JASSS
THE JOURNAL OF ARTIFICIAL SOCIETIES AND SOCIAL SIMULATION

ODD Description: MoundSim LandUse

Below we provide an ODD (Overview + Design + Detail) description of MoundSim LandUse (Grimm et al. 2006, 2010). This description is amended with the recommendations of Müller et al. (2013) to better account for features of the model that are associated with human decision-making.

Purpose

MoundSim LandUse is an agent-based model that aims to explore the nature of landscape modification for maize-based agriculture by the pre-Columbian Casarabe Culture in the Monumental Mound Region (MMR) of the Llanos de Moxos (LM), located in Amazonian Bolivia. The model serves as a 'virtual laboratory' (Magliocca & Ellis 2016), enabling the user to examine how population size, resource demands, and agricultural behavioural practices influence the extent and spatial distribution of modified patches of land in a variety of different 'What if?' scenarios. By comparing these outputs to empirical data as it becomes available, the model aims to consolidate our understanding of the mechanisms and motivations underlying human-environment interactions in the MMR. MoundSim LandUse is implemented in NetLogo Version 6.3.0 (Wilensky 1999).

Entities, State Variables, and Scales:

MoundSim LandUse possesses four agent types: Square land patches, pre-Columbian household units, settlements, and community ties. The model uses imported, georeferenced (GIS) data to recreate a 5020 km^2 of the MMR, including the 4500km^2 region studied in Lombardo and Prümers (2010). This is split into four quadrants of equal size, with the model able to display one quadrant at once. The simulated landscape consists of a 359x416 grid of patches, each representing one hectare of land. Each patch possesses immutable environmental variables determined by imported environmental data (e.g., Elevation, Productivity), as well as variables to determine whether the Casarabe culture's earthworks have been found there. They also possess a variable to determine its state of land cover (land-use), which is initially determined by imported data, but can change over time.

The primary agent of interest within the model represents a pre-Columbian household, reflecting a nuclear family unit of two adults and three children. Households are immobile, assumed to be permanently situated atop a mound settlement agent, to which they are attached by a community link. While the settlement agent influences household behaviour, community link agents possess no variables of their own and simply exist to connect these two agent types. These households are also characterised by a set of variables that describe their demand for a variety of resources considered necessary by both the pre-Columbian inhabitants of the MMR (the 'Casarabe culture') and members of contemporary Indigenous Amazonian communities (Maize, Foraged Tree crops, Fuelwood, Palm Leaves, and Animal protein).

To acquire the resources necessary for their survival, households possess behaviour to claim, retain, modify, and extract resources from terrestrial patches of land. These processes take place within a user-defined distance from the attached settlement agent, stored in three different variables (settlement-radius, forage-radius, fishing-radius). When land is claimed, households possess a variable to record and track the identity of the patch (household-territory), with a similar variable used by each patch to track its owner household (owner).

If an insufficient quantity of each resource is produced by the patches within a household's territory, that household can utilise two alternative methods to obtain them. Firstly, they can obtain resources from communal stores supplied by other households that are attached to the same settlement. Secondly, they can extract resources from unclaimed patches within a set radius that are controlled by their settlement (stored in the settlement-owner patch variable). If the household fails to produce a sufficient quantity of resources, the modeller can decide whether it either becomes stressed (discussed below), or dies.

Global parameters are initially used to store imported environmental information (elevation, productivity, land-use) earthwork (canal? causeway? mound?) data before it is mapped to a number of patch variables. Similar variables also track the starting household population (starting-pop), as well as the number of households that die as a result of resource shortages. Each timestep in the model equates to one year of activity, with simulations running for one thousand timesteps to match current radiocarbon data that indicates the mounds within the central MMR were continuously occupied between 400-1400 CE (Jaimes Betancourt 2012, 2015; Prümers 2015:82-87).

Process, Overview and Scheduling

MoundSim LandUse simulates multiple key processes. Upon initialisation, settlement agents spawn in locations that match where the Casarabe culture's mounds have been identified on the real landscape (Lombardo & Prümers 2010). Each mound then generates a user-defined number of households, which are subsequently attached to it. The model possesses no mechanism to increase the number of households during each simulation, given that it also aims to examine the effects of regional population size on the extent and spatial distribution of altered land. The number of households can, however, decrease during a simulation based on user preference.

The first procedure executed during each timestep represents the conversion of land. During this procedure, households make decisions around whether and how much additional land to convert for maize-based agriculture. If a household deems its supply insufficient for the following year, it will claim and convert additional land patches. Terrestrial land from anywhere on the landscape can be claimed for resource production, as long as it lies within an appropriate distance range. However, households can only claim and convert a limited number of patches during a single timestep, reflecting limitations on time and labour. This procedure updates the landscape asynchronously to reflect the competition for finite resources, as well as to ensure a patch cannot be owned by more than one agent.

The second procedure reflects the production of resources. Patches possess a number of variables to store the supplies of various resources. This procedure allocates a suitable quantity of each resource type to patches within the territory of household agents.

The third procedure, decrease population, determines whether household agents have produced sufficient quantities of each resource to avoid becoming stressed, as well as which (if any) households disband. For simplicity, we assume that any households that disband due to age-related effects are immediately replaced by new households, maintaining a fixed population. However, a household may also disband due to a lack of resources, assumed to represent the death of its inhabitants. Decrease population manages the alternative methods of acquiring

JASSS
THE JOURNAL OF ARTIFICIAL SOCIETIES AND SOCIAL SIMULATION

resources should a household's territory prove insufficient. If the quantity of resources remains insufficient, it determines whether households become stressed, or disband.

The final procedure, vegetation-growth, accounts for the abandonment of cropland and the regeneration of fallowed land into forest and savanna. Initially, it determines whether a patch is abandoned based on the number of years it has remained in cultivation. Subsequently, the procedure stochastically governs the transition of fallowed land into regenerated savanna or forest based on whether it was forested at the start of the simulation.

Design Concepts

Theoretical and Empirical Background

MoundSim LandUse has been built with the aim of exploring a range of 'What If?' scenarios associated with the extent and spatial distribution of pre-Columbian land-use in the MMR. The paucity of available data makes this aim heuristic by nature; it seeks to explore a range of plausible scenarios rather than attempting to approximate reality (Lake 2014). The mechanisms within the model take inspiration from Malthusian theory (1798) in that, provided aggregated household demand never exceeds the total quantity of locally accessible resource supplies, reproduction can proceed unhindered. However, total resource supplies are finite and, when insufficient, households may become stressed or die.

