

A Simulation of Evolving Sustainable Technology Through Social Pressure

Daniel Rush

University of Michigan
danrush@umich.edu

Abstract

In this paper we develop a model to simulate the evolution of a pollution-free resource gathering technology that is initially less efficient but ultimately reaches parity with polluting technology. We find that for low levels of pollution, pressure exerted by society can indeed encourage the development and use of non-polluting technology, with greater pressure being associated with faster achievement of efficiency parity and lower overall pollution. However, greater pressure is also associated with lower populations and at the highest levels of pressure there are significant risks of population crashes. We find that these results hold for both localized pollution and globalized pollution, with globalized pollution encouraging faster achievement of efficiency parity. For high levels of pollution we find that introducing societal pressure significantly increases the occurrence of population crashes, and thus the strategy is only effective under certain conditions.

Introduction

As the population of the world increases humans continue to create technological tools and processes that allow for the greater production of usable resources, including energy and food. Fertilizers and drought-resistant crops enabled a green revolution to feed additional millions while dams and power plants have enabled flood control and electrification that raised the standard of living for billions. However, these technologies and processes are often accompanied by negative consequences. Fertilizers washed into rivers can feed oxygen-depleting algae blooms that damage fisheries, dams can interfere with the reproduction cycles of both plants and animals, and greenhouse gas emissions from power plants can contribute to climate change. Additionally, because these consequences may not be suffered by the persons who

caused them (and perhaps may not even be visible to them), they may not receive due consideration in the deliberations of a self-interested individual deciding how to implement technology to best grow food or provide power for a growing population. This interplay of individual decisions having an impact, in aggregate, on a much larger population suggests that this issue would be well examined as a complex system. Such systems are "characterized by the fact that they were made of simple interacting elements that produced through their aggregate behavior a global emergent order unpredictable simply through analysis of low level interactions" (Helmreich and Stefan 1998). The balance of the immediate needs of an individual with those of others that are spatially or chronologically distant is also at the crux of sustainability, which the WCED defines as "[meeting] the needs of the present without compromising the ability of future generations to meet their own needs"(WCED 1987). While sometimes the lack of a sustainable technology is a barrier to adoption, this is not always the case. Even when sustainable technology is available, a difficulty with many sustainable practices is that the benefit of their use is realized by others (either in different localities or perhaps in future generations) while the costs of their use (such as reduced efficiency or yield) is borne by the people deciding to use them. One mechanism society can use to encourage choices that balance needs that are current and local against needs that are elsewhere and future is to develop social standards that impact the person making the current and local decision. In this way, society can create the conditions under which people have the capacity to meet their own needs, but are encouraged to consider the needs of others once their own needs are accounted for. In this paper we build an agent based model with simple social pressures and use it to investigate how the conditions under which individual agents respond to those pressures result in the gradual improvement of sustainable

(non-polluting) technology until it is as efficient as traditional (polluting) technology.

As we start exploring this question, it may be useful to think of technology as an attribute or endowment unique to each agent that allows it to convert resources from its physical environment into a usable form. While in reality there are often many alternative technologies that can be applied to achieve a particular end (e.g. producing electricity can be achieved via a variety of methods from a number of natural resources), in this simulation model there are only polluting and non-polluting technologies and both are applied to harvesting food. For simplicity, every agent is endowed with the most efficient, lowest cost polluting technology as a 100% efficient default option and one sustainable technology that they could utilize instead of the default if they chose to. Applying this concept to a hypothetical example, an agent's choice in our model might be analogous to a consumer choosing whether to use their entire electricity budget to purchase either 1,500 kilowatt-hours of electricity produced by a natural-gas combusting power-plant that produces some amount of pollution or 750 kilowatt-hours of electricity produced by windmills that give off negligible pollution. Under classical economics, the self-interested agent will maximize its utility by choosing the technology that produces the most output for the lowest cost. In our example, this is the natural-gas derived electricity, which offers twice as much power for the same cost. Instead of attempting to attach a price to the pollution and forcing the agent to consider that cost when choosing which technology to utilize (a classical price-based strategy to account for externalities), the model described in this paper attempts to determine if we can create simple social pressures that would encourage the use of non-polluting technology. To account for the variety of sustainable technologies and resources they can be used to harvest that exist in the real world, we let the non-polluting technology possessed by each agent to potentially differ from those held by other agents along the single dimension of efficiency when applied to harvesting resources. In our earlier example, wind power had 50% of the efficiency (defined as kWh per \$) of the polluting technology. It could also be possible that people near a coast do not have access to windmill power, but have the option to buy power generated by tidal booms that were only 15% as efficient as the baseline polluting natural-gas technology. Similarly, others near a desert might have access to solar power that is 25% as efficient as the baseline technology. Thus in the model, one agent may have a sustainability technology that is 15% efficient in converting the environment's resources into food, while another's may be 65% efficient (to mimic a wide variety of technologies, our model will grant initially generated agents sustainable technologies with efficiencies uniformly distributed across the range of (0, 1) - see iteration 4 in the next section for

