately, the deviations from nominal model behavior that we detected here do not affect the reported results.

IGUT (carnivory) access function shape

The access function determines how much of a resource that is accessed with the fitness of the targeting institution as argument. This access function is a simple linearly increasing function of fitness, which, in turn, is a linear function of its complexity (in terms of how many linked component loci that it consists of.) So, in effect, complexity is fitness in the used setup.

We have tested two additional cases (see Figure B21). One where maximum access is reached already at a fitness level of 5 and one where it is maxed out at a level of 100 (which is considerably more than can be achieved). With the maximum being fixed, this means that we alter the slope of the access function.

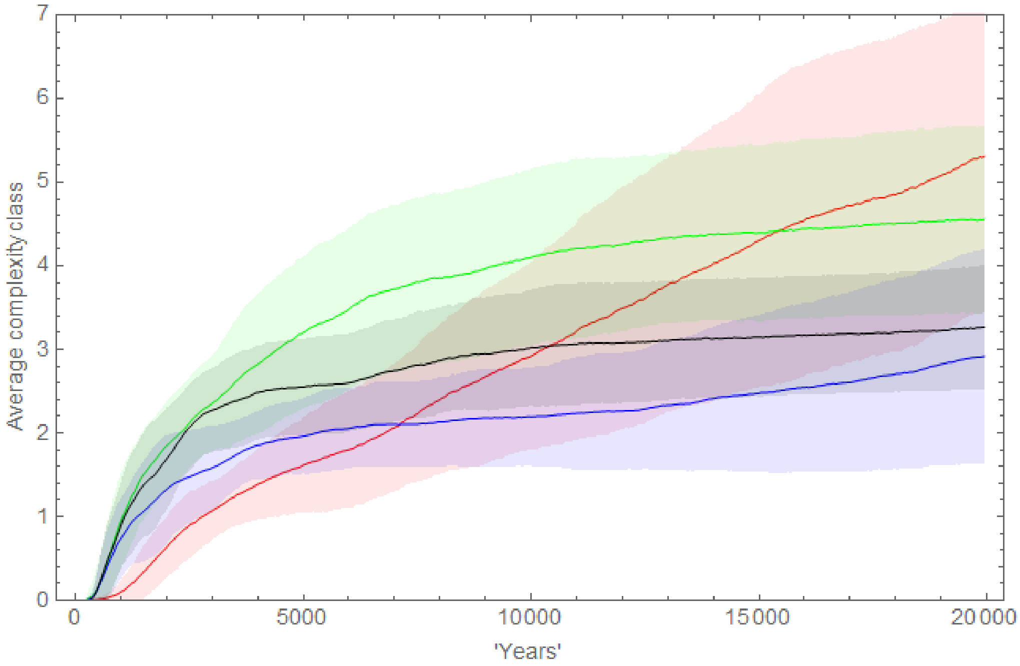

Figure B20 – We here see average complexity class of the IGUT institution in the BASE scenario, varying the size of the IGUT (carnivory) resource. The model is robust to variations - since both systematic deviations seen here are due to explainable model malfunctions that do not affect the results. (Size=2,000,000, blue. Size=1,000,000, black. Size=500,000, green. Size=50,000, red. Black plot is the BASE value of the parameter.

