Extreme Environments Perpetuate Cooperation

Brandon Gower-Winter
Department of Computer Science
University of Cape Town
Cape Town, South Africa
GWRBRA001@myuct.ac.za

Geoff Nitschke
Department of Computer Science
University of Cape Town
Cape Town, South Africa
gnitschke@cs.uct.ac.za

***Abstract*—We investigate whether environmental stress positively impacts the emergence of cooperative behaviour in socially stratified societies. We achieve this by utilizing NeoCOOP, an Agent-based Model that uses artificial evolution as adaptive mechanisms to simulate the emergence and evolution of altruistic and selfish behaviour in Neolithic-inspired agents. We perform scenario experimentation whereby we monitor the resource trading preferences of these agents by varying the frequency of environmental stress and the initial beliefs of said agents. Our results indicate that in extreme conditions, altruism is preferred. Furthermore, our results suggest that the degree of social stratification of a population is positively related to its ability to maintain logistic-like growth while remaining susceptible to environmental stress.**

***Index Terms*—Agent-based Modelling, Cooperation, Social Stratification, Environmental Stress**

I. Introduction

The scale at which humans exhibit cooperative behaviour is unlike any other social mammal on the planet. Central to this behaviour lies the dichotomy of altruism and selfishness [1]. No time in ancient history demonstrates humanity's capacity for both selfish and altruistic acts more clearly than the transition from the Paleolithic to the Neolithic whereby egalitarian, hunter-gatherer, groups transitioned into sedentary agrarian societies exhibiting varying degrees of social stratification [2].

Despite archaeologists best efforts, the underlying dynamics that caused this transitory period are not entirely known. In fact, the mechanisms that led to the agricultural revolution are likely both multifaceted [3] and region-specific but, environmental stress is theorized to have played a significant role in the evolution of cooperative behaviour [4].

Agent-Based Models (ABMs) are often used to investigate emergent complex social phenomena and resource availability, as a function of environmental stress, on emergent cooperative-behaviour [5]–[7]. Also, ABMs are frequently used to study the emergence of social stratification (the grouping of people based on socioeconomic factors such as wealth, race and social status) in ancient societies [2], [8]. However, research combining these topics is scarce meaning the impact of environmental stress on cooperative-behaviour in socially stratified societies remains unknown. We address whether environmental stress (resource scarceness) positively impacts resource sharing (altruism) in socially stratified societies. We used an ABM called *NeoCOOP*, with artificial evolution as an adaptive mechanism to simulate emergent altruistic and selfish behaviour in Neolithic-inspired households. Our experiments examine agent resource trading preferences under varying degrees of environmental stress. This is supplemented by experiments using agent populations initialized to prefer selfish or altruistic behaviour.

Based on the findings of Ember et al.'s study of the resource sharing behaviours of societies in the Standard Cross-Cultural Sample [9], we hypothesize that environmental stress duration positively impacts agent resource trading preferences with longer periods of stress resulting in more altruistic behaviour in comparison to shorter, more frequent, periods of stress. We further hypothesize that clear evidence of social stratification will be present with agents exhibiting altruistic behaviour towards their peers and selfish behaviour towards their subordinates.

II. Related Work

ABMs typically implement cooperative behaviour in one of three ways (although hybrid solutions do exist):

1) Cooperative versus Defective: Agents are categorized as either purely selfish (defective) or purely altruistic (cooperative) and the emergent phenomena that arise from both homogenized and mixed agent populations are compared. These models are typically older and more exploratory [10]. Imitation or mimicking rules may also be added to these models to allow the agents to change their behaviour from cooperative to defective (or vice versa) over time [11].

2) Network-Based Cooperation: This is the modelling of agent-to-agent interaction and cooperation as a directed graph that acts as a form of social network [7], [8]. In order for two agents to interact directly, they must be connected within this network. ABMs implementing network-based cooperation are less common than the other methods of introducing cooperative behaviour with their existence heavily-reliant on the partitioning of agents along one or more metrics. Network-based solutions provide agents with the ability to specialize their behaviour more than other cooperation systems at the cost of removing an agent's ability to generalize.

978-1-6654-8768-9/22/$31.00 ©2022 IEEE

3) Probability-Based Cooperation: These agents are an extension to the cooperative or defective agents described above where the likelihood of agents exhibiting cooperative or defective behaviour is recorded as some probability p [5], [12]. These ABMs typically include some form of learning allowing agents to adapt their p value in accordance with a predefined set or rules or fitness-based algorithms such as Evolutionary Algorithms [13]. Probability-based cooperation ABMs are the 'middle-ground' approach between the highly generalized cooperate-defect systems and the highly specialized network-based systems.

ABM research directly related to ours includes Angourakis et al. [6] who studied the emergence of cooperative behaviour in scenarios with varying degrees of food storage efficiency, Pereda et al. [4] who studied the emergence of cooperation under varying degrees of environmental stress and Aktipis et al. [5] who compared need-based and account-keeping cooperation dynamics as they related to the Maasai of East Africa. More generally, Axelrod and Hamilton's [10] seminal work on the evolution of cooperation, Chliaoutakis and Chalkiadakis' [8] self-organizing agent hierarchies and Molin, Kanwal, and Stone's [7], study of emergent cooperation in spatially explicit environments are relevant to the study presented here.