details). By allowing the % efficiency to vary for each agent, we enable the model to roughly encompass the diversity of technologies that may exist. We also create the conditions for each agent's decision on whether to use that technology or the polluting technology to be unique to themselves. Two agents facing the same amount of physical resources from their environment and the same needs may make different choices on whether to use their non-polluting technology based on the % efficiency associated with the technology they have access to. This makes sense, as a person who needs 800 kWh and has an 85% efficient non-polluting technology can meet that need through spending their fixed budget on 1000 kWh polluting electricity or by using their sustainable technology to yield 850 kWh for the same price. However, another person with only 65% efficient non-polluting technology may prefer to use the fully efficient polluting technology because otherwise their needs will not be met.

Now that we have agents with non-polluting technology of varying efficiencies, we consider how society can encourage the use of those technologies and how they might be improved over time so that they can eventually be as efficient or more-efficient than the polluting technologies. A simple mechanism for society to encourage the use of sustainable technology is for it to develop a social standard that agents apply while finding a mate. Such a standard might say that those who pollute more are less attractive for reproduction than those who pollute less. This could be similar to a person being more inclined to mate with a person who is usually well over someone who is more often sick. However, instead of being concerned with the potential mate's positive contribution to an offspring's ability to fight off sickness, an agent is concerned with the potential mate's negative contribution (pollution) to the world that the offspring will live in. Of course, everyone gets sick sometimes, and everyone may need to pollute at some point, too. This implies that an important part of the social standard is a level below which there is no reproductive impact. We will see such a standard and level implemented in iteration 3 in the next section.

Turning to the question of how technology could improve over time, we want to establish a way for an agent's sustainable technology to be inherited by its offspring. Consider the context where agents are harvesting a natural resource, say food, using the tools that they have available: either the traditional, 100% efficient way that gives them all the food available but creates pollution (e.g. by requiring a forest clear-cut and burn), or a less-efficient, but lower-yielding approach that does not create pollution (e.g. by forgoing the cut and burn and instead utilizing natural breaks in the forest augmented by pruning and bringing in compost for nutrients). It may make sense to think that while everyone shares a common

knowledge of the traditional (standard-practice) harvesting approach, only the children of parents who have a particular sustainable technology will inherit the training and all the relevant details necessary to utilize it. This inheritability is implemented in iteration 4 of the next section. Finally, by combining a form of inherited sustainable technology with society's pressure to remain attractive for reproduction by not polluting, we have the conditions necessary for the 'evolution' of sustainable technology over generations and can formally lay out our simulation.

An Iteratively-Built Model

The model is based on Epstein and Axtell's Sugarscape presented in *Growing Artificial Societies (GAS)* with pollution and sexual reproduction (Epstein and Axtell 1996). It is built in NetLogo (Wilenski 1999) as a modification of the University of Leicester's implementation of the Sugarscape (Weaver 2009), which was built to encompass almost all the rules discussed in GAS. As we have implemented it in this paper, only one resource grows on the grid (it could be conceived of as sugar or energy), but can be harvested by agents, which need it to survive, either with pollution or without.

For replicability and understanding, the model will be presented as five iterations built on top of one another. Iteration 1 describes the base model and reproduces the results Epstein and Axtell achieved with one version of their Sugarscape. Iterations 2-4 introduce successive alterations to the model that provide the environmental and agent traits that serve as building blocks for our eventual final model. That model, described in iteration 5, is where agents choose between traditional polluting and sustainable non-polluting technology. These aggregate choices impact the pollution produced, the viability of the population as a whole, whether the sustainable technology becomes efficient and the speed at which that occurs. Multiple experiments are then performed with the model from Iteration 5 to identify interesting parameter combinations and to establish the robustness of the identified results. Results given for iterations 2-4 are to deepen the reader's understanding of the build process and, while instructive, should not be considered robust. For iterations 2-5, unless otherwise noted, the settings from the previous iteration are retained.

Iteration 1: Base case with sexual reproduction

We first create a docking case that reproduces the results of Epstein and Axtell's single good (sugar-only) Sugarscape with sexual reproduction. This provides a starting point and will set the ground rules that will either continue in

future iterations or be modified as described in later iterations that lead up to our final simulation.