These concepts are encoded as rules governing household agent behaviour. Rather than being able to claim, convert, and extract resources from anywhere on the landscape, households are restricted in terms of accessibility. They can only procure resources within a set radius of their attached settlement, can only convert a set number of patches per timestep, and may not select the optimal location for their farms.

Individual Decision-making

Decision-making within the model is conducted at the household level. Households are satisficers that utilise inductive reasoning to determine which and how many patches to convert (Arthur 1994). This is demonstrated when converting land, as households claim additional patches based on whether the expected supply of resources from their existing territory is sufficient for the forthcoming timestep. They have further preferences for certain environmental characteristics defined by the user, which influences the patches selected for cultivation. However, this comes with restrictions. They can only convert a set amount of land during each timestep, and can only choose between a subset of the available patches due to imperfect information.

Due to the paucity of available data, few socio-cultural values are considered beyond the assumption of maize as the staple crop (Bruno 2010; Dickau et al. 2012; Prümers 2015; Whitney et al. 2013), and households being sedentary, located atop habitation mounds (Lombardo & Prümers 2010; Prümers et al. 2022).

Learning

Agents do not learn as part of the decision-making process.

Individual Sensing

Households employ various sensing mechanisms to inform their decisions. When cultivating patches, they can discern the essential characteristics of each patch, including its elevation, and its land cover characteristics. They will only attempt to convert unclaimed land, in keeping with the practices of contemporary Indigenous groups. The model, however, does not account for territorial disputes.

When choosing where to cultivate, households are assumed to possess the ability to recognise fallow land, determining when it became fallow by recognising distinctive vegetation characteristics (Huanca 1999). This is encoded in the model by allowing agents to sense the land-use variable, with agents choosing only to farm areas of regenerated forest/savanna even if abandoned fallow land is available. This assumption is a practical, if simplified, representation of the real world, as humans may recognise fallowed land but not necessarily its owner.

Individual Prediction

Prediction plays a crucial role for households in deciding whether to cultivate additional land. This is based on whether households expect to obtain sufficient resources from patches within their territory to meet their needs for the following year. Their decision relies upon imperfect information, utilising average resource yields derived from productive patches of cropland, agroforestry, and fallow. It is assumed that the pre-Columbian inhabitants of the MMR were able to assess average resource yields and anticipate expected sources of yield reduction such as crop loss. However, they cannot determine the level of unexpected variability for the current timestep.

Households within the model assume that a single hectare of land can produce an expected quantity of maize. Following patterns observed within the Tsimane population, we assume the MMR's pre-Columbian population consumed this maize throughout the year (Ringhofer 2010). Acknowledging limitations in storage options and pest control, the model assumes that any surplus from cultivation cannot be stored interannually. Additionally, households proactively make cultivation decisions, considering a predetermined 15% loss in maize yield as a pre-emptive measure (Ringhofer 2010).

Interaction

Households within the same settlement collective engage in direct and indirect interactions that are facilitated by the community tie link. Households may only share resources between members of the same community. All households are assumed to reside on a mound settlement, cultivating within the range determined by the settlement agent. This fosters indirect competition for desirable land patches.

Collectives

Each household is compelled to establish a community link to a pre-Columbian settlement collective agent. This ensures every household belongs to a collective, even if that collective contains only the household itself. These collectives play a pivotal role by influencing the ability of member households to convert patches on the landscape. Patches can only be converted within a given range of a settlement. A collective agent can also exist without attached household members (we treat this as an abandoned settlement). Settlement collectives influence the sharing of resources, and dictate the radius within which agents can extract resources from unclaimed patches.

Heterogeneity

Agents are not heterogeneous with regards to the decision-making process

Stochasticity

Stochasticity is involved in selecting patches for cultivation. Only a random subset of the available patches within range are chosen. The size of this subset is defined by the user, enabling variations in agent knowledge to be tested. The user is also provided a NULL model choice, which selects patches within range entirely at random.

Stochasticity is involved in resource production and vegetation growth. In the former, stochasticity acts to vary production around a mean value defined by constants within the model. In the latter, it acts to vary the rate at which certain patches become reforested/regenerated after abandonment.

Observation

During each timestep the view updates to display the number and location of household and settlement agents on the model landscape. Monitors track: the current population; percentage of remaining people relative to the start of the simulation; the amount of land transformed (total and accessible); the percentage of the landscape converted for cropland and agroforestry; and the minimum number of additional households per settlement that could be supported given leftover resources.

Plots on the interface further track: the cumulative number of people that have died; the number of people stressed due to shortages of maize and protein; the percentage of the landscape altered; the number of cropland and agroforestry plots; and the percentage of accessible land altered.

Initialisation

Prior to initialisation, the user is able to specify which quadrant (North-East, North-West, South-East, South-West) of the MMR to display. When initialised, the model imports data from multiple georeferenced datasets: elevation; land-use; productivity; and the presence of earthworks (mounds, causeways, and canals). This data is assigned to patch variables, with patches adjusting their colour to reflect their land cover. Patches of “Old Growth Forest” are flagged as potentially able to become forested after abandonment.

The procedure then spawns settlements in locations where mounds are known to exist on the real landscape (Lombardo & Prümers 2010). Each settlement defines suitable patches within range for cultivation and extraction (farming, foraging, fishing) by its member households. The radius for these activities is defined by the user. Agriculture must take place on terrestrial land, but can be further restricted solely to forested or non-forested land. Fishing may only take place on land with an elevation value of <-0.5875. Foraging and fishing may only take place on patches that are closer to the settlement than to any other, as tracked by the settlement-owner variable.

Finally, each settlement spawns a user-defined number of households that are attached to it via a community tie link. A variable (disband-age) tracks the number of years this household remains active before disbanding. This is instantiated when the household spawns for the sake of simplicity. Household lifespan varies on a normal distribution based upon the mean of a user-defined constant. We set this value to thirty for the basis of our experiments.

Details

Convert-Land

In the convert-land procedure, households claim and convert patches on the virtual landscape to satisfy their resource demands. The procedure begins with defensive programming to prevent agents from trying to manage agroforestry when only savanna patches are available for cultivation. Following this, the procedure resets the demand values of each household from the previous timestep.

Households subsequently determine whether they need to claim additional patches to satisfy their resource demands. This is done by assessing whether they expect their existing territory to be able to produce a sufficient quantity of each resource. Households will always assess the amount of maize they expect to produce from their territory in kilograms (M_{es}):

$$M_{es} = 0.85 * n_c * M_c \tag{1}$$

The above calculation accounts for both the number of patches within their territory (n_{pct}) and the mean productivity of maize per hectare of suitable land (M_c), a constant determined by a slider on the user interface. We assume that households are accurately able to determine the average maize yield per hectare of land within the MMR, as well as anticipating a 15% loss in yield to pests (Ringhofer 2010:99). However, this knowledge is imperfect, and they cannot account for random interannual variability (see resource-production).