III. Methodology

NeoCOOP (*Neolithic Agent Cooperation Model*) is an iterative ABM[1], that simulates evolving altruistic and selfish behaviour in a Neolithic inspired artificial society.

A. Agent Definition

In *NeoCOOP*, each agent represents a Neolithic *household*. The motivation for this is that typical Neolithic households were managed by a single patriarchal figure responsible for making all of the family's decisions as well as managing their resources [14]. Additionally, *NeoCOOP* uses *settlements* (Figure 1) to keep track of one or more *households*. A settlement's primary purpose is to store the coordinates of all the agents contained within that settlement.

Unlike most cooperation-based ABMs, *NeoCOOP* allows agents to make decisions based on their social status and the social status of the agents they are interacting with. We define social status as the sum of an agent's available resources and its *load*, where *load* is the amount of resources the agent has donated to other households over a period of time. To facilitate social stratification, we use the self-organization scheme described by Chliaoutakis and Chalkiadakis [8] whereby a relationship type can be determined for every agent pair by comparing their social status. We define each of the relationship types as follows:

[1]Source Code, ODD+D Description and supplementary material available at https://www.comses.net/codebase-release/f29e197a-81c4-4242-be3a-7300ba9e81e8/

$$is\_peer(h_1, h_2) = \frac{|h_2.status - h_1.status|}{max(h_1.status, h_2.status)} < L \quad (1)$$

$$is\_auth(h_1, h_2) = \frac{(h_2.status - h_1.status)}{max(h_1.status, h_2.status)} > L \quad (2)$$

$$is\_sub(h_1, h_2) = is\_auth(h_2, h_1) \quad (3)$$

Where is_peer, is_auth and is_sub describe whether household h_2 has a peer, authority or subordinate relationship with household h_1 respectively. $h_n.status$ is a household's social status. L, the *load_difference* $\in [0, 1]$ input parameter, defines how much more social status an agent requires to be considered an authority over another agent. In order for a peer, authority or subordinate relationship to be formed, the two households must be from the same settlement.

Lastly, the model facilitates Household adaptation in the form of two *Evolutionary Algorithms* (EA). The genotype used by both EAs comprises four gene values (all constrained within the $[0, 1]$ range):

1) **Peer Transfer:** Probability an agent accepts a resource transfer request from a peer agent.
2) **Subordinate Transfer:** Probability an agent accepts a resource transfer request from a subordinate agent.
3) **Conformity (σ):** The degree to which an agent accepts cultural influence.
4) **Attachment (α):** How much an agent values its current settlement. An agent with a high degree of attachment is less likely to migrate even if the environmental conditions suggest that it should.

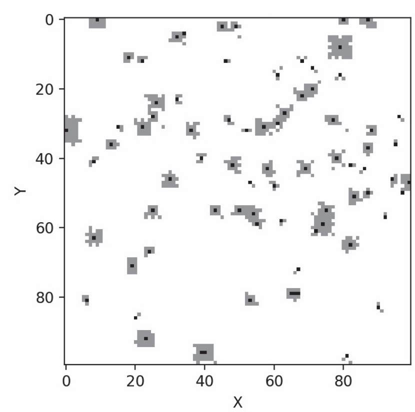

Fig. 1: Visualization of *NeoCOOP* at an arbitrary timestep. Black pixels are settlements, grey pixels are claimed resource patches and white pixels are unclaimed land.

B. Environment

NeoCOOP places agents on a $n \times n$ grid-world. Each cell on the grid contains resources $\in [0, 1]$ that are assigned to it every iteration. Stress is applied to these cells by varying the amount of collectable resources received each iteration according to *sine* waves of different frequencies. Denoting f as the desired number of stress waves, we linearly interpolate (Equation 5) every iteration i between two predefined ranges called $max\_resources = [0.4, 1.0]$ and $min\_resources = [0.0, 0.6]$ using the output of the sine waves (Equation 4) at timestep t/M as the mixing parameter x. This approach blends work by Molin, Kanwal and Stone [7] and Angourakis et al. [6] where environmental stress is induced periodically and between two predefined ranges respectively. This approach allows us to simulate a wide variety of stress scenarios ranging from short, but frequent, periods of stress (at high f) to longer, infrequent, periods of stress (at low f). Averaged over an entire simulation run, a household is expected to receive a total of 0.5 resources per iteration. An example of what the result of this process looks like can be seen in Figure 2.

$$s(x) = 0.5 \times sin(2\pi \times x \times f) + 0.5 \tag{4}$$

$$lerp(r_{min}, r_{max}, x) = r_{min} + s(x) * (r_{max} - r_{min}) \tag{5}$$

$$resources(x) = random(lerp(0.0, 0.6, x), lerp(0.4, 1.0, x)) \tag{6}$$

The motivation for choosing a spatially explicit environment is because even ideal environments have a carrying capacity. Most spatially implicit ABMs do not consider population carrying capacity which limits their capabilities of accessing cooperative behaviour dynamics between two distinct population groups (those with and without direct access to resources).