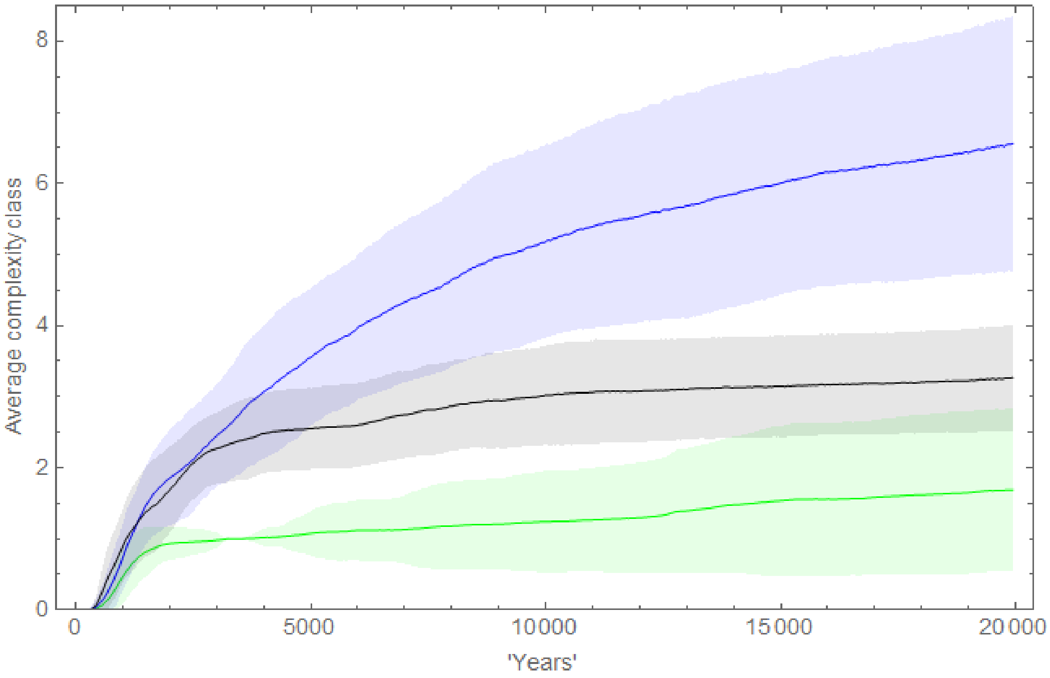

Figure B21 – We here see average complexity class of the IGUT institution in the BASE scenario, varying the slope of the IGUT (carnivory) resource access function (function of institutional complexity). (Low slope function with max at complexity=100, blue. BASE case, black. High slope function with max at complexity=5, green.

Lowering the slope (blue plot) means that less additional access to the resource is gained with each increase in institutional complexity. Intuitively slow down the dynamics, but our runs show that due to the formulation of the model, this is not quite what happens. The sweeps occur unabated in the scenario with a slow increase per unit of institutional complexity. Why?

The basic reason is likely that the edge will still be large enough that an attacker with an IGUT institution with higher complexity will achieve likelihoods of success similar to with a stronger edge. A small edge is sufficient. This may well not be realistic, not least as there may be all manners of additional types of benefits of better access to resources.

In terms of robustness to this parameter, it seems worrying that lowering the slope of the access function itself causes the model to achieve substantially *higher* levels of complexity than in the base case. It is worrying since it is not immediately obvious how this goes with the story of the results, which is that fidelity controls the maximum complexity that is reachable. Indeed, if we increase the slope (green plot), complexity maxes out very quickly.

Somehow, lower slope must yield stronger selection? The answer appears to be that in the low slope case, (a) the increases in efficiency are sufficient for outcompeting other territories, while (b) territories are considerably larger (lower access at the same level of complexity), so they suffer less from the imperfection in the splitting mechanism that we have already discussed (see *Varying the spatial resolution*, and *Size of the IGUT (carnivory) resource*). This imperfection seems to add to the loss of fidelity more than we expected.

This is a finding that must be further explored in future models. However, this effect would simply add another (unwanted) source of error in transmission, beside those that are internal to the model and the SPH (i.e., related to the likelihood of all needed loci surviving a split). Thereby, it does not invalidate the results. The behavior still varies as predicted as we vary the settings according to the logic of the SPH. But it does tell us that we could isolate the effect more efficiently in a model where there is no side-effect of finite resolution that also produces the phenomenon that we are looking for. More specifically, as we have already noted, this has to do with that the splitting algorithm works less and less well the fewer discrete geographical cells that the communities are made up of.

So while this is an artifact of the model formulation, there is no reason to suspect that amending it would invalidate the results. The model also behaves qualitatively in the same way as we vary this parameter – we do not need to find a narrow range of values to make it work. In particular, the complexity still tends toward an asymptote in the same way. This means that the results would just have different numerical levels of complexity at equilibrium if a different slope was used. Also, if the splitting function would not deteriorate due to finite resolution, fidelity would still decrease with increasing complexity.