M_{es} is subsequently compared to the amount of maize demanded by the household (M_d) to determine whether the household suffers from regular maize shortages. As each household is assumed to comprise of two adults and three children, this demand value is fixed and determined by a constant within the model.

$$\text{Shortage} = M_d * O - M_{es} > 0 \tag{2}$$

$$\text{Severe Shortage} = (M_d * O - (M_{es} + M_c)) > 0 \tag{3}$$

THE JOURNAL OF ARTIFICIAL SOCIETIES AND SOCIAL SIMULATION

An expected maize shortage occurs where expected maize supply is unable to account for demand. Expected supplies are considered severely short if this disparity is larger than the supply of Maize that an additional hectare of cropland could provide. Please also note that households attempt to acquire more than their base demand to account for potential shortfalls. This is encoded as an overproduction modifier (O), selected by the user on a slider on the interface.

If households are set to intentionally cultivate the surrounding forest (determined by the intentional-agroforestry slider), they will assess whether their territory is expected to produce sufficient quantities of Tree crops (F_{es}), Fuelwood (W_{es}) and Palm Leaves (Pa_{es}):

$$F_{es} = N_{pf} * F_c \quad (4)$$

$$W_{es} = N_{pf} * W_c \quad (5)$$

$$Pa_{es} = N_{pf} * Pa_c \quad (6)$$

These calculations function similarly to maize, though are instead calculated based on land that is able to become forested (N_{pf}). This includes both active agroforestry and cropland that will subsequently become agroforestry after it is abandoned, preventing agents from repeatedly cultivating land until their territory begins to produce forest resources. Instead, they are assumed to forage for resources on unclaimed land until their plots become productive. When intentional-agroforestry is enabled, it is impossible for agents to own fallowed land that can become forested, though it is possible for them to own fallows in the open savanna. As with maize, households cannot account for natural variations in the productivity of these resources, but can accurately estimate their average yield per hectare. Agents do not expect to produce sufficient animal protein from their territory; it must be procured over a larger area.

Households expecting a resource shortage are separated into three categories: those severely short of maize, those short of maize and no other resource, and those with a shortage in at least one forest product. These categories are hierarchical and mutually exclusive. A severe maize shortage is prioritised over all other shortages. If a household has a shortage of forest products, it will be designated as short of forest resources even if it also lacks maize. Each category undergoes a different process of land conversion.

Following this, household agents determine whether to resolve their resource shortages either by reactivating fallowed land within their existing territory, or by choosing to claim new land. This is only possible when converting land to produce maize, as land agroforestry cannot immediately be reactivated.

The choice to reactivate fallowed land is modelled stochastically, with the chance based upon the probability-of-reactivation slider on the user interface. Households can only reactivate land that was fallowed no less than a user-defined number of timesteps ago, selected via the fallow-period slider. Land cannot be reactivated if it more than 15 years old, mirroring the reestablishment of woody plants (Finegan 1984; Peña-Claros 2003; Ringhofer 2010). If multiple patches are reactivated, they must be selected from the open savanna, mirroring the increased difficulty of clearing forested land. If no land can be reactivated, the choice is skipped. If the reactivated land is sufficient to satisfy household demand, no further patches will be claimed.

THE JOURNAL OF ARTIFICIAL SOCIETIES AND SOCIAL SIMULATION

If no fallowed land is available, or the household has decided against reactivation, they will instead select (a) new patch(es) to claim and cultivate. To do this, households will assess a subset of the patches within a set area of their settlement mound. Both the number of patches and the radius of eligibility are user-defined, chosen via the percent-patches-considered and agricultural-radius sliders respectively. The number of patches considered is proportional to the size of this radius. Not all patches within range are eligible; those with a land use category of 'water', or those already owned and cultivated by other households are automatically excluded. In addition, the user can further restrict households to only cultivate patches of open savanna, or alternatively only those that can become forested.

MoundSim LandUse offers three different methods for agents to select patches. A null model ("NULL") is available, which forces households to randomly choose one patch from the subset. The second, "CANAL-CAUSEWAY", is an iterative procedure that selects the patch closest to known earthworks on the landscape. This uses imported vector data denoting the locations of canals and causeways associated with the Casarabe culture. In the third model, "VARIABLES", households select patches by calculating a utility score (U_{sc}) for each potential site based on multiple characteristics:

$$U_{sc} = P_{sc} + E_{sc} + Fl_{sc} - D_{sc} - Df_{sc} + A_{sc} \quad (7)$$

Patches calculate individual component scores for: Productivity (P_{sc}); Elevation (E_{sc}); Flooding (Fl_{sc}); Distance (D_{sc}); Clearance Difficulty (Df_{sc}); and Aggregation (A_{sc}). Each variable is calculated independently and can be weighted, meaning the user is able to select whether each factor weighs heavily, lightly, or is entirely factored out of decision making. When weighted equally, each component provides the same value to overall utility.

Patches are considered more productive if they lie within the fertile sediment lobe described by Lombardo *et al.* (2012). Patches within this area have a productivity (P) of 1. Those outside have a productivity of 0. A score is calculated based on this value:

$$P_{sc} = P_w * (100 * P) \quad (8)$$

The final value is multiplied by the weighting for this criterion (P_w), selected by the user on the interface.

Elevation (E_{sc}) is the preference for households to select patches in areas less likely to flood regularly, dependent on the patches elevation variable (E). Scores vary between nothing below -2 elevation and maximum points above 2.04063. This reflects the elevation of the lowest mound on the contemporary landscape, a high point at which flooding does not occur and thus no additional benefit can be derived from selecting patches at higher elevation.