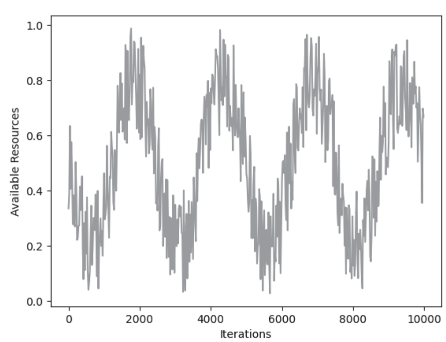

Fig. 2: Example of the available resources (Equation 6) on a single environment cell over the course of the simulation run when $f = 4$.

C. Resource Acquisition, Transfer and Consumption

Every iteration, agents gather resources from a patch of land that they own (Grey pixels in Figure 1). The amount of resources gathered is equal to the full amount of resources available at said patch. These resources are then put into the agent's *storage*. In this work, agents are only allowed to own one patch of land. If an agent does not own any land, it will try to claim some by looking at its settlement's neighbouring cells. An agent that does not own any land will not receive any resources during the resource acquisition phase.

Once acquisition is complete, agents determine if they have enough resources to satisfy their needs for the iteration. An agent needs to consume 0.5 resources per iteration to avoid the risk of dying. If an agent does not have enough resources, it first asks its authority agents if they would be willing to donate some of their excess resources. For each authority asked, a random value $\in [0, 1]$ is generated and compared to the authority agent's *subordinate_transfer* property. If the generated value is less than the *subordinate_transfer* property, the authority agent is willing to grant donations for that iteration. Whenever a donation is granted, the authority agent has its *load* property increased by the resources donated. If an agent has asked all of its authority agents for resources and it will still go hungry, it then repeats this process for its peer relationships with the donating agent using its *peer_transfer* property to determine if the donation succeeds.

If that is still not sufficient, the agent will then ask all of its subordinates for resources. Given that we are modelling Neolithic households, if a subordinate is asked to give any of its excess resources to an authority agent, it does so with 100% certainty. The peer and subordinate transfer properties allow us to simulate agent types that exhibit varying degrees of altruistic and selfish behaviour. For example, an agent may exhibit nepotistic tendencies whereby it is more likely to grant resource donations to its peers (high *peer_transfer*) but less likely to grant the same donations to its subordinates (low *subordinate_transfer*).

When resource transfer is complete, agents consume their resources and determine their *hunger* using Equation 7.

$$hunger(h) = min(\frac{h.resources}{0.5}, 1.0) \tag{7}$$

D. Population Growth, Loss and Migration

Every iteration, households may birth additional households in accordance with the *birth_rate* and their *hunger* (Equation 8). When this occurs, the household is divided into two separate households and resources are split amongst the two new households but load is not. That is, the new household signifies the arrival of a new patriarchal figure in the community and one who must work to gain the same social status as their parent household.

$$birth(h) = random(0.0, 1.0) < h.hunger * birth\_rate \tag{8}$$

Households may lose one or more occupants in accordance with the *death_rate* and their *hunger* (Equation 9). If a household dies of starvation, it is removed from the simulation.

$$death(h) = random(0.0, 1.0) * h.hunger < death\_rate \tag{9}$$

Agents can migrate to another settlement or form a settlement of their own every *yrs_per_move* iterations. This decision is based on the agent's *satisfaction* and its *attachment*. *Satisfaction* is the average *hunger* of the agent over the past *yrs_per_move* iterations. The boolean function for determining if an agent will move is described by Equation 10.

$$\textit{move(h)} = 2\alpha_h * \textit{satisfaction(h)} < \textit{random(0.0, 1.0)} \quad (10)$$

Where α is the attachment of household h. If the *satisfaction* of the agent is low, it is more likely to move. This is partly mediated by the agent's *attachment* which when < 0.5, makes the agent skittish and when > 0.5 makes the agent more likely to stay at a given location regardless of its objective circumstances. In population migration research, the inverse of *satisfaction* is often called *grievance* [15].

When an agent moves, it chooses between all settlements in its vicinity or an unclaimed cell. Typically, an agent will move to the settlement with the most resources. However, if none of the neighbouring settlements have an average resource value ≥ 0.5, the agent will choose to make its own settlement at a new, randomly chosen, location.

E. Agent Adaptation

In this model, agent adaptation uses two Evolutionary Algorithms: a *Genetic Algorithm* (GA) [16] for vertical generational adaptation and a *Cultural Algorithm* (CA) [17] for horizontal generational adaptation. Both the GA and CA utilize the agent genotype described before and a concept called *influence*. *Influence* is used to determine best performing settlements and describes the probability that two settlements will interact with each other. This is done using XTENT [8] (Equation 11):

$$I(s_1, s_2) = (s_2.status)^{\beta} - mD(s_1, s_2) \quad (11)$$

Where, s_1 and s_2 are settlements, $I(s_1, s_2)$ is the influence of s_2 on s_1, $s_2.status$ is the social status of s_2 (the total social status of all households in s_2), $D(s_1, s_1)$ is distance from s_1 to s_2. β and m are coefficients describing the required social status of one settlement to influence another. Calculating the *influence* of every settlement on a given settlement, gives a probability distribution (Equation 12).

$$P(s_1, s_2) = \frac{I(s_1, s_2)}{\sum_{k \in K} I(s_1, s_k)} \quad (12)$$

Where $P(s_1, s_2)$ is the probability of settlement s_2 influencing settlement s_1 and K is the set of neighbouring settlements that have a positive *influence* value $I(s_1, s_k)$ on s_1.