The slope of the function in the runs used for the results is designed to max out at a complexity of 25, which is more than the complexity anticipated in the runs, permitting a rich evolutionary dynamic to play out, with several selective sweeps.

Monopolizability of the IGUT (carnivory) resource

The higher the degree of monopolizability of a resource, the more will the performers of apex alleles (that actually produces a payoff) increase their fitness compared with those performing their supporting component alleles. This severely limits the link between fitness contributions of institutional components and the fitness of the sociont via the biological fitness of the agents.

In reality high monopolizability causes a decreased incentive to cooperate by performing secondary functions that will not position you well to grab the payoff. It is in this context that large carcasses have, since long, been considered important for human evolution. They have a very low degree of monopolizability for a relatively small animal like hominins (Blurton Jones, 1984; Winterhalder, 1996). Hominins would have been neither able nor motivated (unable to consume more than parts regardless) to monopolize large carcasses.

This is not the important effect here, however, since cooperation between agents is not modeled. The important effect here is whether the large amounts of food made available will be efficiently utilized by the group.

Decreasing monopolizability we see that competition on the sociont level increases dramatically, driving up the level of complexity that is reached. This effect is expected, but it would probably be less pronounced in reality. The main reason is that resources are translated fairly directly into increased rates of birth in the model, and the portion that cannot be monopolized is spread to precisely all members of a community equally.

In conclusion, what we see here is potentially important to explore further. A low degree of monopolization should probably be added to the "Importance" criterion of the IGUT as a driver of a cultural ETI. Monopolizability has a very pronounced effect. The degree of monopolizability that we have used in the BASE case does not seem overly low at 50%, and, if anything, the value should be higher. As we see in Figure B22, we clearly have cumulative evolution in the BASE case, but we get much more cumulativity if we had chosen a lower value.