Description

This implementation of the Sugarscape is on the surface of a toroid (think of a flat map rolled and bent to look like a donut so that the top is connected to the bottom and the right and left edges are connected to each other) on which there are 2,601 patches (a 51x51 patch grid) that grow sugar at some rate (we will call this "**Rule G**" for growth) to some set maximum. As in many scenarios in GAS, the patch maximums are arranged to create two sugar 'hills' that grow up to 4 sugar at their peaks with rings around them that grow 3 and then 2 sugars, with most of the areas surrounding the hills only growing up to 1 sugar, and two 'desert' areas growing no sugar. Any patch may be inhabited by an agent, and all agents are endowed with a gender, fertile age range (measured in 'ticks'), metabolism and vision. For each step (tick) of the simulation, agents use their vision endowment to look as many patches in each of the four cardinal directions (this is called their neighborhood) that they can before deciding what patch to move to. Agents can move to any patch they can see and in this iteration will move each tick to the unoccupied patch that has the most sugar available for harvesting (this is called "**Rule M**" for movement). The agents then harvest the sugar available and add it to the store of sugar that they carry around with them. They will then eat as much sugar as their metabolism dictates from their store and, if at any point in the simulation their store falls to zero sugar, they die. If an agent has as much sugar as it initially starts the simulation with (this is determined randomly within parameters set up at the beginning of the simulation) and is of fertile age (also determined randomly within the ranges of 12-60 or 12-40 based on gender) it is free to produce an offspring with a fertile agent of the opposite gender that is in its neighborhood (this is called "**Rule S**" for sexual reproduction). Its partner must also have enough sugar and there must be an unoccupied patch in the agent's neighborhood for the child to be born onto. If all conditions are met, each parent's sugar endowment is decreased by half and a child agent is created near the parent agent, sharing a mix of the two parents' attributes (each inherited attribute is selected at random from one of the parents - i.e. there are no 'dominant' attributes).

Rules	Parameters
Agent: M, S	Initial Sugar: 10-35; Metabolism: 1-4; Agents: 400; Vision: 1-6
Patch: G	Growback Rate: 1; Growback Interval: 1

Table 1: Summary of Rules for Iteration 1

Results

Running the model, we observe an emergent cyclical population that cycles between 800 and 1300 and lives almost exclusively on the sugar hills (leaving the marginal

areas with only 1 sugar unpopulated). The agents are so densely packed together that there is not much space on the mounds for offspring to be born onto. Both higher vision and lower metabolism are selected for so that the average values of these across the population increase and decrease respectively. These results compare to Epstein and Axtell's in almost all respects, including periodic population crashes observed by the combined effects of lowering the initial sugar requirements in tandem with a 10 year reduction in the fertile ages of both males and females. Restoring the fertile ages and initial sugar requirements to the values listed in the parameters above, we observe a stable population (within cyclical boundaries) that selects for better vision and lower metabolism while living on the most fertile portions of the grid.

Iteration 2: Turn on pollution

We next introduce the concept of pollution and alter the movement rules of the agents so that their behavior is affected by it. This introduces core concepts to our model that will eventually allow us to set the stage to try and reduce pollution while evolving technology. It also results in the interesting observation that more agents live on a Sugarscape with pollution than on one without it.

Description

Pollution is implemented for harvesting only, meaning that whenever a sugar is harvested by an agent for placement into its store, pollution is produced. The pollution produced is a user-set percentage of the sugar harvested known as the pollution rate. Collectively these conditions are called "**Rule P1**." This is different from the pollution described in GAS, whose agents created pollution during both harvesting and consumption activities (consumption occurs when the agent eats sugar from its stores as dictated by its metabolism). In modifying the movement rule ("**Rule M1**"), we adopted the convention from GAS where agents simply do not like pollution, and when selecting the patch to move to, will consider the sugar to pollution ratio and select the patch with the highest ratio to move to (instead of simply the most sugar). In this iteration we keep pollution local to the patch it was produced by not turning on diffusion. However, we do introduce the new concept of pollution decay ("**Rule D**") in an attempt to recognize that not all pollution is permanent. In many cases, toxicities decline over time as pollution is broken down by natural processes (like shorter radioactive half-lives) or sequestered by organisms (like some portion of greenhouse gas emissions). To model this, we allow some user-set percentage of the pollution on a patch to disappear each tick. This makes for a more dynamic system.

Rules	Parameters
Agent: M1, S, P1	Same as in iteration 1
Patch: G, [D] ¹	Pollution Rate: .1 (10% of sugar harvested); Decay Rate: .05 (5%) per patch per tick

Table 2: Summary of Rules for Iteration 2

Results

Running the model, we observe a population that cycles between 1300 and 2100 agents and then exhibits ever smaller swings until it stabilizes at around 1700 agents. The population is spread over both sugar mounds as well as the less sugar-dense areas which only grow one sugar per patch. This pattern, that covers more area than the prior iteration where there was no pollution, arises because agents are induced to move into lower sugar regions due to their dislike of pollution (Epstein and Axtell describe these as environmental refugees). This movement into less sugar-dense areas frees up space in the higher-density areas for children to be born onto, giving the environment a greater carrying capacity than it had without the movement induced by the presence of pollution. Introducing decay allows pollution to cycle with the population (slightly lagging it). For sufficiently low decay rates (like the one adopted) pollution increases without bound, though at faster rates when the population is near its peak than when it is near a valley.