The GA executes whenever another household is created. The child agent produced is a combination of the household that split and a second household gotten via *roulette wheel* selection [16]. This selection uses the social status of other agents within the same settlement of the first parent and from other settlements that have enough influence ($I(s_1, s_2) > 0$). The offspring agent is produced using *Uniform crossover* and random mutation.

The CA uses *Knowledge Sources* [18] to diversify how agents are influenced. These are:

- **Normative:** Influence on agent genes: its settlement.
- **Spatial:** Influence on agent genes: another settlement.
- **Domain:** Equivalent to GA mutation function, where domain influence mutates one of the agent's genes.

Every *influence_frequency* iterations, agents are influenced in accordance with the *influence_rate*. If an agent is selected for influencing, a roulette wheel is spun to determine from which knowledge source influence will come from. Influence from the Domain knowledge source occurs at a rate defined by the *mutation_rate* parameter. Influence from the Normative and Spatial knowledge sources occur with varying probability defined by Equations 13 and 14.

$$N(s_h, s_i) = max(\frac{s_h.status}{s_i.status}, 1.0) \quad (13)$$

$$S(s_h, s_i) = 1.0 - N(s_h, s_i) \quad (14)$$

Where, N and S are the probability of choosing the normative and spatial knowledge sources respectively, s_h is the settlement of the agent being influenced, s_i is the settlement that would influence agent h. If the spatial knowledge source is selected, s_i is determined by performing roulette wheel selection on all neighbouring settlements with a positive *influence* on settlement s_h. Roulette wheel weights are determined by the values returned by Equation 12.

Each settlement's beliefs are represented by *Belief Spaces* B_s. Belief Spaces have the same structure as the agent genotype with each property calculated using a weighted average of the corresponding property of all agents within that settlement. The weight an agent contributes to the belief space is determined using its social status relative to the social status of the other agents in the same settlement. If an agent is influenced by the normative knowledge source, the belief space that influences it is the belief space of the settlement the agent belongs to B_{s_h}. If the agent is influenced by the *spatial* knowledge source, the belief space that will influence the agent is the belief space of the settlement selected during roulette wheel selection (B_{s_i}). Agent properties are influenced as follows (Equation 15):

$$G_{h,t+1}(p) = G_{h,t}(p) + \sigma_h(B_{s,t}(p) - G_{h,t}(p)) \times \Phi(h, B_{s,t}) \quad (15)$$

Where, p is the agent property (genes 1-4), t is the timestep, G is the agent's genotype, σ_h is the *conformity* of the agent, B is the selected belief space (B_{s_h} or B_{s_i}) and Φ is the Homophily term which returns a value $\in [0, 1]$ describing how similar the agent's genes are to the belief space that is influencing it. Homophily is a sociological principle that describes the tendency for individuals that are similar, either biologically or culturally, to gather together. The value of Phi is 1.0 for interacting entities that have exactly the same genes, and close to 0.0 for entities whose gene values are further apart. This approach is similar to interaction probability in Axelrod's cultural dissemination model [19]. In our model, Φ limits the degree to which an agent is influenced if the belief space influencing it contains drastically different gene values. Formally, Φ is one minus the average absolute difference between the agent and influencing belief space's genes.

IV. Experiment Design

Before running our experiments, we parameter tuned our model (See Table I) using the same process described in Gower-Winter and Nitschke [20]. A report of the tuning process is included with the source code[1]. Given our goal to find the environmental conditions under which social stratification occur, we ran our experiments as follows.

We first defined initial resource trading belief distributions for the agent types (denoted $[A, S, F]$). For purely altruistic A initialization, agents have their peer and sub transfer properties initialized to 1.0. For purely selfish S initialization, agent peer and sub transfer properties are set to 0.0 and the mixed population F scheme initializes the agents' resources trading beliefs such that half of them follow the A initialization scheme and the other half follow the S initialization scheme. We use differing initialization schemes since the initial resource trading beliefs of an agent population may affect how they evolve over time. Across all agent-types, attachment is initialized $\in [0.0, 1.0]$ for each agent. Unless stated otherwise, all random numbers were drawn from a uniform distribution.

When then defined the *stress scenarios* investigated as follows: $f \in [1, 2, 4, 8, 16, 32, 64, 128]$. We also explore two scenarios where resource availability is confined to the range $[0.4, 1.0]$ and $[0.0, 0.6]$ for the entire simulation run (denoted as the non-existent (N) and perpetual (P) stress scenarios respectively). The motivation for choosing the aforementioned frequencies is based on preliminary experiments where it was observed that selfish behaviour could emerge at low f-values. We then expanded the scope of the experiments to include higher f-values to see if this trend persisted.