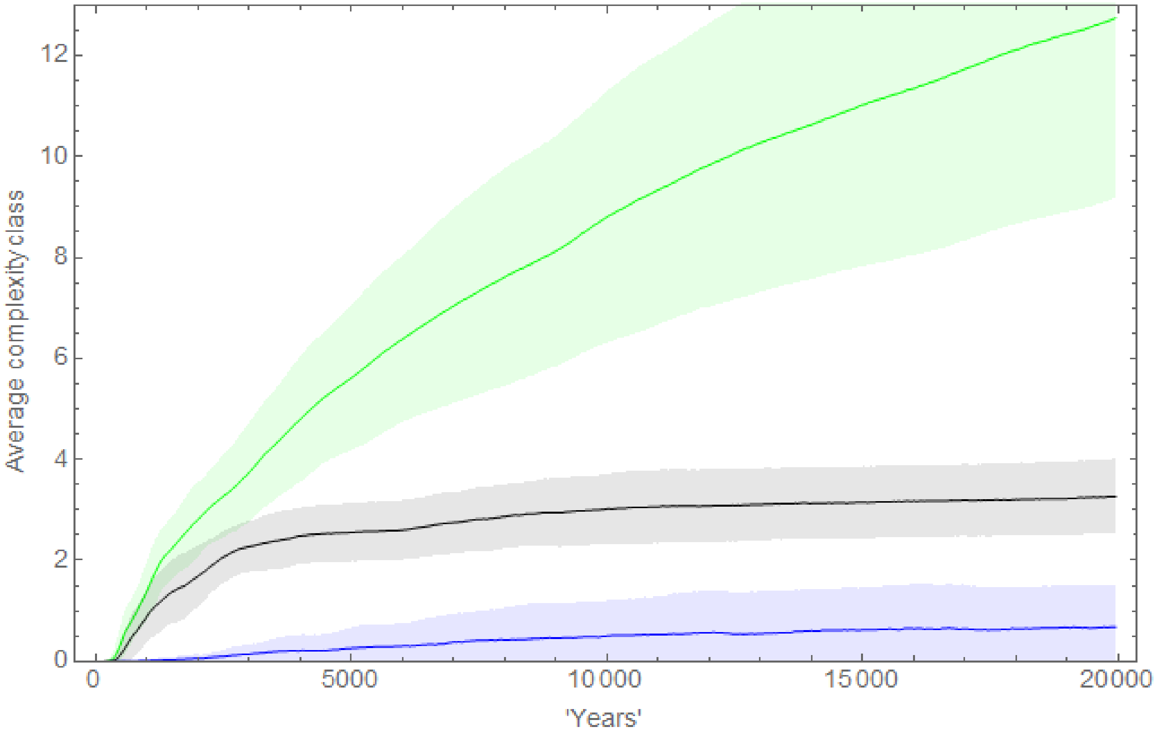

Figure B22 – We here see average complexity class of the IGUT institution in the BASE scenario, varying the degree of monopolizability of the IGUT (carnivory) resource. (High monopolizability=0.75, blue. BASE case monopolizability = 0.5, black. Low monopolizability=0.25, green).

Being of age

Age plays a limited role in the setup used here, but an arbitrary limit of eight "years" has been used as a limit for when reproduction begins and when agents begin contributing to the territorial defense of the community. The reason why such a limit was added was to avoid clearly unrealistic explosions of population, which may lead to instabilities, such as communities with increased harvesting efficiency overshooting in terms of population and going extinct. We see only a small quantitative impact of varying this parameter as long as we avoid such effects (see Figure B23). While this parameter has to do with childhood, which is important in human evolution (e.g. Bogin, 1990; Hawkes et al., 2008; Thompson & Nelson, 2011), the model is not otherwise equipped to investigate the impact since that would demand a more detailed model of the structure of social learning, food distribution, and so on.

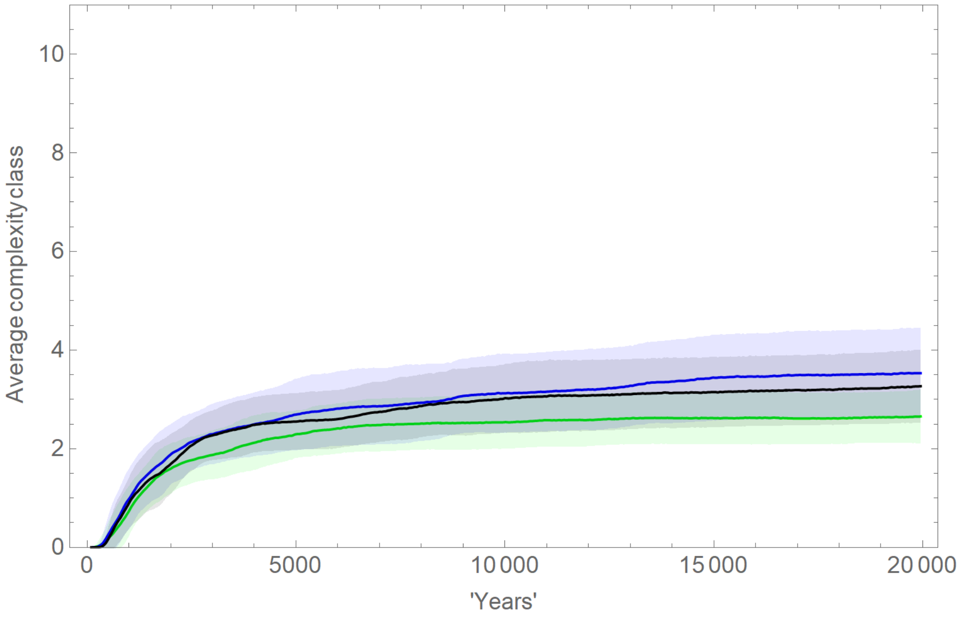

Figure B23 – We here see average complexity class of the IGUT institution in the BASE scenario, varying the age at which agents become fully contributing. (Ten "years": blue. BASE case, eight "years": black. Six "years": green).