$$E_{sc} = E_w * \max(0, \min(100, (100 * (E + 2) / 4.04))) \quad (9)$$

JASSS
THE JOURNAL OF ARTIFICIAL SOCIETIES AND SOCIAL SIMULATION

Flooding (Fl_{sc}) is the preference for households to select patches in areas more likely to flood regularly. At present, it is unknown when members of the Casarabe culture may have cultivated maize. It is possible for this to have occurred during the dry season, when much of the savanna was available. At such a time, access to water may have been crucial. This calculation acts as a reverse of the above slider, with no value derived from land with an elevation value of +2.04063 or more, and no additional benefit gained from land below -2. On this landscape, extremely low-lying areas are considered swamps and thus unsuitable for agriculture:

$$Fl_{sc} = Fl_w * 100 - \max(0, \min(100, (100 * (E + 2) / 4.04))) \quad (10)$$

The distance score (D_{sc}) of a patch decreases as its distance increases from the settlement to which the household is attached. This applies relative to the radius (R_a) within which households are able to cultivate, with no value added when patches are at the outer limit of the radius The loss of value increases non-linearly:

$$D_{sc} = D_w * 100 * (D^2 / R_a^2) \quad (11)$$

The difficulty component (Df_{sc}) decreases as its vegetation score (V) increases. At low levels, the patch is considered easy to clear and thus of increased value:

$$Df_{sc} = Df_w * 100 * (V / V_{Fmax}) \quad (12)$$

This value applies relative to the maximum vegetation score of forested land, as defined by the user (V_{Fmax}). The final component, Aggregation (A_{sc}), increases patch suitability if the patch borders another currently either: i) being actively managed as cropland or agroforestry or ii) is currently fallowed. As with the other mechanics above, it can be weighted (A_w).

$$A_{sc} = A_w * (\text{Border} \rightarrow 100, \neg\text{Border} \rightarrow 0) \quad (13)$$

Under certain circumstances, a household may not be able to find a patch of suitable land. For example, they may fail to find enough space as all patches within range have been claimed. If such a scenario arises, households will instead attempt to reactivate a patch of fallowed land they have access to, or instead sacrifice plots of agroforestry to produce sufficient quantities of maize (as the staple crop, maize is considered more important).

Once a site is selected, it is converted into cropland. Fallowed territory is reactivated by transitioning to the “Cropland” land-use category, resetting the last-converted value to the current tick count, and reducing the patch vegetation score to 0 (to reflect the vegetation being cleared). Resource supply scores are also reset to prevent the patch from supplying incorrect values. When creating new cropland, the land-use, last-converted, vegetation, and resource supply scores are set in an equivalent way. In addition, the patch is added to the household's territory (stored in the household-territory household variable), and has its owner variable set to the households ID number. Transition to agroforestry and fallowed land is managed within the vegetation-growth sub-model.

Resource-Production

In the resource-production procedure, resources are assigned to patches owned by household agents on the virtual landscape. Firstly, the communal resource pools of settlement agents are reset to prevent interannual storage. Resources are then assigned to patches according to their land-use type. Three types of land-use are considered: Cropland, Fallow, and Agroforestry. Only patches owned by household agents receive these resources. These supplies overwrite those remaining from the previous year to prevent storage.

	Land-Use Type	
	Cropland	Forest [Agroforestry/Fallow]
Maize [kg ha^{-1}]	[Varies]	0
Forage [kg ha^{-1}]	0	760
Fuelwood [kg ha^{-1}]	0	1200
Palm [leaves ha^{-1}]	0	206
Protein [kg ha^{-1}]	0	0.23

Table 1: Resources assigned to each patch in the resource-production procedure

Cropland patches are assumed to exclusively produce maize, and we assume no decrease in yield takes place throughout the cultivation period for the sake of simplicity. The productivity of maize per hectare (M_S) is determined by a slider on the user interface (cropland-production; M_c). In contrast, agroforestry, and forested fallow land produces forest resources (F_s, W_s, Pa_s, Pr_s), which are defined by constants (F_c, W_c, Pa_c, Pr_c). These values can also be multiplied by the user-defined forest production modifier (Fpi), assumed to represent forest enrichment due to human intervention. It is also modulated based on the patch vegetation score (V), discussed more under the 'vegetation-growth' procedure. For all resources, the quantity produced varies upon a normal distribution with a standard deviation of 15% of the mean value:

$$M_s \sim M_s * N(M_c, 0.15 * M_c) \quad (14)$$

$$F_s \sim Fpi * V * N(F_c, 0.15 * F_c) \quad (15)$$

$$W_s \sim (Fpi * V * N(W_c, 0.15 * W_c) \quad (16)$$

$$Pa_s \sim (Fpi * V * N(Pa_c, 0.15 * Pa_c) \quad (17)$$

$Pr_s \sim Fpi * V * N(Pr_c, 0.15 * Pr_c)$ (18)

Decrease-Population

This procedure represents the processes associated with the dissolution of households, as well as the death of individuals within them. There are two primary mechanisms of household dissolution in MoundSim LandUse: i) age and ii) insufficient resources.

MoundSim LandUse keeps the number of household agents fixed during a simulation, except in the event of resource shortage. There is no procedure to increase the number of households, and they can only be lost as a result of insufficient resources. This procedure first manages the dissolution of households due to age. Over time, the elderly will die, and children are assumed to marry and start their own households. For simplicity, we represent these processes as the dissolution of households after a certain period of time. This is encoded as a "disband-step" (Y_d) variable, initialised when the household is first created, a value that reflects the timestep during which the household dissolves:

$Y_d \sim T + N(Y_c, 0.1 * Y_c)$ (19)

In the above formula, T reflects the current timestep and Y_c reflects the mean lifespan of a household before dissolving. This is implemented as a constant within the model. The model currently varies this value based upon a normal distribution, with a standard deviation of 10% of the mean value. Within our experiments, this is set to a mean of 30 and standard deviation of 3. This ensures each household will last for 30 timesteps on average, with 95% of households existing for between 24 and 36 timesteps. When a household is dissolved in this way, a new one is automatically spawned in its place. This new household possesses no territory of its own and must claim land in subsequent timesteps. Resource demands of this agent are set to zero for this timestep to ensure it has a chance to act. The territory of the old household is abandoned.

The procedure then determines whether household agents have produced sufficient quantities of each resource. Each household first calculates whether sufficient quantities of Maize, Tree Crops, Fuelwood, Palm Leaves, and Protein have been obtained from its territory:

$M_{ts} = 0.85 * \Sigma\ (i=1 \dots n_{pt})\ M_{si}$ (20)

Here, n_{pt} refers to the number of patches within the household's territory. This is similarly calculated for forest resources:

$F_{ts} = \Sigma\ (i=1 \dots n_{pt})\ F_{si}$ (21)

$W_{ts} = \Sigma\ (i=1 \dots n_{pt})\ W_{si}$ (22)

$Pa_{ts} = \Sigma\ (i=1 \dots n_{pt})\ Pa_{si}$ (23)

$Pr_{ts} = \Sigma\ (i=1 \dots n_{pt})\ Pr_{si}$ (24)

Each household then calculates whether its demands have been satiated by the supplies calculated above, or whether residual demand remains (M_r, F_r, W_r, Pa_r, Pr_r):

$$M_r = M_d - (0.85 * M_{ts}) \quad (25)$$

$$F_r = F_d - F_{ts} \quad (26)$$

$$W_r = W_d - W_{ts} \quad (27)$$

$$Pa_r = Pa_d - Pa_{ts} \quad (28)$$

$$Pr_r = Pr_d - Pr_{ts} \quad (29)$$

As described above, Maize supplies are further multiplied by 0.85 to account for unintended crop loss (Ringhofer 2010:99).