Using the three initialization schemes and 10 stress conditions, 30 scenarios where created. For each scenario, 50 simulations were run for a total of 1500 simulations across all scenarios. Each simulation was initialized with 100 agents and 10 settlements. At initialization, each agent in the model had their *peer_transfer* and *sub_transfer* agent properties set to either 1.0 or 0.0 depending on the initialization scheme (A scenario denoted as $16A$ indicates that the A initialization scheme was used with an f-value of 16). Settlements were randomly placed on the grid-world and the model was run for $M = 10000$ iterations. All stochastic processes utilized a pseudo-random number generator to ensure reproducibility.

TABLE I: NeoCOOP Initialization Parameters.

Property	Value
Iterations (M)	10 000
Initial Households	100
Initial Settlements	10
L	0.6^a
Years Per Move	5^a
Birth Rate	$0.15\%^b$
Death Rate	$0.1\%^b$
β	1.5^a
m	0.005^a
Mutation Rate	0.1^c
Influence Rate	0.1^c
Influence Frequency	15^c
Conformity Range	$\in [0.2, 0.7]^c$

[a]Properties taken from Chliaoutakis and Chalkiadakis [8].
[b]Properties taken from Cardona, Català, and Prats [21] .
[c]Properties that were parameter tuned using *Optuna*.

V. Results

Figures 3, 4 and 5 showcase the evolution of the resource transfer genes for the A, S and F initialization schemes across all stress scenarios. A clear visual distinction between the evolution of the peer and subordinate transfer genes is present. To confirm this, a Wilcoxon rank-sum test ($p = 0.05$) was performed to see if there was a significant difference between the agents' peer and subordinate resource transfer beliefs. As shown in Table II, all scenarios with $f \geq 8$ exhibited statistically significant stratification between the peer and subordinate transfer properties. The magnitude of this difference ranged from 1.32% to 3.74%. However, the trend this difference followed was not consistent across all initialization schemes. Scheme S consistently exhibited differences $\geq 2.65\%$ while schemes A and F exhibited lower differences for $f = [64, 128]$ stress frequencies compared to $f = [8, 16, 32]$ suggesting there may be an optimal range for which greater degrees of social stratification occur.

TABLE II: Summary of the mean difference of the agents peer and subordinate resource transfer beliefs.

(*f*)	Initialization schemes					
	A		F		S	
	diff (%)	p	diff (%)	p	diff (%)	p
P	0.0	N/A	0.0	N/A	0.0	N/A
1	-0.71	0.69	-3.57	0.82	1.21	0.25
2	0.92	0.06	-0.02	0.58	0.22	0.44
4	0.39	0.26	1.13	0.19	1.82	0.06
8	2.30	$\mathbf{3 \times 10^{-5}}$	2.11	$\mathbf{2 \times 10^{-4}}$	2.80	$\mathbf{2 \times 10^{-6}}$
16	3.74	$\mathbf{1 \times 10^{-12}}$	3.33	$\mathbf{1 \times 10^{-8}}$	3.31	$\mathbf{6 \times 10^{-12}}$
32	2.69	$\mathbf{4 \times 10^{-7}}$	2.51	$\mathbf{5 \times 10^{-6}}$	3.45	$\mathbf{3 \times 10^{-11}}$
64	1.39	**0.041**	1.41	**0.011**	2.65	$\mathbf{5 \times 10^{-5}}$
128	1.32	**0.003**	1.95	**0.0007**	3.00	$\mathbf{8 \times 10^{-8}}$
N	0.92	**0.007**	1.23	**0.004**	0.77	0.06

A value $p \leq 0.05$ (bolded) indicates a Wilcoxon rank-sum test determined the stratification of the resource transfer beliefs was significant.

To investigate this observation further, supplementary experiments were performed using the F initialization scheme for $f = [24, 40, 48, 56]$ under the same conditions highlighted in our Experiment Design. Figure 6 showcases the results of these experiments. Again, significant ($p = 0.05$)

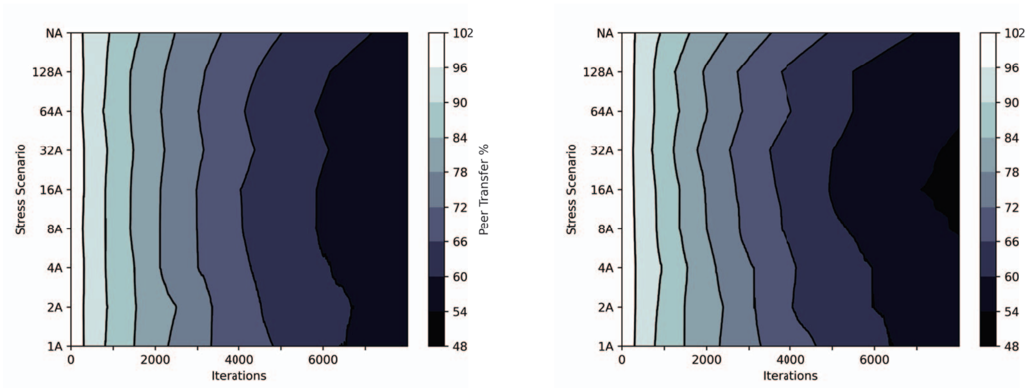

(a) (b)

Fig. 3: The average evolution of the *peer* (a) and *subordinate* (b) transfer agent properties over the course a simulation for all *stress scenarios* investigated using the A-type initialization scheme.