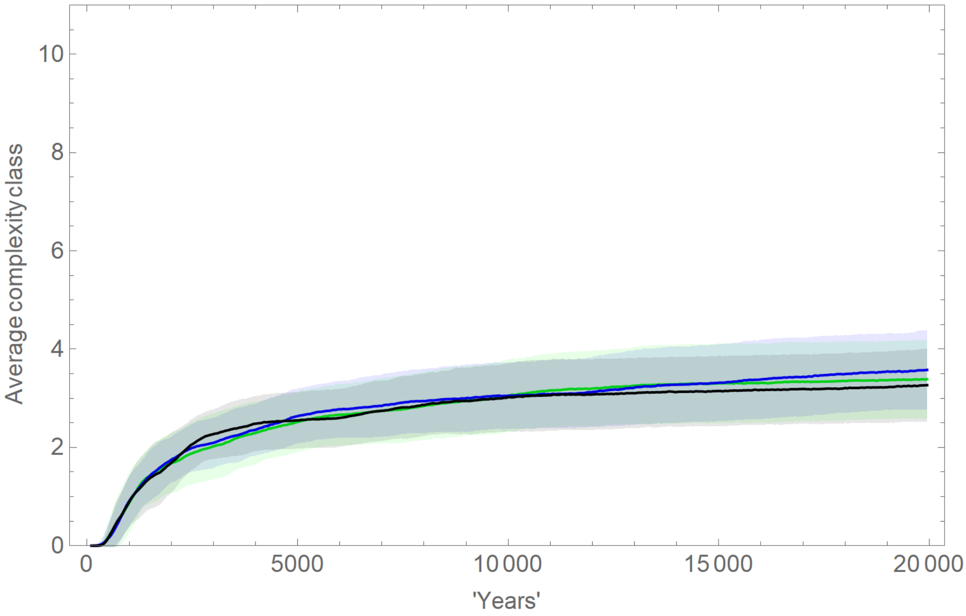

Figure B24 – We here see average complexity class of the IGUT institution in the BASE scenario, varying the minimum time between births. (Three "years": blue. BASE case, two "years": black. One "year": green).

Minimum time between reproductions

Setting a minimum time between reproduction events seems prudent for the same reasons a stated in *Being of age*. In particular since there is no gestation time in the model, only the time it takes to fill the reproduction buffer (see *Reproduction* in Appendix SA). Also here we see no important effects when we vary the parameter (see Figure B24). Birth intervals in general is far from unimportant (e.g. Miller et al., 2019; Nakahashi et al., 2018; Robson & Wood, 2008), but the model is not designed to capture these effects.

Omitted parameters

No bias against component traditions (which do not produce an immediate payoff)

A willingness to learn actions that are not motivated by an immediate payoff is assumed in the model, and it should be noted that this is an assumption that itself needs to be examined. Animals do not willingly learn behaviors that they do not understand that they will be rewarded for.

Since almost all the learning that takes place is of component traditions, such a bias would in the model be almost indistinguishable from generally decreasing the rate of social learning. We have already seen the effect of doing that: it hampers cumulativity.

In particular as more and more complex institutions arise, it seems necessary that this would have entailed evolution of the agents to be willing and able to, also, socially learn activities with indirect and delayed payoffs. This may include, for example, the evolution of social rewards (normative conformity and prestige bias; see e.g. Henrich, 2004).

Mutation rate

Genetic elements in reproduction (see Appendix SA) are subject to a mutation rate. We only use a single genetic feature, in some of the setups, and we have not investigated the impact of tuning this variable since it seems unlikely that results would be affected in non-trivial ways. For example, evolution of this feature is turned off by setting the mutation rate to zero, which causes no variation in the population. As soon as the parameter is increased there will be variation and selection. Unless the fidelity boosting institution is activated, increased cognitive capacity is only costly and will in such cases be selected against. In setups with more genetic inheritance, this parameter may be of more importance, but in the setup used here we have deemed that it simply needs to be of some sufficient magnitude.

Social learning rate

Variation of the basic social learning rate is investigated in the results section of the main article (see Figure 7) and therefore we do not investigate it here.

Summary

In this appendix, we have tested the model for robustness to all choices of parameter values and some model formulation choices. Overall, we think the tests show that the BASE case is representative for the behavior of the model within a large part of the parameter space. Moreover, we think it also shows that the response of the model to these parameter variations are predictable from the basis of the SPH that the model aims to implement.