If a household suffers from a resource shortfall, it may still survive by obtaining resources from i) communal stores owned by its settlement and ii) unclaimed patches of land on the landscape. The former is only possible if the communal pooling of resources is enabled via the communal-resource-pools? switch. Households may only obtain Maize, Fuelwood, and Animal protein in this way. Palm leaves are typically collected on demand, and Foraged tree crops are supplementary to the diet and cannot be stockpiled. In comparison, it is common for maize and fuelwood to be stockpiled for consumption intra-annually, and animal protein is commonly shared in contemporary Amazonian communities (In some cases the lack of sharing can lead to disunity e.g., Good (1987)). These resources are shared freely under the assumption of generalised reciprocity (donor households would expect to receive similar aid in a similar situation). Resources cannot be shared between settlements.

Households may also obtain resources from unclaimed patches of forested land controlled by their settlement. Settlements may directly control patches of territory up to 7km (within daily walking distance (Beckerman 1987; Binford 2001)). However, for it to be controlled, it must be nearer to that settlement than to any other. Only settlements that possess household members may control territory:

THE JOURNAL OF ARTIFICIAL SOCIETIES AND SOCIAL SIMULATION

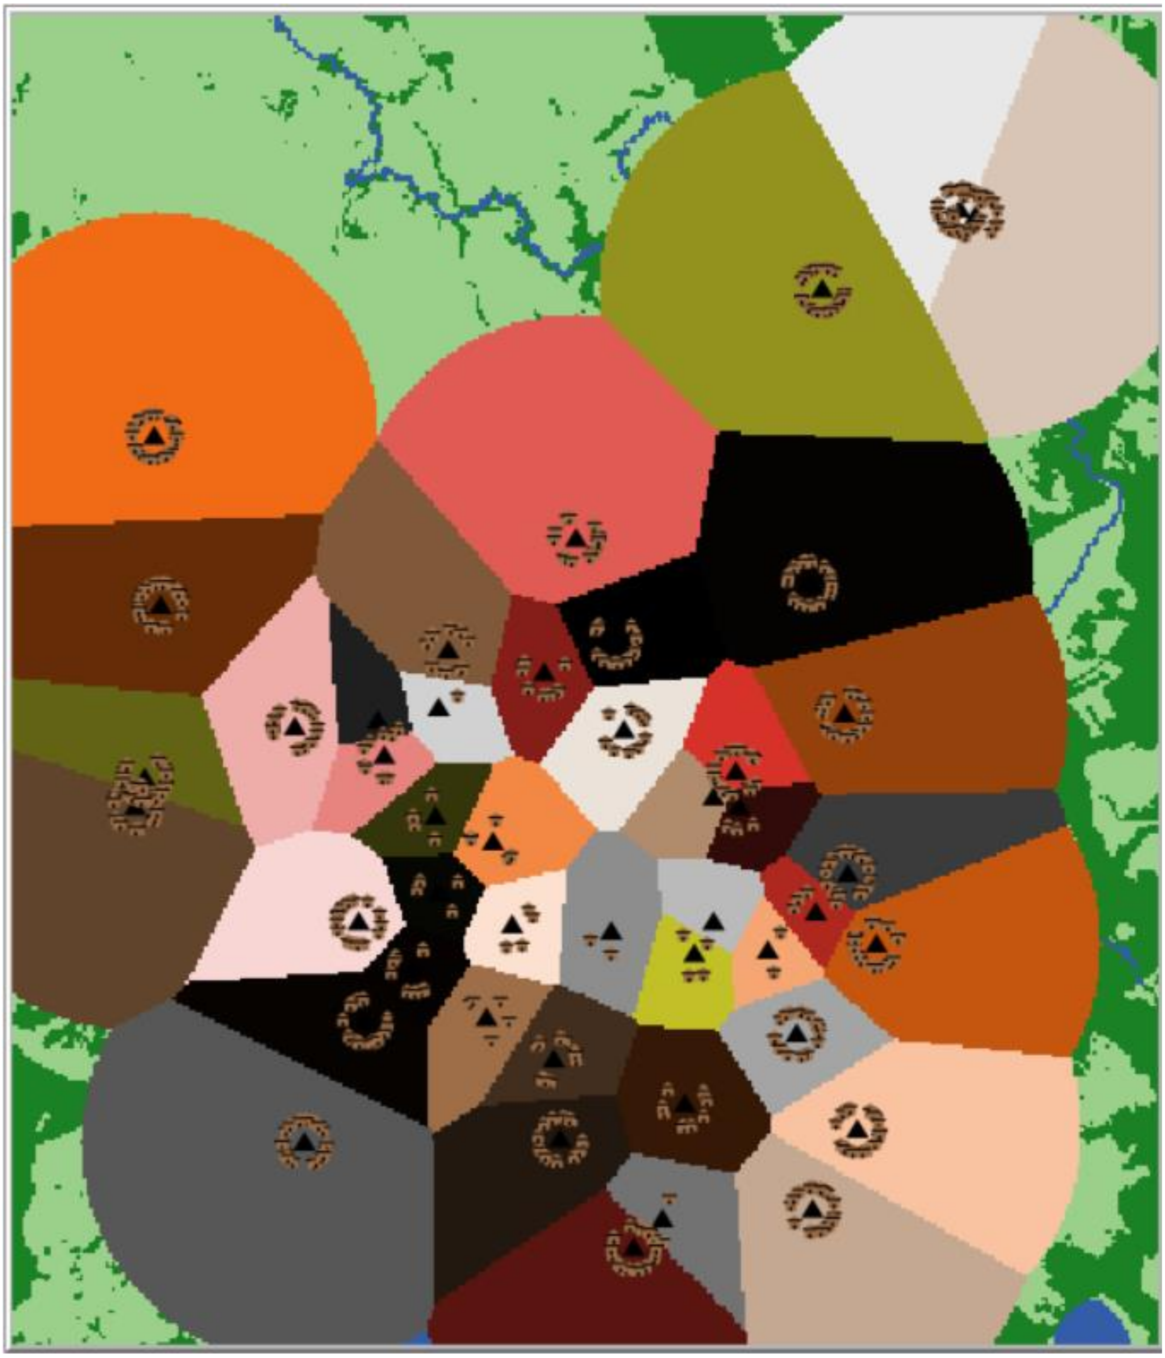

Figure 1: Patches controlled by settlements within the North-East Quadrant of MoundSim LandUse. This initialised upon startup and tracked using the settlement-owner patch variable.