Iteration 3: Pollution impacts reproduction

This iteration sets the stage for selection pressures to be brought to bear on technology by allowing the pollution an agent has created while harvesting to affect its reproductive fitness. Much like better vision allows more effective sugar gathering and lower metabolism enables an agent to live longer from its sugar stores, we want to create conditions where agents that create less pollution through harvesting have a reproductive advantage (thus increasing the opportunities for the technology they carry to be inherited by offspring).

Description

The way we allow an agent's reproductive fitness to be impacted by pollution is to introduce a modification to the sexual reproduction rule. This modification expands the decision rule on whether to reproduce by introducing a social preference that, under certain conditions, will discourage reproduction with agents who have caused above-average amounts of pollution. This can be thought of as a social norm where potential mates will only choose to reproduce with below-average polluters. To forestall agents from avoiding reproducing with an agent that has caused only a trivial amount of pollution but has neighbors

¹ Note: [D] stands for decay, *not* diffusion (which is not enabled) and the brackets indicate that the model was run with and without it.

that haven't created any, we also introduce a socially acceptable level of pollution below which an agent is not subject to comparison to the neighborhood average and is thus automatically eligible for reproduction (presuming all the other criteria are met). **Rule S2** operationalizes the impact of pollution on reproduction and includes the following elements:

1) Agents have a counter that records how much pollution they have created while harvesting during their lifetime

2) This counter is available for inspection by potential mates, who will only reproduce with an otherwise eligible agent if that agent has caused average or less-than-average pollution (with the average being calculated locally among the agents in the agents-in-question's neighborhood, subject to a community standard exception explained below).

3) Based on its tolerance for pollution, society sets a community standard for 'acceptable' pollution below which an agent is not subject to the comparison to its neighbors and pollution therefore does not impact their reproductive chances.

Rules	Parameters
Agent: M1, S2, P1 Patch: G, D	Unless otherwise stated, same as in iteration 2; "Community Standard" critical value: [6, 7, 8]; Initial agents: [400, 600, 700]

Table 3: Summary of Rules for Iteration 3

Results

In this iteration, we find that introducing social norms that affect reproduction can have a profound impact on a population. This effect appears to be linked to the stringency of those norms (as represented by the "community standard" critical value). For a critical value of 8 (the least stringent tested in this iteration), the population cycles observed in prior iterations are still present as most agents are able to gather enough sugar to reproduce without exceeding the critical value by too much. For a critical value of 7, however, those cycles are only observed half the time, and the other half of the time the population does not reproduce often enough to sustain itself and the agent count goes to zero (a population crash). For a critical value of 6, the population crashes every time unless the number of initial agents is increased to between 600 and 700 (after which many, though not all, runs of the simulation maintain a viable population). This seems to indicate that when Rule S2 is in effect, either a higher density of agents or sufficiently low community standards (represented by higher critical values) are necessary to maintain a viable population.

Iteration 4: Technology introduced and can evolve (does not yet impact agent's harvesting ability)

This iteration represents an interim step where we modify the agents themselves so that they carry a technology attribute that has the characteristics necessary for our final objective. Intuitively, we are endowing them with technology and allowing it to change across generations, but not yet allowing it to be selected for by impacting an agent's harvesting ability or reproductive fitness (though we do modify the reproduction rule so that the technology will be inherited by an agent's offspring). As in the prior iteration, community standards and pollution still impact reproduction, but technology just can't do anything about it yet.

Description

Another building block needed for the evolution of technology is for agents to be endowed with the technology and for that endowment to vary across agents and be inheritable by children. However, unlike other agent endowments that are inherited like vision and metabolism, technology should be able to exceed the maximum efficiency that is initially observed in the population. For instance, while the model with current settings constrains maximum vision to 6, we want no such constraint on technology. Intuitively, a later generation may develop a superior technology to what was available to their ancestors. While our model does not specify a mechanism for this, W. Brian Arthur offers one possible explanation in his book *The Nature of Technology*: "novel technologies arise by combination of existing technologies and that (therefore) existing technologies beget further technologies" (Arthur 2009). These three properties (endowment, inheritability and progression beyond initially observed values) are implemented as follows:

Endowment: each initially generated agent receives a technology endowment drawn from the uniform distribution (0, 1). In this iteration, the agents' endowments do not affect their interaction with the environment (that will come in future iterations).

Inheritability: we modify the reproduction rule to create **Rule S3** so that when offspring are produced, they will randomly receive one of their parents' tech endowments (with a probability of receiving something better: see below).

Progression beyond initial values: to allow technology to be strongly (but not completely) bounded by the endowment of the parents, we enable breakthroughs and include a decision tree in the reproduction rule. What follows is a description of this tree in words, as a flowchart, and an example.