We found a number of unwanted model behaviors that we managed to understand when and why they appear.

One such effect was that deviations because of limited spatial resolution will kick in more and more the smaller the territories are. In particular, the fidelity of sociont fission is decreased by poor resolution (the divided territory occupies a small number of discrete territorial cells.) Since territories shrink as efficiency increases from cumulative evolution, this effect may partly be responsible for decreasing fidelity. However, that effect is unlikely to be large in the runs used for the results as we can readily see that much higher levels of cultural complexity are reached in runs with different parameter settings. For the BASE case, territories never become very small.

Another such effect was that if, for some reason, selection is weak on the sociont level, the fact that communities never "die" for other reasons causes almost indefinity accumulation of cultural complexity in a way that has nothing to do with reality. This problem is very likely solvable by simply introducing an external source of risks. However, for the BASE case, and most variations around it, this effect does not appear, and it is not the reason for cumulative evolution in the results. These spurious cases of cumulativity have a telltale signature in the data, since the increasingly complex socionts do not spread but gain complexity in *situ* (see Figure B11b).

A number of suggestions for further work can also be mined from this analysis.

Some of the tests indicate that the SPH can be based on laxer assumptions. In particular:

- The SPH does not seem to need to assume a very high intensity of territorial behavior (see *Probability of attacking cell taking over a neighbor cell*). That is, the hypothesis seems not to hinge on that our early ancestors lay very close to chimpanzees rather than bonobos in terms of territorial behavior.
- The SPH does not appear to hinge strongly on the assumption that community population size is strongly tied to the rate of community fissions (see *Criterion for community fission*).

Some omissions in the model that should be explored include that there is no migration between the communities. The model in fact affords this functionality, subject to a migration rate parameter, but this invokes a host of additional assumptions that must be made. It was therefore left for future work. Runs made with this feature active have not shown that cumulative evolution is particularly sensitive to migration.

Another omission is that cheating is not modeled. This is a major omission and should be subject to studies in the future. Such a functionality has been created in the model but was not used, also for the reason that several additional assumptions were needed. The mechanism is that alleles can turn into "anti-alleles" that are indistinguishable from their real counterparts but have no function and no cost. Adopting them is therefore adaptive. Preliminary runs have shown that institutions that detect and punish these anti-alleles evolve, much like the SLR boosting institution. While the agents are not assumed to necessarily cooperate in the model, cooperation is clearly implied in practice for complex institutions. Cooperation between agents in the model should also be investigated. Both potentials for cheating represent potential problems for the SPH.

We also found that low monopolizability strongly favored cumulative sociont evolution, which is a finding that makes sense logically from the standpoint of the SPH, and from basic ecology. This issue should be studied separately. We also found that there may have been a selection for larger communities (See *Community splitting population threshold* and *Criterion for community splitting*) because larger communities will have higher fidelity.

For future work it would be interesting to investigate the interplay between prestige and conformity in the model. However, in the present setup, traditions never have any undermining effect, so the positive effect of high conformity bias is missing. Such a study would therefore need to introduce that possibility. Doing that, the *Mode of community splitting* could also become important, which as we saw it was not in this setup.

It would also be possible to have a considerably more sophisticated model with regard to how food sharing happens and how coalitions are formed internally. For example, resources are not explicitly directed at juveniles and mothers in this setup, so there are no parameters that can be varied. The impact of institutions and genetic adaptations having to do with collective breeding, childhood, and altriciality could be investigated.

In many cases we discover that we could have used a simpler model with less mechanisms and sub-models. However, it should be noted that if we view the model as a platform, then some of these parameters may become more critical if other investigations are made. Also, say that we had, for example, used a simple step function in the spatial competition model (see Section *Probability of attacking cell taking over a neighbor cell*). The model would have been simpler, but that would also have been an assumption, and it would be much harder to vary the assumption to test it.

Bibliography

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

Bogin, B. (1990). The Evolution of Human Childhood. *American Institute of Biological Sciences*, *40*(1), 16–25. https://doi.org/10.2307/1311235

Byrne, R. W., & Whiten, A. (1989). *Machiavellian Intelligence: Social Expertise and the Evolution of Intellect in Monkeys, Apes, and Humans*. Clarendon Press.