Unclaimed patches of forest (patches with a land-use value of either "Old Growth Forest" or "New Growth Forest") that meet these criteria can provide additional Foraged Tree Crops, Fuelwood, Palm Leaves and Animal protein to the households of that settlement. This is managed through the settlement agents. First, each settlement produces a value reflected the quantity of each resource produced per hectare by unclaimed patches it controls. This is done using the same formulae as those described in equations 14-17. Second, it calculates the total number of old (n_{OldFor}) and new (n_{NewFor}) growth forest patches it controls. From these two values, the calculates the total quantity of each resource that can potentially be extracted from unclaimed forest patches it controls (F_u, W_u, Pa_u, Pr_u):

$$F_u = (Fpi * F_s * n_{NewFor}) + (F_s * n_{OldFor}) \quad (30)$$

$$W_u = (Fpi * W_s * n_{NewFor}) + (W_s * n_{OldFor}) \quad (31)$$

$$Pa_u = (Fpi * Pa_s * n_{NewFor}) + (Pa_s * n_{OldFor}) \quad (32)$$

$$Pr_u = Pr_s * (n_{NewFor} + n_{OldFor}) \quad (33)$$

In addition, each settlement is able to extract additional animal protein from fish produced at low elevations in the open savannas:

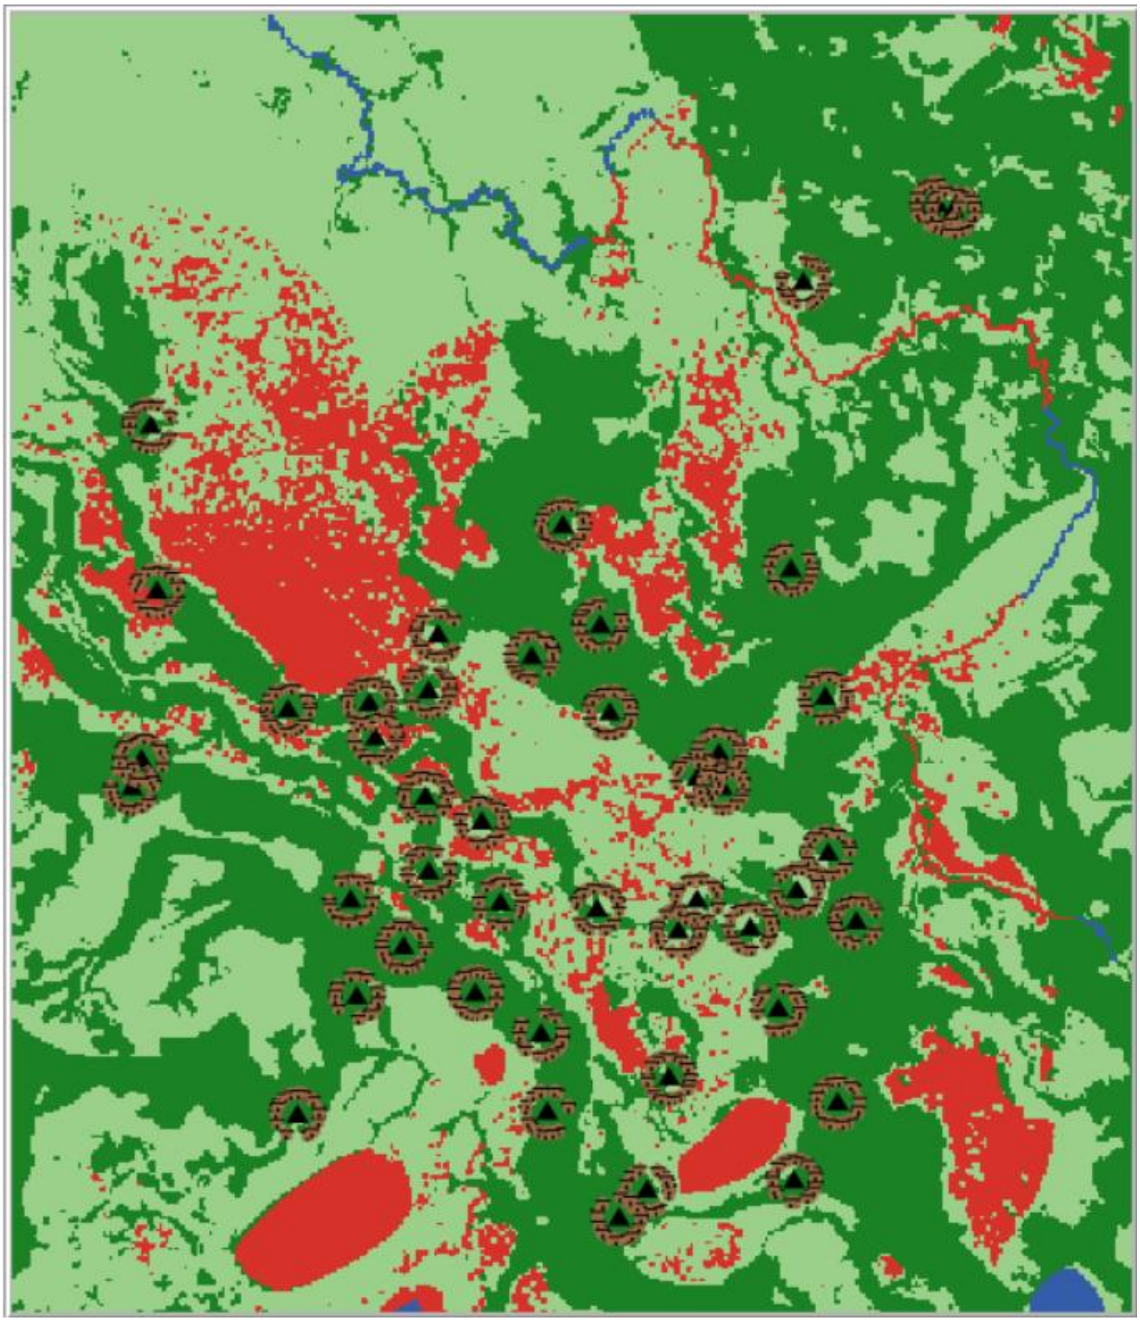

Figure 2: Location of fish producing patches within the northeast quadrant of MoundSim LandUse. To be considered viable, patches must possess an elevation value below -0.5875. They must also be controlled by that settlement (managed by the settlement-owner patch variable).