Description: For each child born, if one of the parents has 'state-of-the-art' technology, then a lottery is held to determine if a breakthrough might occur when the two

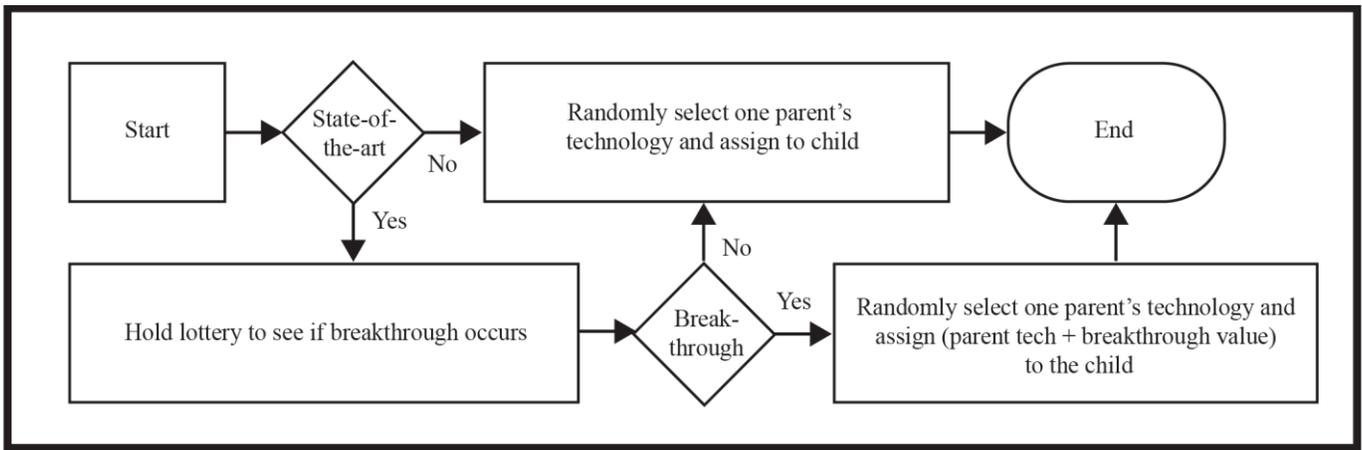


Figure 1: Technology Inheritance with Breakthrough Subject to Probability

parents reproduce. If a breakthrough occurs, the child inherits one of the parents' technology values plus an additional breakthrough value. Intuitively, this represents agents with the best technology being able to (with some randomness involved) advance the technology they pass on to their children by combining the attributes of their technology and their mate's technology. Not all combinations produce an improvement, and not all improvements are made to the better technology. Sometimes, however, a combination is both an improvement and is added to the better technology, representing a true advancement beyond what was available to the child (and perhaps to the whole Sugarscape) before the combination was made.

This process is illustrated in the flowchart shown in figure 1 and following example.

Example: Parent A has technology of .87. Parent B has technology of .52. State-of-the-art is defined as the top 10% of technology (for the purposes of this example, let's assume that the top 10% is calculated by the simulation engine to be anyone with technology above .78), the probability of a breakthrough occurring is 33%, and the value of a breakthrough is .15.

Because Parent A is state-of-the art, a lottery is held. The lottery succeeds and a breakthrough occurs, so the breakthrough value (.15) is added on to one of the parents' technology values (in this case the .52 of Parent B), and the child inherits a technology endowment of .67.

Rules	Parameters
Agent: M1, S3, P1 Patch: G, D	Unless otherwise stated, same as in iteration 3; Initial agents: 400; Agent sustainable technology: randomly generated for initially created agents using a uniform distribution with the range (0,1); "Community Standard" critical value: 8; Breakthrough threshold: .90 (state of the art is top 10 % grid-wide); Breakthrough Value: .15; Breakthrough Probability: .33

Table 4: Summary of Rules for Iteration 4

Results

We experimented with various parameter values until arriving at those described above. Those were chosen because on repeated runs, they resulted in a population with an average technology efficiency value that was stable across generations and was typically observed between .45 and .55 (thereby maintaining the values across generations that our uniformly distributed random assignment of endowments gave us in the initially generated agents - see figure 2). This property was considered desirable because we wanted the technology to be stable (even after introducing the breakthrough concept) before introducing selection pressures. Without breakthrough, natural randomness inherent in the model allowed the average tech observed to vary between .46 and .54. With breakthrough and the parameters above, that same randomness appeared to overwhelm any monotonic increase we might otherwise expect in the average sustainable technology (that is, not enough of the breakthroughs that state-of-the-art parents made propagated far and long enough to meaningfully increase the grid-wide population average value of the technologies held by the agents). This establishes a baseline against which advancements in technology efficiency observed in future iterations can be measured.

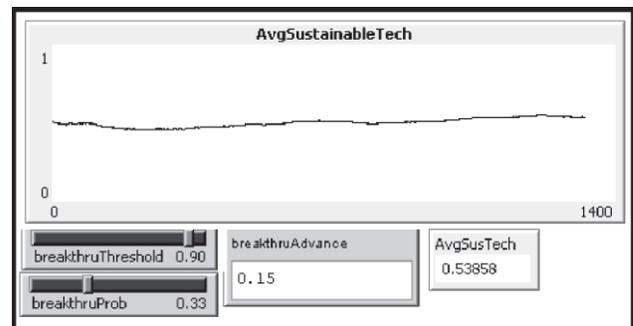


Figure 2: Stable sustainable technology values

Iteration 5: Optional Sustainability

This iteration represents the completion of the base model and will provide the framework for which experiments with various parameters may be conducted.