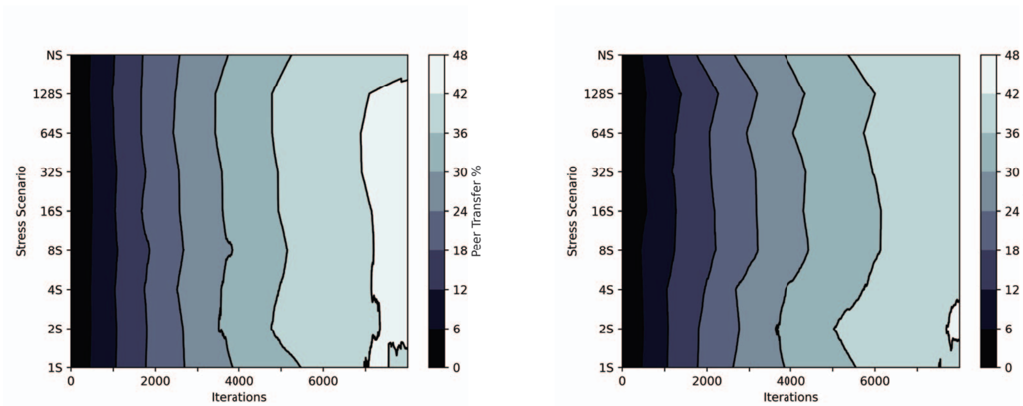

(a) (b)

Fig. 4: The average evolution of the *peer* (a) and *subordinate* (b) transfer agent properties over the course a simulation for all *stress scenarios* investigated using the S-type initialization scheme.

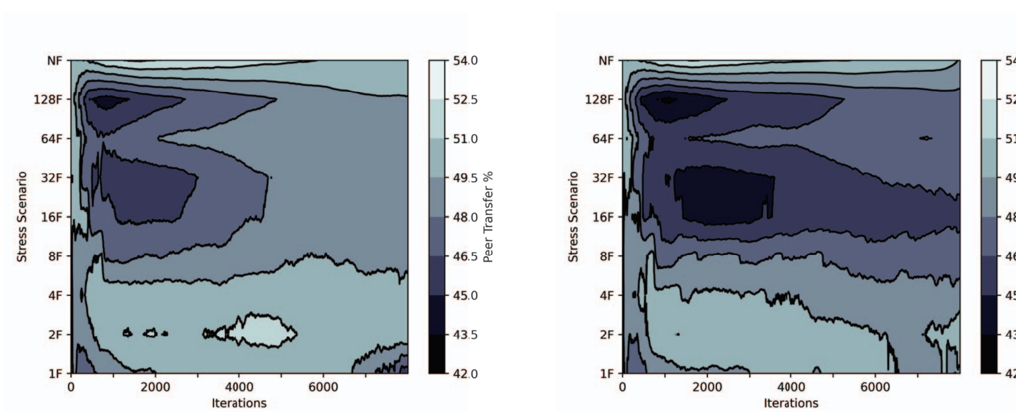

(a) (b)

Fig. 5: The average evolution of the *peer* (a) and *subordinate* (b) transfer agent properties over the course a simulation for all *stress scenarios* investigated using the F-type initialization scheme.

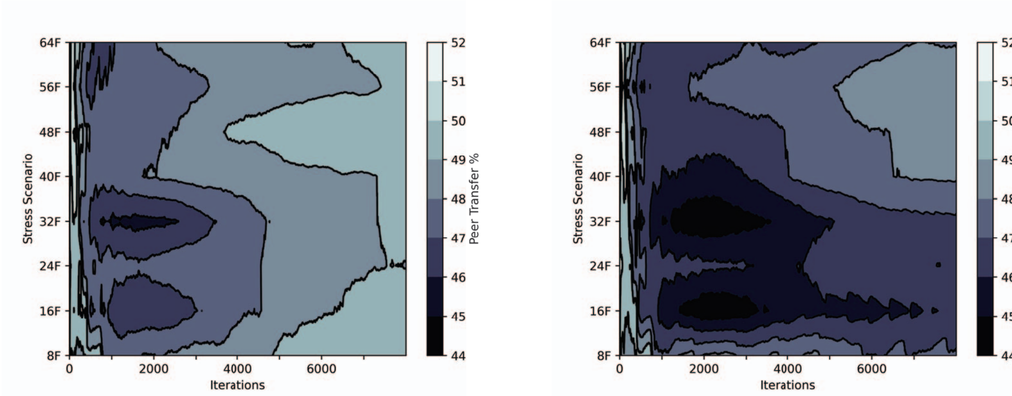

(a) (b)

Fig. 6: The average evolution of the *peer* (a) and *subordinate* (b) transfer agent properties over the course a simulation for all supplementary *stress scenarios* investigated using the F-type initialization scheme.

stratification between the peer and subordinate transfer beliefs was found for all stress scenarios. Additionally, results indicate that the magnitude of the stratification is not uniform with greater values ($> 2.5\%$) found between $8 < f < 40$ compared to the $< 2\%$ difference found for $40 \leq f \leq 64$. This supports the claim that there are optimal conditions under which stratification of resource trading preferences may occur.