Aquatic protein is only accounted for if the fishing-modifier slider is set to above 0.0 on the user interface. The total quantity of unclaimed fishing protein available for extraction depends upon the number of patches that can produce fish within territory controlled by the settlement (n_{Fi}), as well as a user-defined density of fish available within a hectare of eligible land (Den_{Fi}):

$$Fi_u = n_{Fi} * 0.084 * Den_{Fi} \quad (34)$$

Households experiencing a shortage in any of these resources can then obtain unclaimed resources from these areas.

Finally, the procedure identifies household agents that were unable to acquire sufficient quantities of each resources from any of the above methods. The model allows the user to determine whether agents die as a result of these shortfalls. This is determined by the *-die? Switches on the user interface. If a household lacks sufficient quantities of a resource with the respective *-die? switch turned on, that household will abandon its territory and die. However, if the switch is turned off, the household will instead be marked as under resource stress (with the respective *-stressed? variable switched to "True") until the next time decrease-population is executed. This allows the user to track the number of individuals stressed as the result of a particular resource shortage. The number of dissolved households is also tracked separately by resource for this purpose.

When household territory is abandoned, the patch's owner variable is reset. If the territory is actively being farmed as cropland, the land-use variable will be set to "Fallow".

Vegetation-Growth

The vegetation-growth procedure represents the continuous growth of vegetation over time. For the sake of parsimony, we solely focus on secondary succession taking place on land abandoned by household agents, tracked using a boolean variable. The first task managed by this procedure is succession on abandoned land. In the LM, local edaphic and hydrological conditions cause succession to vary between forests on high ground and low-lying savanna (Mayle et al. 2007). This begins after land is first cleared, but is stalled until the end of the cropping cycle by humans through regular weeding (e.g., Beckerman 1987; Ringhofer 2010; Staver 1989), proceeding unhindered thereafter. In MoundSim LandUse, we define the point when land is considered to be regenerated the point at which local Indigenous communities can no longer differentiate between fallow and secondary forest (Huanca 1999:59). The model enables the user to vary this parameter, which is defined by sliders on the user interface.

Within the model, we account for the process of succession using the vegetation patch variable. This is calculated as a normally distributed value with a mean and standard deviation dependent on the current timestep (t) and the timestep during which the patch was last clear (t_{lc}). These two variables can be used to calculate the number of years of secondary succession that has taken place unhindered. This is converted to a vegetation score by transforming it based on the natural logarithm:

$$V \sim N(\log_e(t - t_{lc}), \log_e(0.0025 * \log_e(t - t_{lc}))) \quad (35)$$

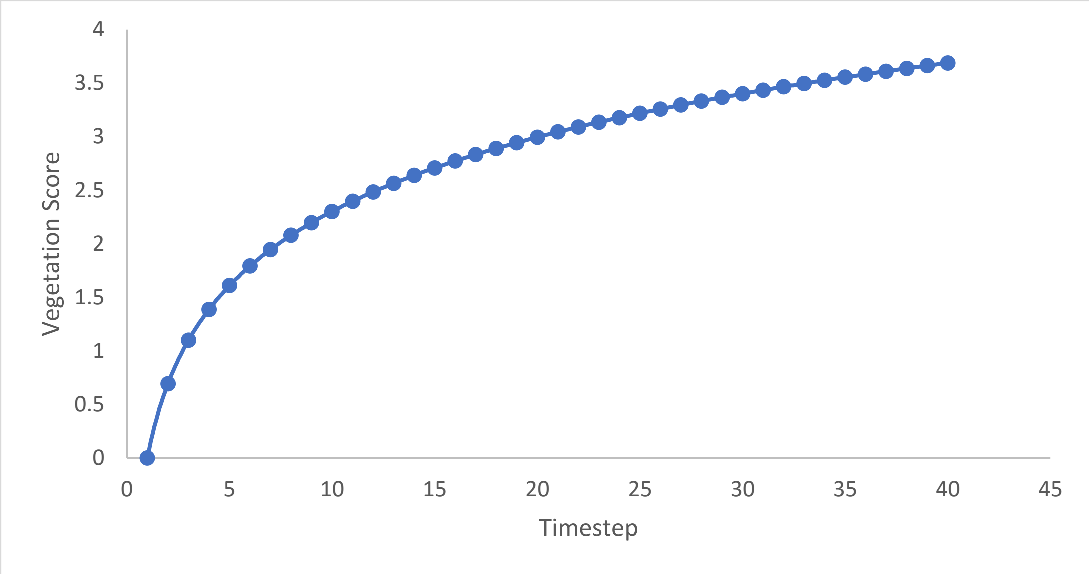

Figure 3: mean vegetation score for patches during timesteps subsequent to abandonment.

Once a patch's vegetation score exceeds a threshold, it will transition to New Growth Savanna or New Growth Forest depending on the relevant succession pathway. This threshold depends on which pathway the patch takes. This value can vary between 10 and 120 years within the model following the results of Poorter *et al.* (2021). There is a lack of information regarding the rate of savanna regeneration in the LM, so we assume a 5-50 year transition rate to account for the faster establishment of herbaceous vegetation. The model follows the Tsimane system of agriculture in that households retain land rights after the land is fallowed (Piland 1991), which is subsequently lost after it becomes indistinguishable from secondary forest/savanna.

The second task managed is the abandonment of unproductive cropland. Contemporary cropping systems typically produce crops for three years before fallowing (e.g., Beckerman 1987; Denevan 2001; Piland 1991; Staver 1989). In MoundSim LandUse, the user is able to determine whether the land is abandoned between one and seven years after first cultivation. The procedure also ensures that households only manage up to six patches of land within their territory at any one time (Beckerman 1987; Piland 1991). In such cases, the oldest patch is abandoned. Finally, the procedure ensures that any patches abandoned by their owner become fallowed and are entered into the succession system.

References

ARTHUR, B. (1994). Inductive Reasoning and Bounded Rationality. *The American Economic Review*, *84*(2), 406–411.

BECKERMAN, S. (1987). Swidden in Amazonia and the Amazon Rim. In B. L. Turner II & Stephen. Brush (Eds.), *Comparative Farming Systems* (pp. 55–94). The Guilford Press.

BINFORD, L. (2001). *Constructing Frames of Reference: An Analytical Method for Archaeological Theory Building Using Ethnographic and Environmental Data Sets* (1st ed.). University of California Press.

BRUNO, M. (2010). Carbonized Plant Remains from Loma Salvatierra, Department of Beni, Bolivia. *Zeitschrift Für Archäologie Außereuropäischer Kulturen*, *3*(151), 153–208.