Description

The basic premise of this iteration is to make the use of sustainable technology optional, allowing the agents to make a series of decisions to determine whether it is in their best interest to harvest sustainably (and not pollute) or harvest regularly (creating pollution). The amount of sugar that an agent receives from harvesting a patch sustainably is determined by the sustainable technology that the particular agent is endowed with. For example, a patch with 2 sugar harvested by an agent endowed with the technology value of .75 would yield 1.5 sugar to the agent (instead of the full 2). We modify the movement rule (**Rule M2**) so that if a patch is harvested sustainably, that action produces no pollution. As with normal harvesting, no sugar would remain on the patch after harvesting. That is, in our example, the .5 sugar the agent does not get the benefit of is presumed to have been waste in the less-polluting, but less-efficient (because the value is <1) sustainable harvest method. While harvesting regularly produces more sugar (getting the agent to the necessary value to reproduce faster), it also creates pollution and incrementally draws the agent closer to the potential of maybe being deemed not fit to reproduce under society's standards (as represented by the 'community standard' critical value beyond which an agent is compared to its peers before being allowed to reproduce). Because the decision to harvest sustainably or traditionally is based in part on a desire to reproduce, it follows that only fertile agents will care to perform the necessary calculus for each harvesting decision and thus they are the only agents that follow this rule. After experimenting with variations, we adopted a model where the youth and the elderly choose to harvest sustainably some percentage of the time based on a pattern that they have observed. For youth the pattern (the probability that they will harvest sustainably) is inherited randomly from one of their parents and represents that parent's lifetime % of harvests that were conducted sustainably up to the point where they reproduced to create that child. For the elderly, the pattern is the lifetime % of harvests that they themselves have conducted sustainably (both as a youth and as a fertile adult). Intuitively, youth will roughly mimic their parents' decisions until they are adults, whereupon they will make their own decisions. Similarly, the elderly will keep acting in the manner that they have chosen to act over the course of their lives.

Rules	Parameters
Agent: M2, S3, P1 Patch: G, D	The same as in iteration 4

Table 5: Summary of Rules for Iteration 5

Decision Rule

The decision rule for agents of fertile age includes two decisions. Recall that in this model, to reproduce, an agent must have at least as much sugar as they started out with (sugar \geq initial sugar). Thus, the initial sugar divided by two represents the halfway point between no sugar (and certain death) and enough sugar to reproduce.

Decision 1: Which is a greater concern to the agent: survival of self or reproductive potential?

IF the agent's sugar is greater than the initial endowment divided by 2

THEN it is closer to reproduction than death and will act on its desire to further its reproductive goals

ELSE the agent will act on its desire to survive and will maximize sugar from the patch by harvesting with the method that yields the most sugar (tie goes to survival). Note that the method that yields the most sugar will create pollution as long as an agent's sustainable technology is < 1. However, after that point, harvesting sustainably will produce as much or more sugar than harvesting regularly, so the agent will choose to harvest with the technology.

Decision 2: If the agent's greatest concern is for reproduction, then determine whether the agent would benefit more from additional wealth or avoiding potential unattractiveness. Then select the method of harvest that makes agent most well off reproductively. This is somewhat like determining the marginal rate of substitution of progress towards sufficient wealth to reproduce (a positive) versus progress towards facing a reproductive comparison (a negative when compared to the sure-thing available to agents who have created less pollution than the 'community standard'). This tradeoff is inherent in the fact that the effect of sustainably harvesting is that an agent gives up an extra accumulation of sugar (the amount will vary with tech endowment) to avoid movement towards potential unattractiveness. Conversely, the effect of harvesting regularly and polluting is that an agent will move closer to unattractiveness to possibly gain additional movement towards reproductive capacity. This implies that when an agent has enough sugar to reproduce, they will always choose to harvest sustainably.

For details regarding this implementation, an appendix is available from the author with equations and pseudo-code, as is the NetLogo model in which the actual implementation occurs and from which the results below are derived.

Experiments

Multiple experiments and a stability check were performed using the model developed in Iteration 5. One experiment was used to identify interesting parameter combinations and consisted of 5 runs of 1000 ticks each for 180 different parameter combinations. The stability check establishes the

long-run stability of the model using 5 runs of 100,000 ticks and validated that all results of interest occurred well before the 1,000 tick limit of the initial runs. A final experiment establishes the robustness of the results using 100 runs of 1000 ticks for 12 parameter combinations, including turning on diffusion which allows us to compare globalized pollution (with diffusion) results versus the localized pollution (no diffusion) results.