For the perpetual stress scenario P, agent populations died out within 2000 iterations across all initialization schemes. Whereas, for the no stress scenario N, agent populations, on average, reached a carrying capacity of approximately 12000 households by iteration 9750. From iteration 9750 and on, stratification began to occur whereby the difference between the peer and sub transfer genes increase by between $[0.77, 1.23]\%$. This indicates that while stratification occurs in harsher environments, it can also occur in environments with no environmental stress.

In terms of the evolution of cooperative behaviour, both the peer and subordinate transfer properties, across all scenarios investigated, converged to values between the $[0.4, 0.6]$ range. Agents evolved away from both extreme altruistic and selfish behaviours with the A and S initialized agents becoming more selfish and altruistic respectively. This 'middle-ground' phenomena has been documented previously [6] in scenarios where cooperative food storage is low and household storage efficiency is high (Both of which are applicable here). Interesting behaviour was observed for the F scheme with altruistic behaviour preferred for lower frequency stress scenarios ($f = [1F, 2F, 4F, 8F]$) and selfish behaviour preferred in higher frequency stress scenarios (all other F scenarios). This is most apparent (See Figure 5) earlier on in a simulation run where agents in higher frequency stress scenarios rapidly evolve, comparatively, selfish behaviour while agents in lower frequency stress scenarios evolve altruistic behaviour. These results indicate that for harsher (low frequency) periods of environment stress, cooperative behaviour is preferred.

VI. Discussion

In this paper we sought to answer whether environmental stress positively impacted resource sharing in socially stratified societies. We hypothesized that the duration of environmental stress impacts agent resource trading preferences with longer periods of stress resulting in more altruistic behaviour in comparison to shorter, more frequent, periods of stress. We further hypothesized that clear evidence of social stratification would be present with agents exhibiting altruistic behaviour towards their peers and selfish behaviour towards their subordinates. For the most part, our results indicate these hypotheses were correct. Low frequency stress scenarios exhibited higher peer and subordinate transfer values and clear stratification between these two resource trading beliefs were found at higher frequencies. However, there is nuance to these statements which require further discussion.

Our results suggest that in extreme conditions, altruism is favoured. When stress is frequent, selfishness is favoured. This makes sense in the context of the model as at both extremes, resource sharing has a clear benefit. In harsh environments, sharing resources acts as a safety mechanism. Helping a household now means they might reciprocate in the future. In environments with no stress, there is no harm in sharing resources because there's no risk you may not have access to more resources in the future. These results are also supported in literature where it has been argued that resource sharing may be reduced after a stress period and that frequent stress may cause selfish behaviour to emerge [9].

Furthermore, our results not only show the emergence of social stratification by way of differing resource trading preferences amongst peers and subordinates, they also suggest that there is an optimal range under which this stratification can occur. This is clearly demonstrated in Figure 6 where it can be seen that the magnitude of this stratification decreases for $f > 40$. While our experiments were not explicitly designed to test for such a relationship, we believe that the degree of stratification exhibited by a society is related to its ability to maintain logistic growth while still being affected by environmental stress (See Supplementary Materials[1] for additional figures supporting this claim). From a historical perspective, some neolithic societies maintained a two-stage logistic-like growth curve [22] despite environmental stress, and in the case of *Tell Halula* in the Middle Euphrates, there is a clear emergence of social classes. This is a region which received frequent environmental stress and, while only indirectly attributed in literature, the presence of both a social class and frequent environmental stress further supports our theory [23].

Lastly, we highlight the social behaviour dynamics that emerge when a population reaches non-environmental stress related carrying capacity. For the N stress scenarios, emergent stratification was observed across all initialization schema once the populations had reached carrying capacity. This type of cultural evolution is distinct from the cultural evolution investigated in this paper because it is not brought about by the application of stress on the collective of agents but rather by the partitioning of the agents into two groups: those with and without access to resource patches. Our experiments did not run at carrying capacity for a long enough period of time for us to make any postulations but, there is opportunity to study this further in future work.

VII. Conclusions and Future Work

We investigated the emergence and evolution of cooperative behaviour in artificial Neolithic societies exposed to varying degrees of environmental stress. We achieved this by using the NeoCOOP ABM, which utilizes EAs as adaptive mechanisms to facilitate emergent social behaviour including resource trading beliefs in households (agents). Results indicate that in extreme scenarios, altruism is favoured with selfish behaviour favoured as the frequency of environmental stress increases. Furthermore, results indicate that the magnitude of social stratification is related to the agent population's capacity to maintain logistic-like growth while remaining susceptible to environmental stress. If a population collapses or completely withstands an environmental stress event, the magnitude of stratification is likely to be lower.

Future work will directly deal with the limitations of this paper in the hopes of affording us a greater understanding of the complex dynamics that gave rise to the Neolithic agricultural revolution. These efforts will focus on the implementation of food storage efficiency [6] and studying the emergence of non-environmental stress related social stratification such as direct access to resources through land availability.