Dunbar, R. I. M. (1992). Neocortex size as a constraint on group size in primates. *Journal of Human Evolution*, *22*(6), 469–493. https://doi.org/10.1016/0047-2484(92)90081-J

Dunbar, R. I. M. (1995). Neocortex size and group size in primates: a test of the hypothesis. *Journal of Human Evolution*, *3*(28), 287–296.

Dunbar, R. I. M. (1998). The Social Brain Hypothesis. *Evolutionary Anthropology*, 178–190. https://doi.org/10.1002/(SICI)1520-6505(1998)6:5<178::AID-EVAN5>3.3.CO;2-P

Epstein, J. M. (2007). *Generative Social Science: Studies in Agent-Based Computational Modeling (Princeton Studies in Complexity)*. Princeton University Press.

Feldblum, J., Manfredi, S., Gilby, I., & Pusey, A. E. (2018). The timing and causes of a unique chimpanzee community fission preceding Gombe's Four Years' War. *American Journal of Physical Anthropology*.

Feldblum, J. T., Manfredi, S., Gilby, I. C., & Pusey, A. E. (2018). The timing and causes of a unique chimpanzee community fission preceding Gombe's "Four-Year War." *American Journal of Physical Anthropology*, *166*(3), 730–744. https://doi.org/10.1002/ajpa.23462

Gurven, M. D., & Gomes, C. M. (2017). Mortality, Sensescence, and Life Span. In M. N. Muller, R. W. Wrangham, & D. R. Pilbeam (Eds.), *Chimpanzees and Human Evolution*. Belknap Press, Harvard University.

Hawkes, K., City, S. L., Spindle, W., & Bogin, B. (2008). The Birth of Childhood. *Science*, *322*, 1040–1043. https://doi.org/10.1126/science.322.5904.1040

Henrich, J. (2004). Cultural group selection, coevolutionary processes and large-scale cooperation.

Journal of Economic Behavior and Organization, *53*, 143–162. https://doi.org/10.1016/S0167-2681(03)00094-5

Henrich, J., & Gil-White, F. J. (2001). The evolution of prestige: Freely conferred deference as a mechanism for enhancing the benefits of cultural transmission. *Evolution and Human Behavior*, *22*(3), 165–196. https://doi.org/10.1016/S1090-5138(00)00071-4

Lehmann, J., & Boesch, C. (2004). To fission or to fusion: effects of community size on wild chimpanzee (Pan troglodytes verus) social organisation. *Behavioral Ecology and Sociobiology*, *56*(3), 207–216. https://doi.org/10.1007/s00265-004-0781-x

Miller, I. F., Churchill, S. E., & Nunn, C. L. (2019). Speeding in the slow lane: Phylogenetic comparative analyses reveal that not all human life history traits are exceptional. *Journal of Human Evolution*, *130*, 36–44. https://doi.org/10.1016/j.jhevol.2018.12.007

Moffett, M. W. (2013). Human identity and the evolution of societies. *Human Nature*, *24*(3), 219–267. https://doi.org/10.1007/s12110-013-9170-3

Nakahashi, W., Horiuchi, S., & Ihara, Y. (2018). Estimating hominid life history: the critical interbirth interval. *Population Ecology*, *60*(1–2), 127–142. https://doi.org/10.1007/s10144-018-0610-0

Read, D. W. (2012). *How culture makes us human: Primate social evolution and the formation of human societies*. Left Coast Press, Incorporated.

Robson, S. L., & Wood, B. (2008). Hominin life history: Reconstruction and evolution. *Journal of Anatomy*, *212*(4), 394–425. https://doi.org/10.1111/j.1469-7580.2008.00867.x

Thompson, J. L., & Nelson, A. J. (2011). Middle Childhood and Modern Human Origins. *Human Nature*, *22*(3), 249–280. https://doi.org/10.1007/s12110-011-9119-3

Whiten, A., & Byrne, R. W. (1997). *Machiavellian Intelligence II*. Cambridge University Press.

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