The most interesting phase of the development of sustainable technology occurs while that technology is less efficient than traditional technologies. After efficiency parity is achieved, society (and each individual agent) can achieve maximal resource production with no negative externalities caused by pollution by abandoning traditional harvesting methods and using the sustainable technology all the time. Thus, we limit our investigation to what occurs leading up to the realization of efficiency. In our model, this happens when the variable "sustech" equals 1. Phenomena of interest that we are looking for in this time include population counts and crashes, overall pollution and speed to sustech=1.

Results

Optional sustainability allows for more stringent community standards without population crashes

We observe that the population no longer crashes all the time for critical 'community standards' values of less than 7 as it did in iteration 3. Instead, we observe that lower critical values result in lower populations, but that crashes do not occur at critical values 4 and above. For the highest level of social pressure (critical value = 0) the population crashes 88% of the time (68% with diffusion) and the average population when sustech=1 in the runs that didn't crash is 365 (524 with diffusion). For the next highest level of social pressure tested (2), the population crashes 83% of the time (61% with diffusion) with an average population of 540 (790 with diffusion) in the non-crashing runs. There were no population crashes observed for critical values of 4 or above. These results imply that by giving agents the individual choice on whether to pollute, we have enabled the population to better sustain the introduction of pollution-based standards. While overall pollution levels were similar at sustech=1 for localized and globalized pollution, holding everything else equal, localized pollution was found to lead to lower populations and more crashes than globalized pollution. Figure 3 shows histograms of the number of ticks required to reach sustech=1 when diffusion is turned on. Any runs for critical values of 0 or 2 that reached sustech=1 but subsequently crashed are excluded from the figure and the percentages calculated above. For instance, sustech=1 was achieved at a critical value of 0 in 43 of the runs, but 11 of

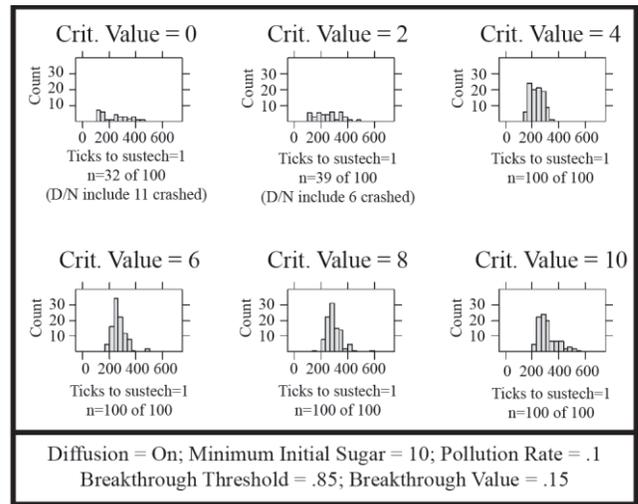


Figure 3: Ticks for sustech=1 across critical values

those subsequently crashed, leaving only 32 of the runs successful in achieving sustech=1 and finishing.

Time to efficient sustainable technologies

We further observe that lower critical 'community standard' values (i.e. stricter social pressures) generally lead to faster achievement of parity between the efficiency of the traditional harvesting technology and the sustainable technology (sustech = 1). This suggests that the more stringent a society's expectations are for its members to not pollute, the faster that desirable attributes enabling more efficient production without causing pollution is selected for. However, these gains must be balanced against the lower populations and much higher risks of population crashes that accompany more stringent social pressures. Interestingly, there is a consistent local maximum in the ticks to sustech=1 at critical value=2 that indicates there may be an additional factor that is important near that value and that for slightly higher critical values (such as 4) a society can achieve sustech=1 just as fast as at a critical value below 2 (see the top left graph in figures 4 and 5 to observe this). We find that these results are true for both localized and global (with diffusion) pollution, with global pollution encouraging 5-11% faster development of efficiency.

See figures 3, 4 and 5 for plots related to these last two results.

Maximized population

For runs that did not crash, we observe a population that begins to cycle like earlier iterations did, but then quickly climbs to a peak (around 1700 agents for the parameters listed in iteration 5) and then cycles (+/- 300 agents in iteration 5) until sustainable technology reaches the same efficiency as traditional harvesting (referred to compactly as "sustech=1"). The population then cycles on an upward

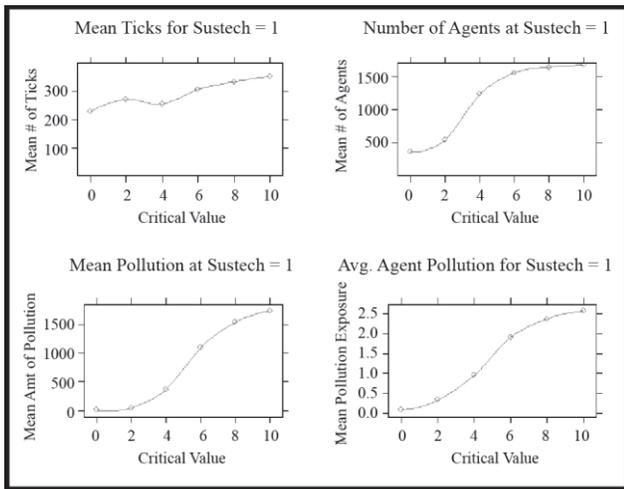


Figure 4: Values of Various Parameters when sustainable tech equaled 1 (Local / No Diffusion)