References

[1] H. Rachlin, "Altruism and selfishness," *Behavioral and brain sciences*, vol. 25, no. 2, pp. 239–250, 2002.

[2] S. Powers and L. Lehmann, "An evolutionary model explaining the neolithic transition from egalitarianism to leadership and despotism," *Proceedings of the Royal Society B: Biological Sciences*, vol. 281, no. 1791, p. 20141349, 2014.

[3] M. Stiner, "Thirty years on the "broad spectrum revolution" and paleolithic demography," *Proceedings of the National Academy of Sciences*, vol. 98, no. 13, pp. 6993–6996, 2001.

[4] M. Pereda, D. Zurro, J. Santos, I. Godino, M. Álvarez, J. Caro, and J. Galán, "Emergence and evolution of cooperation under resource pressure," *Scientific reports*, vol. 7, no. 1, pp. 1–10, 2017.

[5] A. Aktipis, R. De Aguiar, A. Flaherty, P. Iyer, D. Sonkoi, and L. Cronk, "Cooperation in an uncertain world: For the maasai of east africa, need-based transfers outperform account-keeping in volatile environments," *Human Ecology*, vol. 44, no. 3, pp. 353–364, 2016.

[6] A. Angourakis, J. Santos, J. Galán, and A. Balbo, "Food for all: An agent-based model to explore the emergence and implications of cooperation for food storage," *Environmental Archaeology*, vol. 20, no. 4, pp. 349–363, 2015.

[7] L. Molin, J. Kanwal, and C. Stone, "Resource availability and evolution of cooperation in a 3D agent-based simulation," in *Proceedings of the Genetic and Evolutionary Computation Conference*, pp. 93–101, 2021.

[8] A. Chliaoutakis and G. Chalkiadakis, "An agent-based model for simulating inter-settlement trade in past societies," *Journal of Artificial Societies and Social Simulation*, vol. 23, no. 3, 2020.

[9] C. Ember, I. Skoggard, E. Ringen, and M. Farrer, "Our better nature: Does resource stress predict beyond-household sharing?," *Evolution and Human Behavior*, vol. 39, no. 4, pp. 380–391, 2018.

[10] R. Axelrod and W. Hamilton, "The evolution of cooperation," *science*, vol. 211, no. 4489, pp. 1390–1396, 1981.

[11] C. Power, "A spatial agent-based model of n-person prisoner's dilemma cooperation in a socio-geographic community," *Journal of Artificial Societies and Social Simulation*, vol. 12, no. 1, p. 8, 2009.

[12] T. Nhim, A. Richter, and X. Zhu, "The resilience of social norms of cooperation under resource scarcity and inequality—an agent-based model on sharing water over two harvesting seasons," *Ecological Complexity*, vol. 40, p. 100709, 2019.

[13] P. Revay and C. Cioffi-Revilla, "Survey of evolutionary computation methods in social agent-based modeling studies," *Journal of Computational Social Science*, vol. 1, no. 1, pp. 115–146, 2018.

[14] M. Lehner, "Fractal house of pharaoh: Ancient egypt as a complex adaptive system, a trial formulation," *Dynamics in Human and Primate Societies. Agent-based Modeling of Social and Spatial Processes, Oxford*, pp. 275–353, 2000.

[15] A. Srbljinovic, D. Penzar, P. Rodik, K. Kardov, *et al.*, "An agent-based model of ethnic mobilisation," *Journal of Artificial Societies and Social Simulation*, vol. 6, no. 1, p. 1, 2003.

[16] A. Eiben and J. Smith, *Evolutionary Computing*. Berlin, Heidelberg: Springer, 2nd ed., 2015.

[17] R. Reynolds, "An introduction to cultural algorithms," in *Proceedings of the third annual conference on evolutionary programming*, vol. 24, pp. 131–139, World Scientific, 1994.

[18] R. Reynolds and B. Peng, "Cultural algorithms: modeling of how cultures learn to solve problems," in *16th IEEE International Conference on Tools with Artificial Intelligence*, pp. 166–172, IEEE, 2004.

[19] R. Axelrod, "The dissemination of culture: A model with local convergence and global polarization," *Journal of conflict resolution*, vol. 41, no. 2, pp. 203–226, 1997.

[20] B. Gower-Winter and G. Nitschke, "Societies prefer the middle-ground between selfishness and cooperation," in *4th International Workshop on Agent-Based Modelling of Human Behaviour (ABMHuB'22)*, 2022.

[21] P.-J. Cardona, M. Català, and C. Prats, "The origin and maintenance of tuberculosis is explained by the induction of smear-negative disease in the paleolithic," *Pathogens*, vol. 11, no. 3, p. 366, 2022.

[22] M. Bandy, "New world settlement evidence for a two-stage neolithic demographic transition," *Current Anthropology*, vol. 46, no. 5, pp. 109–115, 2005.

[23] J. P. Ferrio, G. Arab, R. Buxó, E. Guerrero, M. Molist, J. Voltas, and J. Araus, "Agricultural expansion and settlement economy in tell halula (mid-euphrates valley): A diachronic study from early neolithic to present," *Journal of Arid Environments*, vol. 86, pp. 104–112, 2012.