DENEVAN, W. (2001). *Cultivated Landscapes of Native Amazonia and the Andes*. Oxford University Press.

DICKAU, R., Bruno, M., Iriarte, J., Prümers, H., Jaimes Betancourt, C., Holst, I., & Mayle, F. (2012). Diversity of cultivars and other plant resources used at habitation sites in the Llanos de Mojos, Beni, Bolivia: Evidence from macrobotanical remains, starch grains, and phytoliths. *Journal of Archaeological Science*, *39*(2), 357–370. https://doi.org/10.1016/j.jas.2011.09.021

FINEGAN, B. (1984). Forest Succession. *Nature*, *312*(5990), 109–114.

GOOD, K. (1987). Limiting Factors in Amazonian Ecology. In M. Harris & E. Ross (Eds.), *Food and Evolution: Towards a Theory of Human Food Habits* (pp. 207–421). Temple University Press. https://archive.org/details/foodevolutiontow0000unse/page/406/mode/2up

GRIMM, V., Berger, U., Bastiansen, F., Eliassen, S., Ginot, V., Giske, J., Goss-Custard, J., Grand, T., Heinz, S. K., Huse, G., Huth, A., Jepsen, J. U., Jørgensen, C., Mooij, W. M., Müller, B., Pe'er, G., Piou, C., Railsback, S. F., Robbins, A. M., … DeAngelis, D. L. (2006). A standard protocol for describing individual-based and agent-based models. *Ecological Modelling*, *198*(1–2), 115–126. https://doi.org/10.1016/j.ecolmodel.2006.04.023

GRIMM, V., Berger, U., DeAngelis, D., Polhill, G., Giske, J., & Railsback, S. (2010). The ODD protocol: A review and first update. *Ecological Modelling*, *221*(23), 2760–2768.

HUANCA, T. (1999). *Tsimane Indigenous Knowledge, Swidden Fallow Management and Conservation*. University of Florida.

JAIMES BETANCOURT, C. (2012). *La cerámica de la Loma Salvatierra, Beni-Bolivia*. Plural Editores.

JAIMES BETANCOURT, C. (2015). La cerámica de la Loma Mendoza. In H. Prümers (Ed.), *Loma Mendoza. Las Excavaciones de los años 1999-2002* (pp. 89–222). Plural Editores.

Lake, M. (2014). Trends in Archaeological Simulation. *Journal of Archaeological Method and Theory*, *21*(2), 258–287. https://doi.org/10.1007/s10816-013-9188-1

LOMBARDO, U., May, J. H., & Veit, H. (2012). Mid- to late-Holocene fluvial activity behind pre-Columbian social complexity in the southwestern Amazon basin. *Holocene*, *22*(9), 1035–1045. https://doi.org/10.1177/0959683612437872

LOMBARDO, U., & Prümers, H. (2010). Pre-Columbian human occupation patterns in the eastern plains of the Llanos de Moxos, Bolivian Amazonia. *Journal of Archaeological Science*, *37*(8), 1875–1885. https://doi.org/10.1016/j.jas.2010.02.011

MAGLIOCCA, N., & Ellis, E. (2016). Evolving human landscapes: a virtual laboratory approach. *Journal of Land Use Science*, *11*(6), 642–671. https://doi.org/10.1080/1747423X.2016.1241314

MALTHUS, T. (1798). *An Essay on the Principle of Population* (1st ed.). Pickering & Chatto Publishers.

MAYLE, F., Langstroth, R., Fisher, R., & Meir, P. (2007). Long-term forest-savannah dynamics in the Bolivian Amazon: Implications for conservation. *Philosophical Transactions of the Royal Society B: Biological Sciences*, *362*(1478), 291–307. https://doi.org/10.1098/rstb.2006.1987

MULLER, B., Bohn, F., Dreßler, G., Groeneveld, J., Klassert, C., Martin, R., Schlüter, M., Schulze, J., Weise, H., & Schwarz, N. (2013). Describing human decisions in agent-based models - ODD+D, an extension of the ODD protocol. *Environmental Modelling and Software*, *48*(1), 37–48. https://doi.org/10.1016/j.envsoft.2013.06.003

PENA-CLAROS, M. (2003). Changes in Forest Structure and Species Composition during Secondary Forest Succession in the Bolivian Amazon. *Biotropica*, *35*(4), 450–461. https://doi.org/10.1111/j.1744-7429.2003.tb00602.x

PILAND, R. (1991). *Traditional Chimane agriculture and its relation to soils of the Beni biosphere reserve.* University of Florida.

POORTER, L., Craven, D., Jakovac, C., van der Sande, M., Amissah, L., Bongers, F., Chazon, R., Farrior, C., Kambach, S., Meave, J., Muñoz, R., Norden, N., Rüger, N., van Breugel, M., Zambrano, A., Amani, B., Andrade, J. L., Brancalion, P., Broadbent, E., … Hérault, B. (2021). Multidimensional tropical forest recovery. *Science*, *374*(6573), 1370–1376.

PRIMERS, H. (2015). Loma Mendoza. Las Excavaciones de los años 1999-2002. In *Kommission für Archäologie Außereuropäischer Kulturen des Deutschen Archäologischen Instituts, 2015*. Plural Editores. http://publications.lib.chalmers.se/records/fulltext/245180/245180.pdf%0A

PRUMERS, H., Jaimes Betancourt, C., Iriarte, J., Robinson, M., & Schaich, M. (2022). Lidar reveals pre-Hispanic low-density urbanism in the Bolivian Amazon. *Nature*, *606*(7913), 325–328. https://doi.org/10.1038/s41586-022-04780-4

RINGHOFER, L. (2010). Fishing, Foraging and Farming in the Bolivian Amazon. In *Fishing, Foraging and Farming in the Bolivian Amazon*. https://doi.org/10.1007/978-90-481-3487-8

STAVER, C. (1989). Why farmers rotate fields in maize-cassava-plantain bush fallow agriculture in the wet Peruvian Amazon. *Human Ecology*, *17*(4), 401–426. https://doi.org/10.1007/BF00889498

WHITNEY, B., Dickau, R., Mayle, F., Soto, D., & Iriarte, J. (2013). Pre-Columbian landscape impact and agriculture in the Monumental Mound region of the Llanos de Moxos, lowland Bolivia. *Quaternary Research (United States)*, *80*(2), 207–217. https://doi.org/10.1016/j.yqres.2013.06.005

WILENSKY, U. (1999). *NetLogo*. Center for Connected Learning and Computer-Based Modeling, Northwestern University, Evanston, IL. http://ccl.northwestern.edu/netlogo/