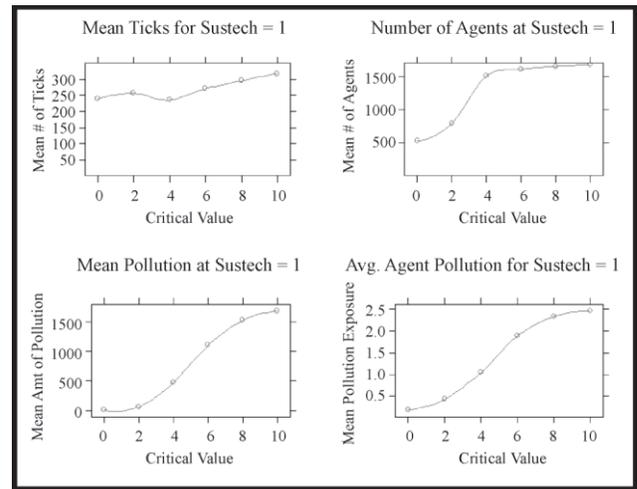


Figure 5: Values of Various Parameters when sustainable tech equaled 1 (Global / With Diffusion)

trend to the maximum number that this Sugarscape can sustain of around 2300 agents. This end pattern holds true for all non-crashing runs.

Higher pollution rates cause population crashes

We also observe that the population crashes almost all the time that pollution rates are increased from 10% to either 30% or 50% of sugar harvested. It is possible that the agents' dislike of pollution drives enough of them to live far enough apart (not clustered in high resource areas) that they do not have enough eligible mates in their vision neighborhoods to reproduce enough to maintain a viable population. This would seem to indicate that for high enough levels of pollution, agents' inherent dislike of pollution overwhelms the conditions needed for reproduction, and likely represents a limitation inherent in the model.

Conclusion

Our model is far too simple to claim to represent any actual society attempting to encourage the development of sustainable technology. Instead, the results of community standards resulting in faster development of sustainable technology at the cost of reduced population during development (and the risk of population crashes in some circumstances) provide a foundation for studying how one type of societal pressure might impact the development of such technology if there were few other forces at work. The most interesting results of this study likely come from future research directions that it suggests. A natural extension of this research would be to establish tribes that have differing critical values representing differing community standards (perhaps one tribe would not even

have a critical value and would always pollute) to see which are more competitive in the long run. The second, and in our opinion, more serious consideration that should be investigated is that if the community standard is so easy to meet (i.e. it is so high that no agent will ever pass it and become potentially unattractive), then it ceases to be realistic to harvest sustainably when an agent has enough wealth to reproduce. This challenge represents a potential for future research based on the modification of the decision rule in iteration 5, perhaps by introducing some amount of foresight as Epstein and Axtell did.

Despite the model's simplicity, however, the results do provide hints of implications for a society that wants to decrease pollution by encouraging the development of sustainable technologies. For such a society to maximize welfare in the sense of achieving the most good (lowest pollution) for the most people, our results imply that to balance population realities and pollution goals, the society may be best served by adopting standards, but making them less stringent. These will allow higher initial pollution, but also support a higher population while developing sustainable technology to efficient levels just as fast (or faster) than more stringent standards. In other words, if a society that wants to eliminate pollution is willing to accept some amount of pollution to begin with, it can encourage technological parity to develop just as quickly without the smaller populations and risk of crashes.

References

Arthur, W. B. 2009. *The nature of technology: What it is and how it evolves*. New York: Free Press.

Epstein, J. M., Axtell, R. 1996. *Growing artificial societies: Social science from the bottom up*. Washington, D.C.; Cambridge, MA: Brookings Institution Press; MIT Press.

Helmreich, S. 1998. *Silicon second nature: Culturing artificial life in a digital world*. Berkeley: University of California Press.

WCED (World Commission on Environment and Development). 1987. *Our common future*. Oxford, UK:

Weaver, I. 2009. Sugarscape NetLogo User Community Model. <http://www2.le.ac.uk/departments/interdisciplinary-science/research/the-sugarscape>. Centre for Interdisciplinary Science, University of Leicester, Leicester, UK. (Retrieved March 9, 2011 from <http://ccl.northwestern.edu/netlogo/models/community/Sugarscape>).

Wilensky, U. 1999. NetLogo. <http://ccl.northwestern.edu/netlogo/>. Center for Connected Learning and Computer-Based Modeling, Northwestern University. Evanston, IL.